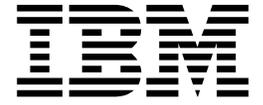


IBM WebSphere MQ



Administering IBM WebSphere MQ

Version 7 Release 1

Note

Before using this information and the product it supports, read the information in "Notices" on page 1709 (*WebSphere MQ V7.1 Installing Guide*).

This edition applies to version 7 release 1 of WebSphere MQ and to all subsequent releases and modifications until otherwise indicated in new editions.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright IBM Corporation 2007, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables xi

Administering IBM WebSphere MQ 1

Local and remote administration.	4
How to use WebSphere MQ control commands.	4
Automating administration tasks	5
Introduction to Programmable Command Formats	6
Using the MQAI to simplify the use of PCFs	18
Introduction to the WebSphere MQ Administration Interface (MQAI)	19
WebSphere MQ Administration Interface (MQAI)	20
Administration using the IBM WebSphere MQ Explorer	59
What you can do with the WebSphere MQ Explorer	60
Setting up the WebSphere MQ Explorer	61
Security on Windows	68
Extending the WebSphere MQ Explorer	71
Using the WebSphere MQ Taskbar application (Windows only)	71
The WebSphere MQ alert monitor application (Windows only)	72
Administering local WebSphere MQ objects	72
Starting and stopping a queue manager	72
Stopping a queue manager manually	74
Performing local administration tasks using MQSC commands	76
Working with queue managers	85
Working with local queues	87
Working with alias queues	93
Working with model queues.	95
Working with administrative topics	96
Working with subscriptions	98
Working with services	102
Managing objects for triggering	108
Administering remote WebSphere MQ objects	110
Channels, clusters, and remote queuing	110
Remote administration from a local queue manager	112
Creating a local definition of a remote queue	118
Using remote queue definitions as aliases	120
Data conversion	121
Administering IBM WebSphere MQ Telemetry	123
Configuring a queue manager for telemetry on Linux and AIX	123
Configuring a queue manager for telemetry on Windows	125
Configure distributed queuing to send messages to MQTT clients	127
MQTT client identification, authorization, and authentication	129
Telemetry channel authentication using SSL	137
Publication privacy on telemetry channels.	139

SSL configuration of MQTT clients and telemetry channels	140
Telemetry channel JAAS configuration	145
WebSphere MQ Telemetry daemon for devices concepts	147
Administering multicast	158
Getting started with multicast	158
WebSphere MQ Multicast topic topology	159
Controlling the size of multicast messages.	160
Enabling data conversion for Multicast messaging	162
Multicast application monitoring	163
Multicast message reliability	164
Advanced multicast tasks	164
Administering HP Integrity NonStop Server	167
Manually starting the TMF/Gateway from Pathway	168
Stopping the TMF/Gateway from Pathway	168
Administering IBM i	168
Managing WebSphere MQ for IBM i using CL commands	169
Alternative ways of administering WebSphere MQ for IBM i	183
Work management	188
Availability, backup, recovery, and restart	195
Quiescing WebSphere MQ for IBM i.	234
Administering IBM WebSphere MQ for z/OS	237
Issuing commands	238
The WebSphere MQ for z/OS utilities	246
Operating WebSphere MQ for z/OS.	248
Writing programs to administer WebSphere MQ	272
Managing WebSphere MQ resources on z/OS	285
Recovery and restart	321
WebSphere MQ and IMS	353

Security 367

Security overview	367
Security concepts and mechanisms	367
IBM WebSphere MQ security mechanisms.	383
Planning for your security requirements	414
Planning identification and authentication.	415
Planning authorization	417
Planning confidentiality	434
Planning data integrity	443
Planning auditing	443
Common Criteria environmental considerations	445
Planning security by topology.	447
Firewalls and Internet pass-thru	461
IBM WebSphere MQ for z/OS security implementation checklist	462
Setting up security	464
Setting up security on Windows, UNIX and Linux systems	464
Setting up security on HP Integrity NonStop Server	491

Setting up security on IBM i	492
Setting up security on z/OS	520
Setting up WebSphere MQ MQI client security	613
Working with SSL or TLS	616
Configuring IBM WebSphere MQ for Common Criteria	685
Identifying and authenticating users.	689
Privileged users	692
Identifying and authenticating users using the MQCSP structure	692
Implementing identification and authentication in security exits	692
Identity mapping in message exits	693
Identity mapping in the API exit and API-crossing exit	694
Working with revoked certificates	695
Authorizing access to objects	704
Controlling access to objects by using the OAM on UNIX, Linux and Windows systems.	704
Granting required access to resources	712
Authority to administer WebSphere MQ on UNIX, Linux and Windows systems.	751
Authority to work with WebSphere MQ objects on UNIX, Linux and Windows systems.	753
Implementing access control in security exits	758
Implementing access control in message exits	760
Implementing access control in the API exit and API-crossing exit	760
Confidentiality of messages	761
Connecting two queue managers using SSL or TLS.	761
Connecting a client to a queue manager securely	768
Specifying CipherSpecs	774
Resetting SSL and TLS secret keys	778
Implementing confidentiality in user exit programs.	780
Data integrity of messages	781
Connecting two queue managers using SSL or TLS.	782
Connecting a client to a queue manager securely	791
Specifying CipherSpecs	797
Auditing	801
Keeping clusters secure	801
Stopping unauthorized queue managers sending messages	801
Stopping unauthorized queue managers putting messages on your queues	802
Authorizing putting messages on remote cluster queues	802
Preventing queue managers joining a cluster	803
Forcing unwanted queue managers to leave a cluster.	805
Preventing queue managers receiving messages SSL and clusters	806
Publish/subscribe security	808
Example publish/subscribe security setup.	815
Subscription security	826
Monitoring and performance	829
Event monitoring	829
Instrumentation events	830

Performance events	845
Configuration events	863
Command events	867
Logger events	869
Sample program to monitor instrumentation events.	877
Message monitoring	885
Activities and operations	886
Message route techniques	888
Activity recording.	889
Trace-route messaging	894
WebSphere MQ display route application	908
Activity report reference.	926
Trace-route message reference	952
Trace-route reply message reference	963
Accounting and statistics messages	966
Accounting messages.	967
Statistics messages.	971
Displaying accounting and statistics information	976
Accounting and statistics message reference	981
Application activity trace	1030
Collecting application activity trace information	1030
Application activity trace message reference	1037
Real-time monitoring	1102
Attributes that control real-time monitoring	1103
Displaying queue and channel monitoring data	1104
Monitoring queues	1106
Monitoring channels.	1108
The Windows performance monitor	1114
Monitoring performance and resource usage.	1114
Introduction to monitoring	1114
Interpreting WebSphere MQ performance statistics.	1122
Interpreting WebSphere MQ accounting data	1137
Troubleshooting and support	1149
Troubleshooting overview	1149
Making initial checks on Windows, UNIX and Linux systems.	1151
Has IBM WebSphere MQ run successfully before?	1152
Have any changes been made since the last successful run?	1153
Are there any error messages or return codes to explain the problem?	1153
Can you reproduce the problem?	1154
Are you receiving an error code when creating or starting a queue manager? (Windows only)	1154
Does the problem affect only remote queues?	1154
Have you obtained incorrect output?	1155
Are some of your queues failing?	1157
Have you failed to receive a response from a PCF command?	1157
Has the application run successfully before?	1158
Is your application or system running slowly?	1160
Does the problem affect specific parts of the network?	1160
Does the problem occur at specific times of the day?	1160
Is the problem intermittent?	1160

Making initial checks on IBM i	1161	Resolving problem: Lost messages in an MQTT application	1215
Problem characteristics	1163	Resolving problem: Telemetry service does not start	1216
Required authority	1166	Resolving problem: JAAS login module not called by the telemetry service	1218
Determining problems with IBM WebSphere MQ for IBM i applications.	1167	Resolving problem: Starting or running the daemon	1221
Making initial checks on z/OS	1169	Resolving problem: MQTT clients not connecting to the daemon	1221
Has IBM WebSphere MQ for z/OS run successfully before?	1171	Troubleshooting channel authentication records	1222
Have you applied any APARs or PTFs?	1172	Multicast troubleshooting	1223
Are there any error messages, return codes or other error conditions?	1172	Testing multicast applications on a non-multicast network	1223
Has your application or WebSphere MQ for z/OS stopped processing work?	1174	Setting the appropriate network for multicast traffic	1223
Is there a problem with the WebSphere MQ queues?	1175	Multicast topic string is too long	1224
Does the problem affect specific parts of the network?	1178	Multicast topic topology issues	1224
Problems that occur at specific times of the day or affect specific users	1178	Using logs	1226
Is the problem intermittent or does the problem occur with all z/OS, CICS, or IMS systems?	1178	Error logs on Windows, UNIX and Linux systems	1227
Has the application run successfully before?	1179	Error logs on HP Integrity NonStop Server	1230
Have any changes been made since the last successful run?	1180	Error logs on IBM i	1231
Do you have a program error?	1181	Using trace.	1234
Has there been an abend?	1182	Using trace on Windows	1234
Have you obtained incorrect output?	1183	Using trace on UNIX and Linux systems	1236
Can you reproduce the problem?	1183	Using trace on HP Integrity NonStop Server	1239
Have you failed to receive a response from an MQSC command?	1184	Using trace on WebSphere MQ server on IBM i	1240
Is your application or IBM WebSphere MQ for z/OS running slowly?	1185	Using trace on WebSphere MQ client on IBM i	1242
Dealing with problems	1186	Using trace for problem determination on z/OS.	1244
Resolving problems with commands	1186	Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions.	1255
Resolving problems with queue managers	1187	Tracing WebSphere MQ classes for JMS programs	1256
Resolving problems with queue manager clusters	1187	Tracing WebSphere MQ classes for Java programs	1260
Resolving problems with undelivered messages	1199	Java diagnostics	1262
Resolving problems with IBM WebSphere MQ MQI clients.	1200	Problem determination on z/OS.	1265
Troubleshooting IBM WebSphere MQ client for HP Integrity NonStop Server.	1201	IBM WebSphere MQ for z/OS performance constraints	1266
Java and JMS troubleshooting	1201	WebSphere MQ for z/OS recovery actions	1268
Troubleshooting JMSSC0108 messages	1201	WebSphere MQ for z/OS abends	1269
Problem determination for the WebSphere MQ resource adapter	1204	Diagnostic information produced on IBM WebSphere MQ for z/OS	1272
Troubleshooting for IBM WebSphere MQ Telemetry	1206	Other sources of information	1274
Location of telemetry logs, error logs, and configuration files	1206	Diagnostic aids for CICS	1275
MQTT v3 Java client reason codes	1208	Diagnostic aids for IMS	1275
Tracing the telemetry service	1209	Diagnostic aids for Db2	1276
Tracing the MQTT v3 Java client	1210	WebSphere MQ dumps.	1276
System requirements for using SHA-2 cipher suites with MQTT channels and clients	1211	Dealing with performance problems on z/OS	1293
Resolving problem: MQTT client does not connect	1212	Dealing with incorrect output	1299
Resolving problem: MQTT client connection dropped.	1214	Resolving problems with channels and DQM	1305
		Error message from channel control	1306
		Using Ping to check the communication link	1307
		Channel problems and the dead-letter queue	1307
		Validation checks.	1308
		In-doubt relationship	1308
		Channel startup negotiation errors	1308
		Shared channel recovery	1308
		When a channel does not run	1308

Retrying the link	1311
Data structures	1312
User exit problems	1312
Disaster recovery	1312
Channel switching	1313
Connection switching	1313
Client problems	1313
Error logs	1314
Message monitoring	1315
First Failure Support Technology (FFST)	1315
FFST: WebSphere MQ for Windows	1315
FFST: WebSphere MQ for UNIX and Linux systems	1318
First Failure Support Technology (FFST)	1320
FFST: WebSphere MQ for HP Integrity NonStop Server	1322
IBM Support Assistant (ISA)	1324
Installing the IBM Support Assistant (ISA)	1324
Updating the IBM Support Assistant (ISA)	1325
Searching knowledge bases	1326
Searching the IBM database for similar problems, and solutions	1327
Contacting IBM Software Support	1346
Determine the effect of the problem on your business	1347
Describe your problem and gather background information	1347

Submit your problem to IBM Software Support	1347
Dealing with the support center	1348
Collecting documentation for the problem	1352
Sending the documentation to the change team	1353
Resolving a problem	1354
Getting product fixes	1355
Recovering after failure	1355
Disk drive failures	1356
Damaged queue manager object	1357
Damaged single object	1358
Automatic media recovery failure	1358
Example recovery scenarios	1358
Reason codes	1379
API completion and reason codes	1380
PCF reason codes	1590
Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes	1672
WCF custom channel exceptions	1676

Index 1685

Notices 1709

Programming interface information	1710
Trademarks	1711

Sending your comments to IBM 1713

Figures

1. Hierarchy of MQAI concepts	49	40. !Remap the same topics for publications flowing in both directions	152
2. Adding data items	52	41. !Remap the same topics for publications flowing in both directions, using both.	152
3. Modifying a single data item.	54	42. Extract from the MQSC command file, myprog.in	185
4. Modifying all data items	55	43. Display MQM Command Server panel	188
5. Truncating a bag	56	44. Sequence of events when updating MQM objects	198
6. Converting bags to PCF messages	56	45. WebSphere MQ for IBM i journaling.	202
7. Converting PCF messages to bag form	56	46. Mirror a queue manager journal	212
8. Deleting a single data item	57	47. Mirrored journal configuration.	219
9. Deleting all data items	57	48. Mirrored journal configuration.	223
10. Nesting	58	49. Transfer a queue manager journal using an independent ASP	224
11. Extract from an MQSC command file	81	50. Transfer a queue manager journal using an independent ASP	227
12. Extract from an MQSC command report file	82	51. Issuing a DISPLAY command from the z/OS console	249
13. Example script for running MQSC commands from a batch file	84	52. Starting the queue manager from a z/OS console	250
14. Typical output from a DISPLAY QMGR command	86	53. Sample start-up procedure	251
15. Typical results from queue browser	92	54. Stopping a queue manager	252
16. Remote administration using MQSC commands	113	55. The WebSphere MQ operations and control initial panel	268
17. Setting up channels and queues for remote administration	114	56. Extract from a load balancing job for a page set	304
18. installMQXRService_unix.mqsc	125	57. Sample job for moving a queue from one CF structure to another	318
19. Set default transmission queue.	126	58. Sample job for moving a non-shared queue to a shared queue	318
20. installMQXRService_win.mqsc	126	59. Sample job for moving a shared queue to a non-shared queue	319
21. Defining a cluster topic on Windows	128	60. Sample job for moving a non-shared queue without messages to a shared queue.	319
22. MQI Object descriptor to send a message to an MQTT v3 client destination.	128	61. Moving messages from a non-shared queue to an existing shared queue.	320
23. JMS destination to send a message to an MQTT v3 client	129	62. Starting the queue manager from a z/OS console	327
24. MQTT Client code snippet	144	63. Sample start-up procedure	328
25. AcceptAllProvider.java	144	64. Stopping a queue manager	329
26. AcceptAllTrustManagerFactory.java	144	65. Sample input statements for CSQJU003	333
27. AcceptAllX509TrustManager.java	145	66. Point in time for recovery for 2 queue managers in a queue-sharing group	335
28. Sample jaas.config file	146	67. Sample ARM policy	340
29. Sample JAASLoginModule.Login() method	146	68. Example restart messages	346
30. Connecting WebSphere MQ Telemetry daemon for devices to WebSphere MQ	148	69. Symmetric key cryptography	370
31. Publish everything to the remote broker	149	70. Asymmetric key cryptography.	370
32. Publish everything to the remote broker - explicit.	150	71. Obtaining a digital certificate	374
33. Publish everything to the local broker	150	72. Chain of trust	375
34. Publish everything from the export topic at the local broker to the import topic at the remote broker	150	73. Overview of the SSL or TLS handshake	378
35. Publish everything to the import topic at the local broker from the export topic at the remote broker	150	74. The digital signature process	381
36. Publish everything from the 1884/ mount point to the remote broker with the original topic strings.	151	75. Security in a client/server connection	416
37. Separate the topic spaces of clients connected to different daemons	151	76. Link level security and application level security	435
38. Separate the topic spaces of clients connected to the same daemon	151		
39. Remap different topics for publications flowing in both directions	152		

77. Security, message, send, and receive exits on a message channel	441	119. Requesting activity reports, Diagram 3	918
78. Flows for session level authentication	455	120. Requesting activity reports, Diagram 4	919
79. WebSphere MQ support for conversation level authentication	457	121. Requesting a trace-route reply message, Diagram 1	920
80. Checking for subsystem security	525	122. Requesting a trace-route reply message, Diagram 2	921
81. Checking for queue manager level security	526	123. Requesting a trace-route reply message, Diagram 3	921
82. Checking for queue-sharing group level security	526	124. Requesting a trace-route reply message, Diagram 4	922
83. Logic flow for WebSphere MQ security caching	579	125. Delivering activity reports to the system queue, Diagram 1	923
84. Typical output from the DISPLAY SECURITY command	581	126. Delivering activity reports to the system queue, Diagram 2	923
85. Sample output from RACFRW showing RESLEVEL general audit records	587	127. Diagnosing a channel problem	926
86. Example security scenario	601	128. Part of an SMF record 115 showing the header and self-defining sections.	1125
87. Sample LDIF file for a Certificate Authority. This might vary from implementation to implementation.	698	129. SMF record 115, subtype 1	1126
88. Example of an LDAP Directory Information Tree structure	698	130. SMF record 115, subtype 2	1127
89. Configuration resulting from this task	762	131. Part of an SMF record 116 showing the header and self-defining sections.	1140
90. Configuration resulting from this task	764	132. Part of an SMF record 116 showing the header and self-defining sections.	1143
91. Queue managers allowing one-way authentication	767	133. Example SMF type 116, subtype zero record	1146
92. Configuration resulting from this task	769	134. Example SMF type 116, subtype 1 record	1147
93. Configuration resulting from this task	771	135. Example SMF type 116, subtype 2 record	1148
94. Client and queue manager allowing anonymous connection	773	136. Example SMF type 116, subtype 1 record with no WQST data records	1148
95. Configuration resulting from this task	785	137. Telemetry configuration directory on Windows	1207
96. Configuration resulting from this task	787	138. Telemetry configuration directory on AIX or Linux.	1207
97. Queue managers allowing one-way authentication	790	139. Sample service.env for Windows	1208
98. Configuration resulting from this task	792	140. <i>WMQ Installation directory\data\mqmgrs\qMgrName\mqxr\mqxr_win.properties</i>	1219
99. Configuration resulting from this task	794	141. <i>WMQ Installation directory\data\mqmgrs\qMgrName\mqxr\jaas.config</i>	1219
100. Client and queue manager allowing anonymous connection	796	142. <i>WMQ Installation directory\data\mqmgrs\qMgrName\service.env</i>	1219
101. Publish/subscribe security relationships	808	143. Classpath output from runMQXRService.bat	1219
102. Topic object access example.	815	144. <i>WMQ Installation directory\data\mqmgrs\qMgrName\errors\mqxr.log</i>	1219
103. Example of granting access to a topic within a topic tree	817	145. Exception thrown connecting com.ibm.mq.id.PubAsyncRestartable	1220
104. Granting access to specific topics within a topic tree	818	146. mqxr.log - error loading JAAS configuration	1220
105. Example of granting access control to avoid additional messages.	820	147. Sample WebSphere MQ error log	1228
106. Granting publish access to a topic.	821	148. Sample WebSphere MQ error log	1231
107. Granting publish access to a topic within a topic tree	823	149. Extract from a IBM WebSphere MQ error log	1233
108. Granting access for publishing and subscribing	824	150. Example startup of GTF to use with the WebSphere MQ trace	1246
109. Understanding instrumentation events	831	151. Example of GTF Stop command to use with the WebSphere MQ trace	1246
110. Monitoring queue managers across different platforms, on a single node	833	152. Formatting the GTF output in batch	1248
111. Understanding queue service interval events	847	153. Sample com.ibm.mq.commonservices properties file	1264
112. Queue service interval events - example 1	851	154. Sample symptom string.	1273
113. Queue service interval events - example 2	853	155. Dumping the WebSphere MQ queue manager and application address spaces	1278
114. Queue service interval events - example 3	854	156. Dumping the WebSphere MQ queue manager address space	1278
115. Queue depth events (1)	860	157. Dumping the channel initiator address space	1278
116. Queue depth events (2)	862		
117. Requesting activity reports, Diagram 1	916		
118. Requesting activity reports, Diagram 2	917		

158. Dumping the WebSphere MQ queue manager and channel initiator address spaces	1278	164. Sample beginning of a SYSUDUMP	1291
159. Dumping a coupling facility structure	1279	165. Sample WebSphere MQ for Windows First Failure Symptom Report	1317
160. IPCS MVS Dump Component Data Analysis panel	1279	166. FFST report for WebSphere MQ for UNIX systems	1319
161. Sample JCL for printing dumps through IPCS in the z/OS environment	1288	167. Sample FFST data.	1323
162. Sample SVC dump title.	1289	168. High-level flowchart of various sets of keywords	1332
163. Dump title with PSW and ASID	1290	169. Sample problem reporting sheet	1349

Tables

1. Windows, HP OpenVMS Alpha, HP Integrity NonStop Server, IBM i, UNIX and Linux systems - object authorities	16
2. Name resolution of an MQTT queue manager alias	127
3. Name resolution of an MQTT client remote queue definition	129
4. WebSphere MQ tasks..	189
5. Work management objects	190
6. Work management objects created for a queue manager	190
7. Remote journaling options	213
8. Summary of the main MQSC and PCF commands by object type	241
9. Sources from which to run MQSC commands	242
10. Valid operations and control panel actions for WebSphere MQ objects	264
11. Archiving and logging parameters that can be changed while a queue manager is running	288
12. Differences between FIPS mode and non-FIPS mode algorithm support..	397
13. Suite B security levels with allowed CipherSpecs and digital signature algorithms .	399
14. Relationships between CipherSpecs and digital certificates	405
15. Supported NIST elliptic curves.	408
16. PCF commands and their equivalent OAM commands	429
17. The user ID used by a server-connection channel	433
18. Security authorization needed for MQCONN calls.	473
19. Security authorization needed for MQOPEN calls.	473
20. Security authorization needed for MQPUT1 calls.	473
21. Security authorization needed for MQCLOSE calls.	474
22. Authorizations for MQI calls	500
23. Authorizations for context calls	500
24. Authorizations for MQSC and PCF calls	500
25. Authorizations for generic operations	500
26. Security authorization needed for MQCONN calls.	503
27. Security authorization needed for MQOPEN calls.	503
28. Security authorization needed for MQPUT1 calls.	504
29. Security authorization needed for MQCLOSE calls.	504
30. RACF classes used by WebSphere MQ	521
31. Switch profiles for subsystem level security	525
32. Switch profiles for queue-sharing group or queue manager level security	526
33. Valid security switch combinations for queue manager level security	527
34. Valid security switch combinations for queue-sharing group level security	527
35. Valid security switch combinations for queue manager and queue-sharing group level security	527
36. Other valid security switch combinations that switch both levels of checking on	528
37. Switch profiles for resource checking	528
38. Access levels for queue security using the MQOPEN or MQPUT1 calls	533
39. Access levels for queue security using the MQSUB call	535
40. Access levels for close options on permanent dynamic queues	537
41. RACF authority to the dead-letter queue and its alias	539
42. Access required to the SYSTEM queues by WebSphere MQ	539
43. MQOPEN, MQPUT1, MQSUB, and MQCLOSE options and the security authorization required	540
44. Access level required for topic security to subscribe	543
45. Access level required to profiles for topic security for closure of a subscribe operation	544
46. Access level required to profiles for topic security for a publish operation	544
47. Access required to the SYSTEM topics	545
48. Access levels for process security	545
49. Access levels for namelist security	546
50. Access required to the SYSTEM namelists by WebSphere MQ	547
51. Access levels for alternate user security	548
52. Access levels for context security	549
53. Access required to the SYSTEM queues for context operations	551
54. MQSC commands, profiles, and their access levels	552
55. PCF commands, profiles, and their access levels	556
56. Checks made at different RACF access levels for batch connections	563
57. Checks made at different RACF access levels for CICS connections	564
58. Checks made at different RACF access levels for IMS connections	565
59. Checks made at different RACF access levels for channel initiator connections	566
60. Checks made at different RACF access levels for the intra-group queuing agent.	567
61. User ID checking against profile name for batch connections	569
62. User ID checking against profile name for CICS-type user IDs	569
63. User ID checking against profile name for IMS-type user IDs	570

64. How the second user ID is determined for the IMS connection	571	101. Event statistics summary for queue depth events (example 1).	861
65. User IDs checked against profile name for TCP/IP channels	571	102. Summary showing which events are enabled	861
66. User IDs checked against profile name for LU 6.2 channels	572	103. Event statistics summary for queue depth events (example 2).	863
67. User IDs checked against profile name for LU 6.2 and TCP/IP server-connection channels	574	104. Summary showing which events are enabled	863
68. User IDs checked against profile name for intra-group queuing	576	105. TraceRoute PCF group	900
69. RACF access to data sets associated with a queue manager	583	106. Activity report format.	928
70. RACF access to data sets associated with distributed queuing	583	107. Trace-route message format	954
71. CICS bridge monitor security	595	108. Trace-route reply message format	965
72. CICS bridge task security	596	109. Detail level of channel statistics information collection	974
73. Security profiles for the example scenario	604	110. MQI accounting message structure	983
74. Sample security profiles for the batch application on queue manager QM1	605	111. Array indexed by object type	1029
75. Sample security profiles for queue manager QM2 using TCP/IP and not SSL	606	112. Array indexed by persistence value	1030
76. Sample security profiles for queue manager QM2 using LU 6.2	606	113. Monitoring levels	1103
77. Sample security profiles for queue manager QM2 using TCP/IP and SSL	607	114. Substates seen with status binding or requesting	1110
78. Sample security profiles for the CICS application on queue manager QM1	607	115. Sender and receiver MCA substates	1112
79. Security profiles for the example scenario	610	116. Setting and displaying attributes to control online monitoring.	1118
80. Sample security profiles for the batch application on queue manager QM1	611	117. SMF record 115 header description	1123
81. Sample security profiles for queue manager QM2 using TCP/IP	612	118. Offsets to self-defining sections	1124
82. Sample security profiles for queue manager QM2 using LU 6.2	612	119. Db2 statistics record (Q5ST)	1132
83. Privileged users by platform	692	120. Coupling facility statistics record (QEST)	1136
84. Example topic object authorities	811	121. Topic manager statistics record (QTST)	1137
85. User IDs used for security checks for commands	815	122. SMF record header description	1138
86. Example topic object access	815	123. Offsets to self-defining sections	1139
87. Access requirements for example topics and topic objects	817	124. SMF record header description	1141
88. Access requirements for example topics and topic objects	818	125. Offsets to self-defining sections	1141
89. Example publish access requirements	822	126. Structure of the common WebSphere MQ SMF header record QWHS.	1144
90. Example publish access requirements	823	127. Structure of the thread cross reference for a CICS system	1145
91. Example publishing and subscribing access requirements.	824	128. Structure of the thread cross reference for an IMS system	1145
92. Complete list of access authorities resulting from security examples	825	129. Meaning of channel names.	1147
93. Default publication context information	826	130. MQTT v3 Java client reason codes	1209
94. Enabling queue manager events using MQSC commands	841	131. Symptom table.	1221
95. Enabling channel and bridge events using MQSC commands	842	132. Queue manager error log directory	1229
96. Performance event statistics.	846	133. System error log directory	1229
97. Enabling queue service interval events using MQSC	849	134. Client error log directory	1230
98. Event statistics summary for example 1	852	135. <i>TraceType</i> values	1243
99. Event statistics summary for example 2	854	136. Control blocks traced for WebSphere MQ MQI calls	1249
100. Event statistics summary for example 3	855	137. CICS adapter trace entries	1252
		138. Abend completion codes	1269
		139. Types of dump used with WebSphere MQ for z/OS.	1277
		140. Keywords for the WebSphere MQ for z/OS dump formatting control statement	1284
		141. Resource manager dump formatting keywords	1284
		142. Summary dump keywords for the WebSphere MQ for z/OS dump formatting control statement	1286
		143. IPCS subcommands used for dump analysis	1286
		144. Keywords for the IPCS VERBEXIT CSQXDPRD	1287
		145. Keywords defined for use in a free format search.	1329

146. Keywords defined for use in a structured format search	1330	151. SDB format symptom-to-keyword cross-reference	1343
147. Types of WebSphere MQ for z/OS failures	1334	152. WebSphere MQ component and resource manager identifiers	1345
148. Message prefixes	1339	153. Example recovery scenarios	1358
149. INCORROUT modifier keywords: RECOVERY.	1342	154. SSL return codes	1672
150. INCORROUT modifier keywords: UTILITY	1342	155. Certificate validation errors	1674

Administering IBM WebSphere MQ

Administering queue managers and associated resources includes the tasks that you perform frequently to activate and manage those resources. Choose the method you prefer to administer your queue managers and associated resources.

You can administer IBM® WebSphere® MQ objects locally or remotely, see “Local and remote administration” on page 4.

There are a number of different methods that you can use to create and administer your queue managers and their related resources in IBM WebSphere MQ. These methods include command-line interfaces, a graphical user interface, and an administration API. See the sections and links in this topic for more information about each of these interfaces.

There are different sets of commands that you can use to administer IBM WebSphere MQ depending on your platform:

- “IBM WebSphere MQ control commands”
- “IBM WebSphere MQ Script (MQSC) commands”
- “Programmable Command Formats (PCFs)” on page 2
- “IBM i Control Language (CL)” on page 2

There are also the other following options for creating and managing IBM WebSphere MQ objects:

- “The IBM WebSphere MQ Explorer” on page 3
- “The Windows Default Configuration application” on page 3
- “The Microsoft Cluster Service (MSCS)” on page 3

For information about the administration interfaces and options on IBM WebSphere MQ for z/OS®, see “Administering IBM WebSphere MQ for z/OS” on page 237.

You can automate some administration and monitoring tasks for both local and remote queue managers by using PCF commands. These commands can also be simplified through the use of the WebSphere MQ Administration Interface (MQAI) on some platforms. For more information about automating administration tasks, see “Automating administration tasks” on page 5.

IBM WebSphere MQ control commands

Control commands allow you to perform administrative tasks on queue managers themselves.

IBM WebSphere MQ for Windows, UNIX and Linux systems provides the *control commands* that you issue at the system command line.

The control commands are described in  Creating and managing queue managers (*WebSphere MQ V7.1 Installing Guide*). For the command reference for the control commands, see  WebSphere MQ Control commands (*WebSphere MQ V7.1 Reference*).

IBM WebSphere MQ Script (MQSC) commands

Use MQSC commands to manage queue manager objects, including the queue manager itself, queues, process definitions, namelists, channels, client connection channels, listeners, services, and authentication information objects.

You issue MQSC commands to a queue manager by using the **runmqsc** command. You can do this interactively, issuing commands from a keyboard, or you can redirect the standard input device (stdin) to run a sequence of commands from an ASCII text file. In both cases, the format of the commands is the same.

You can run the **runmqsc** command in three modes, depending on the flags set on the command:

- *Verification mode*, where the MQSC commands are verified on a local queue manager, but are not run
- *Direct mode*, where the MQSC commands are run on a local queue manager
- *Indirect mode*, where the MQSC commands are run on a remote queue manager

Object attributes specified in MQSC commands are shown in this section in uppercase (for example, RQMNAME), although they are not case-sensitive. MQSC command attribute names are limited to eight characters.

MQSC commands are available on all platforms, including IBM i, and z/OS. MQSC commands are summarized in  Comparing command sets (*WebSphere MQ V7.1 Reference*).

On Windows, UNIX or Linux, you can use the MQSC as single commands issued at the system command line. To issue more complicated, or multiple commands, the MQSC can be built into a file that you run from the Windows, UNIX or Linux system command line. MQSC can be sent to a remote queue manager. For full details, see  MQSC reference (*WebSphere MQ V7.1 Reference*).

On IBM i, to issue the commands on an IBM i server, create a list of commands in a Script file, and then run the file by using the STRMQMMQSC command.

Note: On IBM i, MQSC responses to commands issued from a script file are returned in a spool file.

“Script (MQSC) Commands” on page 77 contains a description of each MQSC command and its syntax.

See “Performing local administration tasks using MQSC commands” on page 76 for more information about using MQSC commands in local administration.

Programmable Command Formats (PCFs)

Programmable Command Formats (PCFs) define command and reply messages that can be exchanged between a program and any queue manager (that supports PCFs) in a network. You can use PCF commands in a systems management application program for administration of IBM WebSphere MQ objects: authentication information objects, channels, channel listeners, namelists, process definitions, queue managers, queues, services, and storage classes. The application can operate from a single point in the network to communicate command and reply information with any queue manager, local, or remote, using the local queue manager.

For more information about PCFs, see “Introduction to Programmable Command Formats” on page 6.

For definition of PCFs and structures for the commands and responses, see  Programmable command formats reference (*WebSphere MQ V7.1 Reference*).

IBM i Control Language (CL)

This language can be used to issue administration commands to IBM WebSphere MQ for IBM i. The commands can be issued either at the command line or by writing a CL program. These commands perform similar functions to PCF commands, but the format is different. CL commands are designed

exclusively for servers and CL responses are designed to be human-readable, whereas PCF commands are platform independent and both command and response formats are intended for program use.

For full details of the IBM i Control Language (CL), see  WebSphere MQ for IBM i CL commands (*WebSphere MQ V7.1 Reference*).

The IBM WebSphere MQ Explorer

Using the IBM WebSphere MQ Explorer, you can perform the following actions:

- Define and control various resources including queue managers, queues, process definitions, namelists, channels, client connection channels, listeners, services, and clusters.
- Start or stop a local queue manager and its associated processes.
- View queue managers and their associated objects on your workstation or from other workstations.
- Check the status of queue managers, clusters, and channels.
- Check to see which applications, users, or channels have a particular queue open, from the queue status.

On Windows and Linux systems, you can start IBM WebSphere MQ Explorer by using the system menu, the MQExplorer executable file, or the **strmqcfig** command.

On Linux, to start the IBM WebSphere MQ Explorer successfully, you must be able to write a file to your home directory, and the home directory must exist.

You can use IBM WebSphere MQ Explorer to administer remote queue managers on other platforms including z/OS, for details and to download the SupportPac MSOT, see  <http://www.ibm.com/support/docview.wss?uid=swg24021041>.

See “Administration using the IBM WebSphere MQ Explorer” on page 59 for more information.

The Windows Default Configuration application

You can use the Windows Default Configuration program to create a *starter* (or default) set of IBM

WebSphere MQ objects. A summary of the default objects created is listed in  Objects created by the Windows default configuration application.

The Microsoft Cluster Service (MSCS)

Microsoft Cluster Service (MSCS) enables you to connect servers into a *cluster*, giving higher availability of data and applications, and making it easier to manage the system. MSCS can automatically detect and recover from server or application failures.

It is important not to confuse clusters in the MSCS sense with IBM WebSphere MQ clusters. The distinction is:

IBM WebSphere MQ clusters

are groups of two or more queue managers on one or more computers, providing automatic interconnection, and allowing queues to be shared among them for load balancing and redundancy.

MSCS clusters

Groups of computers, connected together and configured in such a way that, if one fails, MSCS performs a *failover*, transferring the state data of applications from the failing computer to another computer in the cluster and reinitiating their operation there.

 Supporting the Microsoft Cluster Service (MSCS) (*WebSphere MQ V7.1 Installing Guide*) provides detailed information about how to configure your IBM WebSphere MQ for Windows system to use MSCS.

Related concepts:

 WebSphere MQ technical overview (*WebSphere MQ V7.1 Product Overview Guide*)

“Administering local WebSphere MQ objects” on page 72

“Administering remote WebSphere MQ objects” on page 110

“Administering IBM i” on page 168

“Administering IBM WebSphere MQ for z/OS” on page 237

 Planning (*WebSphere MQ V7.1 Installing Guide*)

 Planning your IBM WebSphere MQ environment on z/OS (*WebSphere MQ V7.1 Installing Guide*)

 Configuring (*WebSphere MQ V7.1 Installing Guide*)

 Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*)

Local and remote administration

You can administer WebSphere MQ objects locally or remotely.

Local administration means carrying out administration tasks on any queue managers you have defined on your local system. You can access other systems, for example through the TCP/IP terminal emulation program **telnet**, and carry out administration there. In WebSphere MQ, you can consider this as local administration because no channels are involved, that is, the communication is managed by the operating system.

WebSphere MQ supports administration from a single point of contact through what is known as *remote administration*. This allows you to issue commands from your local system that are processed on another system. For example, you can issue a remote command to change a queue definition on a remote queue manager. You do not have to log on to that system, although you do need to have the appropriate channels defined. The queue manager and command server on the target system must be running.

Some commands cannot be issued in this way, in particular, creating or starting queue managers and starting command servers. To perform this type of task, you must either log onto the remote system and issue the commands from there or create a process that can issue the commands for you.

“Administering remote WebSphere MQ objects” on page 110 describes the subject of remote administration in greater detail.

How to use WebSphere MQ control commands

This section describes how to use the WebSphere MQ control commands.

If you want to issue control commands, your user ID must be a member of the mqm group. For more information about this, see “Authority to administer WebSphere MQ on UNIX, Linux and Windows systems” on page 751. In addition, note the following environment-specific information:

WebSphere MQ for Windows

All control commands can be issued from a command line. Command names and their flags are not case sensitive: you can enter them in uppercase, lowercase, or a combination of uppercase and lowercase. However, arguments to control commands (such as queue names) are case sensitive.

In the syntax descriptions, the hyphen (-) is used as a flag indicator. You can use the forward slash (/) instead of the hyphen.

WebSphere MQ for UNIX and Linux systems

All WebSphere MQ control commands can be issued from a shell. All commands are case-sensitive.

A subset of the control commands can be issued using the WebSphere MQ Explorer.

For more information, see  WebSphere MQ Control commands (*WebSphere MQ V7.1 Reference*)

Automating administration tasks

You might decide that it would be beneficial to your installation to automate some administration and monitoring tasks. You can automate administration tasks for both local and remote queue managers using programmable command format (PCF) commands. This section assumes that you have experience of administering WebSphere MQ objects.

PCF commands

WebSphere MQ programmable command format (PCF) commands can be used to program administration tasks into an administration program. In this way, from a program you can manipulate queue manager objects (queues, process definitions, namelists, channels, client connection channels, listeners, services, and authentication information objects), and even manipulate the queue managers themselves.

PCF commands cover the same range of functions provided by MQSC commands. You can write a program to issue PCF commands to any queue manager in the network from a single node. In this way, you can both centralize and automate administration tasks.

Each PCF command is a data structure that is embedded in the application data part of a WebSphere MQ message. Each command is sent to the target queue manager using the MQI function **MQPUT** in the same way as any other message. Providing the command server is running on the queue manager receiving the message, the command server interprets it as a command message and runs the command. To get the replies, the application issues an **MQGET** call and the reply data is returned in another data structure. The application can then process the reply and act accordingly.

Note: Unlike MQSC commands, PCF commands and their replies are not in a text format that you can read.

Briefly, these are some of the things needed to create a PCF command message:

Message descriptor

This is a standard WebSphere MQ message descriptor, in which:

- Message type (*MsgType*) is MQMT_REQUEST.
- Message format (*Format*) is MQFMT_ADMIN.

Application data

Contains the PCF message including the PCF header, in which:

- The PCF message type (*Type*) specifies MQCFT_COMMAND.
- The command identifier specifies the command, for example, *Change Queue* (MQCMD_CHANGE_Q).

For a complete description of the PCF data structures and how to implement them, see “Introduction to Programmable Command Formats” on page 6.

PCF object attributes

Object attributes in PCF are not limited to eight characters as they are for MQSC commands. They are shown in this guide in italics. For example, the PCF equivalent of RQMNAME is *RemoteQMgrName*.

Escape PCFs

Escape PCFs are PCF commands that contain MQSC commands within the message text. You can use PCFs to send commands to a remote queue manager. For more information about escape PCFs, see

 [Escape \(WebSphere MQ V7.1 Reference\)](#).

Introduction to Programmable Command Formats

Programmable Command Formats (PCFs) define command and reply messages that can be exchanged between a program and any queue manager (that supports PCFs) in a network. PCFs simplify queue manager administration and other network administration. They can be used to solve the problem of complex administration of distributed networks especially as networks grow in size and complexity.

The Programmable Command Formats described in this product documentation are supported by:

- IBM WebSphere MQ for AIX®
- IBM WebSphere MQ for HP-UX
- IBM WebSphere MQ for IBM i
- IBM WebSphere MQ for Linux
- IBM WebSphere MQ for Solaris
- IBM WebSphere MQ for Windows
- IBM WebSphere MQ for z/OS
- IBM WebSphere MQ for HP Integrity NonStop Server v5.3
- IBM WebSphere MQ for HP OpenVMS, V5.3

The problem PCF commands solve

The administration of distributed networks can become complex. The problems of administration continue to grow as networks increase in size and complexity.

Examples of administration specific to messaging and queuing include:

- Resource management.
For example, queue creation and deletion.
- Performance monitoring.
For example, maximum queue depth or message rate.
- Control.
For example, tuning queue parameters such as maximum queue depth, maximum message length, and enabling and disabling queues.
- Message routing.
Definition of alternative routes through a network.

WebSphere MQ PCF commands can be used to simplify queue manager administration and other network administration. PCF commands allow you to use a single application to perform network administration from a single queue manager within the network.

What are PCFs?

PCFs define command and reply messages that can be exchanged between a program and any queue manager (that supports PCFs) in a network. You can use PCF commands in a systems management application program for administration of WebSphere MQ objects: authentication information objects, channels, channel listeners, namelists, process definitions, queue managers, queues, services, and storage classes. The application can operate from a single point in the network to communicate command and reply information with any queue manager, local, or remote, using the local queue manager.

Each queue manager has an administration queue with a standard queue name and your application can send PCF command messages to that queue. Each queue manager also has a command server to service the command messages from the administration queue. PCF command messages can therefore be processed by any queue manager in the network and the reply data can be returned to your application, using your specified reply queue. PCF commands and reply messages are sent and received using the normal Message Queue Interface (MQI).

For a list of the available PCF commands, including their parameters, see  Definitions of the Programmable Command Formats (*WebSphere MQ V7.1 Reference*).

Using Programmable Command Formats

You can use PCFs in a systems management program for WebSphere MQ remote administration.

This section includes:

- “PCF command messages”
- “Responses” on page 10
- “Extended responses” on page 12
-  Rules for naming IBM WebSphere MQ objects (*WebSphere MQ V7.1 Product Overview Guide*)
- “Authority checking for PCF commands” on page 14

PCF command messages:

PCF command messages consist of a PCF header, parameters identified in that header and also user-defined message data. The messages are issued using Message Queue interface calls.

Each command and its parameters are sent as a separate command message containing a PCF header followed by a number of parameter structures (see  MQCFH - PCF header (*WebSphere MQ V7.1 Reference*)). The PCF header identifies the command and the number of parameter structures that follow in the same message. Each parameter structure provides a parameter to the command.

Replies to the commands, generated by the command server, have a similar structure. There is a PCF header, followed by a number of parameter structures. Replies can consist of more than one message but commands always consist of one message only.

On platforms other than z/OS, the queue to which the PCF commands are sent is always called the SYSTEM.ADMIN.COMMAND.QUEUE. On z/OS, commands are sent to SYSTEM.COMMAND.INPUT, although SYSTEM.ADMIN.COMMAND.QUEUE can be an alias for it. The command server servicing this queue sends the replies to the queue defined by the *ReplyToQ* and *ReplyToQMgr* fields in the message descriptor of the command message.

How to issue PCF command messages

Use the normal Message Queue Interface (MQI) calls, MQPUT, MQGET, and so on, to put and retrieve PCF command and response messages to and from their queues.

Note:

Ensure that the command server is running on the target queue manager for the PCF command to process on that queue manager.

For a list of supplied header files, see  WebSphere MQ COPY, header, include, and module files (*WebSphere MQ V7.1 Reference*).

Message descriptor for a PCF command

The WebSphere MQ message descriptor is fully documented in  MQMD – Message descriptor (*WebSphere MQ V7.1 Reference*).

A PCF command message contains the following fields in the message descriptor:

Report

Any valid value, as required.

MsgType

This field must be MQMT_REQUEST to indicate a message requiring a response.

Expiry

Any valid value, as required.

Feedback

Set to MQFB_NONE

Encoding

If you are sending to IBM i, Windows, UNIX or Linux systems, set this field to the encoding used for the message data; conversion is performed if necessary.

CodedCharSetId

If you are sending to IBM i, Windows, UNIX or Linux systems, set this field to the coded character-set identifier used for the message data; conversion is performed if necessary.

Format

Set to MQFMT_ADMIN.

Priority

Any valid value, as required.

Persistence

Any valid value, as required.

MsgId

The sending application can specify any value, or MQMI_NONE can be specified to request the queue manager to generate a unique message identifier.

CorrelId

The sending application can specify any value, or MQCI_NONE can be specified to indicate no correlation identifier.

ReplyToQ

The name of the queue to receive the response.

ReplyToQMGr

The name of the queue manager for the response (or blank).

Message context fields

These fields can be set to any valid values, as required. Normally the Put message option MQPMO_DEFAULT_CONTEXT is used to set the message context fields to the default values.

If you are using a version-2 MQMD structure, you must set the following additional fields:

GroupId

Set to MQGI_NONE

MsgSeqNumber

Set to 1

Offset

Set to 0

MsgFlags

Set to MQMF_NONE

OriginalLength

Set to MQOL_UNDEFINED

Sending user data

The PCF structures can also be used to send user-defined message data. In this case the message descriptor *Format* field must be set to MQFMT_PCF.

Sending and receiving PCF messages in a specified queue:

Sending PCF messages to a specified queue

To send a message to a specified queue, the `mqPutBag` call converts the contents of the specified bag into a PCF message and sends the message to the specified queue. The contents of the bag are left unchanged after the call.

As input to this call, you must supply:

- An MQI connection handle.
- An object handle for the queue on which the message is to be placed.
- A message descriptor. For more information about the message descriptor, see  MQMD – Message descriptor (*WebSphere MQ V7.1 Reference*).
- Put Message Options using the MQPMO structure. For more information about the MQPMO structure, see  MQPMO – Put-message options (*WebSphere MQ V7.1 Reference*).
- The handle of the bag to be converted to a message.

Note: If the bag contains an administration message and the `mqAddInquiry` call was used to insert values into the bag, the value of the MQIASY_COMMAND data item must be an INQUIRE command recognized by the MQAI.

For a full description of the `mqPutBag` call, see  `mqPutBag` (*WebSphere MQ V7.1 Reference*).

Receiving PCF messages from a specified queue

To receive a message from a specified queue, the `mqGetBag` call gets a PCF message from a specified queue and converts the message data into a data bag.

As input to this call, you must supply:

- An MQI connection handle.
- An object handle of the queue from which the message is to be read.
- A message descriptor. Within the MQMD structure, the *Format* parameter must be MQFMT_ADMIN, MQFMT_EVENT, or MQFMT_PCF.

Note: If the message is received within a unit of work (that is, with the MQGMO_SYNCPOINT option) and the message has an unsupported format, the unit of work can be backed out. The message is then reinstated on the queue and can be retrieved using the MQGET call instead of the mqGetBag call. For more information about the message descriptor, see  MQGMO – Get-message options (*WebSphere MQ V7.1 Reference*).

- Get Message Options using the MQGMO structure. For more information about the MQGMO structure, see  MQMD – Message descriptor (*WebSphere MQ V7.1 Reference*).
- The handle of the bag to contain the converted message.

For a full description of the mqGetBag call, see  mqGetBag (*WebSphere MQ V7.1 Reference*).

Responses:

In response to each command, the command server generates one or more response messages. A response message has a similar format to a command message.

The PCF header has the same command identifier value as the command to which it is a response (see  MQCFH - PCF header (*WebSphere MQ V7.1 Reference*) for details). The message identifier and correlation identifier are set according to the report options of the request.

If the PCF header type of the command message is MQCFT_COMMAND, standard responses only are generated. Such commands are supported on all platforms except z/OS. Older applications do not support PCF on z/OS; the WebSphere MQ Windows Explorer is one such application (however, the Version 6.0 or later WebSphere MQ Explorer does support PCF on z/OS).

If the PCF header type of the command message is MQCFT_COMMAND_XR, either extended or standard responses are generated. Such commands are supported on z/OS and some other platforms. Commands issued on z/OS generate only extended responses. On other platforms, either type of response might be generated.

If a single command specifies a generic object name, a separate response is returned in its own message for each matching object. For response generation, a single command with a generic name is treated as multiple individual commands (except for the control field MQCFC_LAST or MQCFC_NOT_LAST). Otherwise, one command message generates one response message.

Certain PCF responses might return a structure even when it is not requested. This structure is shown in the definition of the response ( Definitions of the Programmable Command Formats (*WebSphere MQ V7.1 Reference*)) as *always returned*. The reason that, for these responses, it is necessary to name the objects in the response to identify which object the data applies.

Message descriptor for a response

A response message has the following fields in the message descriptor:

MsgType

This field is MQMT_REPLY.

MsgId

This field is generated by the queue manager.

CorrelId

This field is generated according to the report options of the command message.

Format

This field is MQFMT_ADMIN.

Encoding

Set to MQENC_NATIVE.

CodedCharSetId

Set to MQCCSI_Q_MGR.

Persistence

The same as in the command message.

Priority

The same as in the command message.

The response is generated with MQPMO_PASS_IDENTITY_CONTEXT.

Standard responses:

Command messages with a header type of MQCFT_COMMAND, standard responses are generated. Such commands are supported on all platforms except z/OS.

There are three types of standard response:

- OK response
- Error response
- Data response

OK response

This response consists of a message starting with a command format header, with a *CompCode* field of MQCC_OK or MQCC_WARNING.

For MQCC_OK, the *Reason* is MQRC_NONE.

For MQCC_WARNING, the *Reason* identifies the nature of the warning. In this case the command format header might be followed by one or more warning parameter structures appropriate to this reason code.

In either case, for an inquire command further parameter structures might follow as described in the following sections.

Error response

If the command has an error, one or more error response messages are sent (more than one might be sent even for a command that would normally have only a single response message). These error response messages have MQCFC_LAST or MQCFC_NOT_LAST set as appropriate.

Each such message starts with a response format header, with a *CompCode* value of MQCC_FAILED and a *Reason* field that identifies the particular error. In general, each message describes a different error. In addition, each message has either zero or one (never more than one) error parameter structures following the header. This parameter structure, if there is one, is an MQCFIN structure, with a *Parameter* field containing one of the following:

- MQIACF_PARAMETER_ID

The *Value* field in the structure is the parameter identifier of the parameter that was in error (for example, MQCA_Q_NAME).

- MQIACF_ERROR_ID

This value is used with a *Reason* value (in the command format header) of MQRC_UNEXPECTED_ERROR. The *Value* field in the MQCFIN structure is the unexpected reason code received by the command server.

- MQIACF_SELECTOR

This value occurs if a list structure (MQCFIL) sent with the command contains a duplicate selector or one that is not valid. The *Reason* field in the command format header identifies the error, and the *Value* field in the MQCFIN structure is the parameter value in the MQCFIL structure of the command that was in error.

- MQIACF_ERROR_OFFSET

This value occurs when there is a data compare error on the Ping Channel command. The *Value* field in the structure is the offset of the Ping Channel compare error.

- MQIA_CODED_CHAR_SET_ID

This value occurs when the coded character-set identifier in the message descriptor of the incoming PCF command message does not match that of the target queue manager. The *Value* field in the structure is the coded character-set identifier of the queue manager.

The last (or only) error response message is a summary response, with a *CompCode* field of MQCC_FAILED, and a *Reason* field of MQRCCF_COMMAND_FAILED. This message has no parameter structure following the header.

Data response

This response consists of an OK response (as described earlier) to an inquire command. The OK response is followed by additional structures containing the requested data as described in  Definitions of the Programmable Command Formats (*WebSphere MQ V7.1 Reference*).

Applications must not depend upon these additional parameter structures being returned in any particular order.

Extended responses:

Commands issued on z/OS generate extended responses.

There are three types of extended response:

- Message response, with type MQCFT_XR_MSG
- Item response, with type MQCFT_XR_ITEM
- Summary response, with type MQCFT_XR_SUMMARY

Each command can generate one, or more, sets of responses. Each set of responses comprises one or more messages, numbered sequentially from 1 in the *MsgSeqNumber* field of the PCF header. The *Control* field of the last (or only) response in each set has the value MQCFC_LAST. For all other responses in the set, this value is MQCFC_NOT_LAST.

Any response can include one, or more, optional MQCFBS structures in which the *Parameter* field is set to MQBACF_RESPONSE_SET, the value being a response set identifier. Identifiers are unique and identify the set of responses which contain the response. For every set of responses, there is an MQCFBS structure that identifies it.

Extended responses have at least two parameter structures:

- An MQCFBS structure with the *Parameter* field set to MQBACF_RESPONSE_ID. The value in this field is the identifier of the set of responses to which the response belongs. The identifier in the first set is arbitrary. In subsequent sets, the identifier is one previously notified in an MQBACF_RESPONSE_SET structure.
- An MQCFST structure with the *Parameter* field set to MQCACF_RESPONSE_Q_MGR_NAME, the value being the name of the queue manager from which the set of responses come.

Many responses have additional parameter structures, and these structures are described in the following sections.

You cannot determine in advance how many responses there are in a set other than by getting responses until one with MQCFC_LAST is found. Neither can you determine in advance how many sets of responses there are as any set might include MQBACF_RESPONSE_SET structures to indicate that additional sets are generated.

Extended responses to Inquire commands

Inquire commands normally generate an item response (type MQCFT_XR_ITEM) for each item found that matches the specified search criteria. The item response has a *CompCode* field in the header with a value of MQCC_OK, and a *Reason* field with a value of MQRC_NONE. It also includes other parameter structures describing the item and its requested attributes, as described in  Definitions of the Programmable Command Formats.

If an item is in error, the *CompCode* field in the header has a value of MQCC_FAILED and the *Reason* field identifies the particular error. Additional parameter structures are included to identify the item.

Certain Inquire commands might return general (not name-specific) message responses in addition to the item responses. These responses are informational, or error, responses of the type MQCFT_XR_MSG.

If the Inquire command succeeds, there might, optionally, be a summary response (type MQCFT_XR_SUMMARY), with a *CompCode* value of MQCC_OK, and a *Reason* field value of MQRC_NONE.

If the Inquire command fails, item responses might be returned, and there might optionally be a summary response (type MQCFT_XR_SUMMARY), with a *CompCode* value of MQCC_FAILED, and a *Reason* field value of MQRCCF_COMMAND_FAILED.

Extended responses to commands other than Inquire

Successful commands generate message responses in which the *CompCode* field in the header has a value of MQCC_OK, and the *Reason* field has a value of MQRC_NONE. There is always at least one message; it might be informational (MQCFT_XR_MSG) or a summary (MQCFT_XR_SUMMARY). There might optionally be additional informational (type MQCFT_XR_MSG) messages. Each informational message might include a number of additional parameter structures with information about the command; see the individual command descriptions for the structures that can occur.

Commands that fail generate error message responses (type MQCFT_XR_MSG), in which the *CompCode* field in the header has a value of MQCC_FAILED and the *Reason* field identifies the particular error. Each message might include a number of additional parameter structures with information about the error: see the individual error descriptions for the structures that can occur. Informational message responses might be generated. There might, optionally, be a summary response (MQCFT_XR_SUMMARY), with a *CompCode* value of MQCC_FAILED, and a *Reason* field value of MQRCCF_COMMAND_FAILED.

Extended responses to commands using CommandScope

If a command uses the *CommandScope* parameter, or causes a command using the *CommandScope* parameter to be generated, there is an initial response set from the queue manager where the command was received. Then a separate set, or sets, of responses is generated for each queue manager to which the command is directed (as if multiple individual commands were issued). Finally, there is a response set from the receiving queue manager which includes an overall summary response (type MQCFT_XR_SUMMARY). The MQCACF_RESPONSE_Q_MGR_NAME parameter structure identifies the queue manager that generates each set.

The initial response set has the following additional parameter structures:

- MQIACF_COMMAND_INFO (MQCFIN). Possible values in this structure are MQCMDI_CMDSCOPE_ACCEPTED or MQCMDI_CMDSCOPE_GENERATED.
- MQIACF_CMDSCOPE_Q_MGR_COUNT (MQCFIN). This structure indicates the number of queue managers to which the command is sent.

Authority checking for PCF commands:

When a PCF command is processed, the *UserIdentifier* from the message descriptor in the command message is used for the required WebSphere MQ object authority checks. Authority checking is implemented differently on each platform as described in this topic.

The checks are performed on the system on which the command is being processed; therefore this user ID must exist on the target system and have the required authorities to process the command. If the message has come from a remote system, one way of achieving the ID existing on the target system is to have a matching user ID on both the local and remote systems.

WebSphere MQ for IBM i

In order to process any PCF command, the user ID must have *dsp* authority for the WebSphere MQ object on the target system.

In addition, WebSphere MQ object authority checks are performed for certain PCF commands, as shown in Table 1 on page 16.

In most cases these checks are the same checks as those checks performed by the equivalent WebSphere MQ CL commands issued on a local system. See the “Setting up security on IBM i” on page 492, for more information about the mapping from WebSphere MQ authorities to IBM i system authorities, and the authority requirements for the WebSphere MQ CL commands. Details of security concerning exits are given in the “Link level security using a security exit” on page 438 documentation.

To process any of the following commands the user ID must be a member of the group profile QMQMADM:

- Ping Channel
- Change Channel
- Copy Channel
- Create Channel
- Delete Channel
- Reset Channel
- Resolve Channel
- Start Channel
- Stop Channel
- Start Channel Initiator
- Start Channel Listener

WebSphere MQ for Windows, UNIX and Linux systems

In order to process any PCF command, the user ID must have *dsp* authority for the queue manager object on the target system. In addition, WebSphere MQ object authority checks are performed for certain PCF commands, as shown in Table 1 on page 16.

To process any of the following commands the user ID must belong to group *mqm*.

Note: For Windows **only**, the user ID can belong to group *Administrators* or group *mqm*.

- Change Channel
- Copy Channel
- Create Channel
- Delete Channel
- Ping Channel
- Reset Channel
- Start Channel
- Stop Channel
- Start Channel Initiator
- Start Channel Listener
- Resolve Channel
- Reset Cluster
- Refresh Cluster
- Suspend Queue Manager
- Resume Queue Manager

WebSphere MQ for HP OpenVMS and HP Integrity NonStop Server

In order to process any PCF command, the user ID must have *dsp* authority for the queue manager object on the target system. In addition, WebSphere MQ object authority checks are performed for certain PCF commands, as shown in Table 1 on page 16.

To process any of the following commands the user ID must belong to group *mqm*:

- Change Channel
- Copy Channel
- Create Channel
- Delete Channel
- Ping Channel
- Reset Channel
- Start Channel
- Stop Channel
- Start Channel Initiator
- Start Channel Listener
- Resolve Channel
- Reset Cluster
- Refresh Cluster
- Suspend Queue Manager
- Resume Queue Manager

WebSphere MQ Object authorities

Table 1. Windows, HP OpenVMS Alpha, HP Integrity NonStop Server, IBM i, UNIX and Linux systems - object authorities

Command	WebSphere MQ object authority	Class authority (for object type)
Change Authentication Information	dsp and chg	n/a
Change Channel	dsp and chg	n/a
Change Channel Listener	dsp and chg	n/a
Change Client Connection Channel	dsp and chg	n/a
Change Namelist	dsp and chg	n/a
Change Process	dsp and chg	n/a
Change Queue	dsp and chg	n/a
Change Queue Manager	chg <i>see Note 3 and Note 5</i>	n/a
Change Service	dsp and chg	n/a
Clear Queue	clr	n/a
Copy Authentication Information	dsp	crt
Copy Authentication Information (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> chg	crt
Copy Channel	dsp	crt
Copy Channel (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> chg	crt
Copy Channel Listener	dsp	crt
Copy Channel Listener (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> chg	crt
Copy Client Connection Channel	dsp	crt
Copy Client Connection Channel (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> chg	crt
Copy Namelist	dsp	crt
Copy Namelist (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> dsp and chg	crt
Copy Process	dsp	crt
Copy Process (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> chg	crt
Copy Queue	dsp	crt
Copy Queue (Replace) <i>see Note 1</i>	<i>from:</i> dsp <i>to:</i> dsp and chg	crt
Create Authentication Information	(<i>system default authentication information</i>) dsp	crt
Create Authentication Information (Replace) <i>see Note 1</i>	(<i>system default authentication information</i>) dsp <i>to:</i> chg	crt
Create Channel	(<i>system default channel</i>) dsp	crt
Create Channel (Replace) <i>see Note 1</i>	(<i>system default channel</i>) dsp <i>to:</i> chg	crt
Create Channel Listener	(<i>system default listener</i>) dsp	crt
Create Channel Listener (Replace) <i>see Note 1</i>	(<i>system default listener</i>) dsp <i>to:</i> chg	crt
Create Client Connection Channel	(<i>system default channel</i>) dsp	crt
Create Client Connection Channel (Replace) <i>see Note 1</i>	(<i>system default channel</i>) dsp <i>to:</i> chg	crt
Create Namelist	(<i>system default namelist</i>) dsp	crt
Create Namelist (Replace) <i>see Note 1</i>	(<i>system default namelist</i>) dsp <i>to:</i> dsp and chg	crt

Table 1. Windows, HP OpenVMS Alpha, HP Integrity NonStop Server, IBM i, UNIX and Linux systems - object authorities (continued)

Command	WebSphere MQ object authority	Class authority (for object type)
Create Process	(system default process) dsp	crt
Create Process (Replace) <i>see Note 1</i>	(system default process) dsp to: chg	crt
Create Queue	(system default queue) dsp	crt
Create Queue (Replace) <i>see Note 1</i>	(system default queue) dsp to: dsp and chg	crt
Create Service	(system default queue) dsp	crt
Create Service (Replace) <i>see Note 1</i>	(system default queue) dsp to: chg	crt
Delete Authentication Information	dsp and dlt	n/a
Delete Authority Record	(queue manager object) chg <i>see Note 4</i>	<i>see Note 4</i>
Delete Channel	dsp and dlt	n/a
Delete Channel Listener	dsp and dlt	n/a
Delete Client Connection Channel	dsp and dlt	n/a
Delete Namelist	dsp and dlt	n/a
Delete Process	dsp and dlt	n/a
Delete Queue	dsp and dlt	n/a
Delete Service	dsp and dlt	n/a
Inquire Authentication Information	dsp	n/a
Inquire Authority Records	<i>see Note 4</i>	<i>see Note 4</i>
Inquire Channel	dsp	n/a
Inquire Channel Listener	dsp	n/a
Inquire Channel Status (for ChannelType MQCHT_CLSSDR)	inq	n/a
Inquire Client Connection Channel	dsp	n/a
Inquire Namelist	dsp	n/a
Inquire Process	dsp	n/a
Inquire Queue	dsp	n/a
Inquire Queue Manager	<i>see note 3</i>	n/a
Inquire Queue Status	dsp	n/a
Inquire Service	dsp	n/a
Ping Channel	ctrl	n/a
Ping Queue Manager	<i>see note 3</i>	n/a
Refresh Queue Manager	(queue manager object) chg	n/a
Refresh Security (for SecurityType MQSECTYPE_SSL)	(queue manager object) chg	n/a
Reset Channel	ctrlx	n/a
Reset Queue Manager	(queue manager object) chg	n/a
Reset Queue Statistics	dsp and chg	n/a
Resolve Channel	ctrlx	n/a
Set Authority Record	(queue manager object) chg <i>see Note 4</i>	<i>see Note 4</i>
Start Channel	ctrl	n/a

Table 1. Windows, HP OpenVMS Alpha, HP Integrity NonStop Server, IBM i, UNIX and Linux systems - object authorities (continued)

Command	WebSphere MQ object authority	Class authority (for object type)
Stop Channel	ctrl	n/a
Stop Connection	(queue manager object) chg	n/a
Start Listener	ctrl	n/a
Stop Listener	ctrl	n/a
Start Service	ctrl	n/a
Stop Service	ctrl	n/a
Escape	see Note 2	see Note 2
<p>Note:</p> <ol style="list-style-type: none"> 1. This command applies if the object to be replaced does exist, otherwise the authority check is as for Create, or Copy without Replace. 2. The required authority is determined by the MQSC command defined by the escape text, and it is equivalent to one of the previous commands. 3. In order to process any PCF command, the user ID must have dsp authority for the queue manager object on the target system. 4. This PCF command is authorized unless the command server has been started with the -a parameter. By default the command server starts when the Queue Manager is started, and without the -a parameter. See the System Administration Guide for further information. 5. Granting a user ID <i>chg</i> authority for a queue manager gives the ability to set authority records for all groups and users. Do not grant this authority to ordinary users or applications. 		

WebSphere MQ also supplies some channel security exit points so that you can supply your own user exit programs for security checking. Details are given in  Displaying a channel (*WebSphere MQ V7.1 Installing Guide*) manual.

WebSphere MQ for z/OS

See  Task 1: Identify the z/OS system parameters (*WebSphere MQ V7.1 Installing Guide*) for information about authority checking on z/OS.

Using the MQAI to simplify the use of PCFs

The MQAI is an administration interface to WebSphere MQ that is available on the AIX, HP-UX, IBM i, Linux, Solaris, and Windows platforms.

The MQAI performs administration tasks on a queue manager through the use of *data bags*. Data bags allow you to handle properties (or parameters) of objects in a way that is easier than using PCFs.

Use the MQAI in the following ways:

To simplify the use of PCF messages

The MQAI is an easy way to administer WebSphere MQ; you do not have to write your own PCF messages, avoiding the problems associated with complex data structures.

To pass parameters in programs written using MQI calls, the PCF message must contain the command and details of the string or integer data. To do this, you need several statements in your program for every structure, and memory space must be allocated. This task can be long and laborious.

Programs written using the MQAI pass parameters into the appropriate data bag and you need only one statement for each structure. The use of MQAI data bags removes the need for you to handle arrays and allocate storage, and provides some degree of isolation from the details of the PCF.

To handle error conditions more easily

It is difficult to get return codes back from PCF commands, but the MQAI makes it easier for the program to handle error conditions.

After you have created and populated your data bag, you can send an administration command message to the command server of a queue manager, using the **mqExecute** call, which waits for any response messages. The **mqExecute** call handles the exchange with the command server and returns responses in a *response bag*.

For more information about the MQAI, see “Introduction to the WebSphere MQ Administration Interface (MQAI).”

Introduction to the WebSphere MQ Administration Interface (MQAI)

WebSphere MQ Administration Interface (MQAI) is a programming interface to WebSphere® MQ. It performs administration tasks on a WebSphere MQ queue manager using data bags to handle properties (or parameters) of objects in a way that is easier than using Programmable Command Formats (PCFs).

MQAI concepts and terminology

The MQAI is a programming interface to WebSphere® MQ, using the C language and also Visual Basic for Windows®. It is available on platforms other than z/OS®.

It performs administration tasks on a WebSphere MQ queue manager using data bags. Data bags allow you to handle properties (or parameters) of objects in a way that is easier than using the other administration interface, Programmable Command Formats (PCFs). The MQAI offers easier manipulation of PCFs than using the MQGET and MQPUT calls.

For more information about data bags, see “Data bags” on page 48. For more information about PCFs, see “Introduction to Programmable Command Formats” on page 6

Use of the MQAI

You can use the MQAI to:

- Simplify the use of PCF messages. The MQAI is an easy way to administer WebSphere MQ; you do not have to write your own PCF messages and thus avoid the problems associated with complex data structures.
- Handle error conditions more easily. It is difficult to get return codes back from the WebSphere MQ script (MQSC) commands, but the MQAI makes it easier for the program to handle error conditions.
- Exchange data between applications. The application data is sent in PCF format and packed and unpacked by the MQAI. If your message data consists of integers and character strings, you can use the MQAI to take advantage of WebSphere MQ built-in data conversion for PCF data. This avoids the need to write data-conversion exits. For more information on using MQAI to administer WebSphere MQ and to exchange data between applications, see “Using the MQAI to simplify the use of PCFs” on page 18.

Examples of using the MQAI

The list shown gives some example programs that demonstrate the use of MQAI. The samples perform the following tasks:

1. Create a local queue. "Creating a local queue (amqsaicq.c)"
2. Display events on the screen using a simple event monitor. "Displaying events using an event monitor (amqsaiem.c)" on page 25
3. Print a list of all local queues and their current depths. "Inquiring about queues and printing information (amqsailq.c)" on page 42
4. Print a list of all channels and their types. "Inquire channel objects (amqsaicl.c)" on page 35

Building your MQAI application

To build your application using the MQAI, you link to the same libraries as you do for WebSphere MQ.

For information on how to build your WebSphere MQ applications, see  Building a WebSphere MQ application (*WebSphere MQ V7.1 Programming Guide*).

Hints and tips for configuring WebSphere MQ using MQAI

The MQAI uses PCF messages to send administration commands to the command server rather than dealing directly with the command server itself. Tips for configuring WebSphere MQ using the MQAI can be found in "Hints and tips for configuring WebSphere MQ" on page 47

WebSphere MQ Administration Interface (MQAI)

WebSphere MQ for Windows, AIX, IBM i, Linux, HP-UX, and Solaris support the WebSphere MQ Administration Interface (MQAI). The MQAI is a programming interface to WebSphere MQ that gives you an alternative to the MQL, for sending and receiving PCFs.

The MQAI uses *data bags* which allow you to handle properties (or parameters) of objects more easily than using PCFs directly by way of the MQAI.

The MQAI provides easier programming access to PCF messages by passing parameters into the data bag, so that only one statement is required for each structure. This access removes the need for the programmer to handle arrays and allocate storage, and provides some isolation from the details of PCF.

The MQAI administers WebSphere MQ by sending PCF messages to the command server and waiting for a response.

The MQAI is described in the second section of this manual. See the  Using Java (*WebSphere MQ V7.1 Programming Guide*) documentation for a description of a component object model interface to the MQAI.

Creating a local queue (amqsaicq.c)

```

/*****
/*
/* Program name: AMQSAICQ.C
/*
/* Description: Sample C program to create a local queue using the
/*               WebSphere MQ Administration Interface (MQAI).
/*
/* Statement:   Licensed Materials - Property of IBM
/*
/*               84H2000, 5765-B73
/*               84H2001, 5639-B42
/*               84H2002, 5765-B74
/*               84H2003, 5765-B75
/*               84H2004, 5639-B43
/*
/*               (C) Copyright IBM Corp. 1999, 2019
/*
*****/

```

```

/*****/
/* */
/* Function: */
/* AMQSAICQ is a sample C program that creates a local queue and is an */
/* example of the use of the mqExecute call. */
/* */
/* - The name of the queue to be created is a parameter to the program. */
/* */
/* - A PCF command is built by placing items into an MQAI bag. */
/* These are:- */
/* - The name of the queue */
/* - The type of queue required, which, in this case, is local. */
/* */
/* - The mqExecute call is executed with the command MQCMD_CREATE_Q. */
/* The call generates the correct PCF structure. */
/* The call receives the reply from the command server and formats into */
/* the response bag. */
/* */
/* - The completion code from the mqExecute call is checked and if there */
/* is a failure from the command server then the code returned by the */
/* command server is retrieved from the system bag that is */
/* embedded in the response bag to the mqExecute call. */
/* */
/* Note: The command server must be running. */
/* */
/* */

/*****/
/* */
/* AMQSAICQ has 2 parameters - the name of the local queue to be created */
/* - the queue manager name (optional) */
/* */

/*****/
/* Includes */
/*****/
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <ctype.h>

#include <cmqc.h> /* MQI */
#include <cmqfc.h> /* PCF */
#include <cmqbc.h> /* MQAI */

void CheckCallResult(MQCHAR *, MQLONG , MQLONG );
void CreateLocalQueue(MQHCONN, MQCHAR *);

int main(int argc, char *argv[])
{
    MQHCONN hConn; /* handle to WebSphere MQ connection */
    MQCHAR QMName[MQ_Q_MGR_NAME_LENGTH+1]=""; /* default QMgr name */
    MQLONG connReason; /* MQCONN reason code */
    MQLONG compCode; /* completion code */
    MQLONG reason; /* reason code */

    /*****/
    /* First check the required parameters */
    /*****/
    printf("Sample Program to Create a Local Queue\n");

```

```

if (argc < 2)
{
    printf("Required parameter missing - local queue name\n");
    exit(99);
}

/*****
/* Connect to the queue manager */
/*****
if (argc > 2)
    strncpy(QMName, argv[2], (size_t)MQ_Q_MGR_NAME_LENGTH);
    MQCONN(QMName, &hConn, &compCode, &connReason);

/*****
/* Report reason and stop if connection failed */
/*****
    if (compCode == MQCC_FAILED)
    {
        CheckCallResult("MQCONN", compCode, connReason);
        exit( (int)connReason);
    }

/*****
/* Call the routine to create a local queue, passing the handle to the */
/* queue manager and also passing the name of the queue to be created. */
/*****
    CreateLocalQueue(hConn, argv[1]);

/*****
/* Disconnect from the queue manager if not already connected */
/*****
    if (connReason != MQRC_ALREADY_CONNECTED)
    {
        MQDISC(&hConn, &compCode, &reason);
        CheckCallResult("MQDISC", compCode, reason);
    }
    return 0;
}

/*****
/*
/* Function:      CreateLocalQueue */
/* Description:  Create a local queue by sending a PCF command to the command */
/*               server. */
/*
/*
/*****
/*
/* Input Parameters:  Handle to the queue manager */
/*                   Name of the queue to be created */
/*
/* Output Parameters: None */
/*
/* Logic: The mqExecute call is executed with the command MQCMD_CREATE_Q. */
/*        The call generates the correct PCF structure. */
/*        The default options to the call are used so that the command is sent*/
/*        to the SYSTEM.ADMIN.COMMAND.QUEUE. */
/*        The reply from the command server is placed on a temporary dynamic */
/*        queue. */
/*        The reply is read from the temporary queue and formatted into the */

```

```

/*      response bag.                                          */
/*      */
/*      The completion code from the mqExecute call is checked and if there */
/*      is a failure from the command server then the code returned by the */
/*      command server is retrieved from the system bag that is          */
/*      embedded in the response bag to the mqExecute call.              */
/*      */
/*****/
void CreateLocalQueue(MQHCONN hConn, MQCHAR *qName)
{
    MQLONG reason;                /* reason code          */
    MQLONG compCode;              /* completion code      */
    MQHBAG commandBag = MQHB_UNUSABLE_HBAG; /* command bag for mqExecute */
    MQHBAG responseBag = MQHB_UNUSABLE_HBAG; /* response bag for mqExecute */
    MQHBAG resultBag;            /* result bag from mqExecute */
    MQLONG mqExecuteCC;          /* mqExecute completion code */
    MQLONG mqExecuteRC;          /* mqExecute reason code  */

    printf("\nCreating Local Queue %s\n\n", qName);

    /*****/
    /* Create a command Bag for the mqExecute call. Exit the function if the */
    /* create fails.                                                         */
    /*****/
    mqCreateBag(MQCBO_ADMIN_BAG, &commandBag, &compCode, &reason);
    CheckCallResult("Create the command bag", compCode, reason);
    if (compCode !=MQCC_OK)
        return;

    /*****/
    /* Create a response Bag for the mqExecute call, exit the function if the */
    /* create fails.                                                         */
    /*****/
    mqCreateBag(MQCBO_ADMIN_BAG, &responseBag, &compCode, &reason);
    CheckCallResult("Create the response bag", compCode, reason);
    if (compCode !=MQCC_OK)
        return;

    /*****/
    /* Put the name of the queue to be created into the command bag. This will */
    /* be used by the mqExecute call.                                         */
    /*****/
    mqAddString(commandBag, MQCA_Q_NAME, MQBL_NULL_TERMINATED, qName, &compCode,
                &reason);
    CheckCallResult("Add q name to command bag", compCode, reason);

    /*****/
    /* Put queue type of local into the command bag. This will be used by the */
    /* mqExecute call.                                                         */
    /*****/
    mqAddInteger(commandBag, MQIA_Q_TYPE, MQQT_LOCAL, &compCode, &reason);
    CheckCallResult("Add q type to command bag", compCode, reason);

    /*****/
    /* Send the command to create the required local queue.                   */
    /* The mqExecute call will create the PCF structure required, send it to */
    /* the command server and receive the reply from the command server into */
    /* the response bag.                                                       */
    /*****/

```

```

mqExecute(hConn,          /* WebSphere MQ connection handle      */
          MQCMD_CREATE_Q, /* Command to be executed                */
          MQHB_NONE,     /* No options bag                        */
          commandBag,    /* Handle to bag containing commands     */
          responseBag,   /* Handle to bag to receive the response*/
          MQHO_NONE,     /* Put msg on SYSTEM.ADMIN.COMMAND.QUEUE*/
          MQHO_NONE,     /* Create a dynamic q for the response  */
          &compCode,     /* Completion code from the mqExecute   */
          &reason);     /* Reason code from mqExecute call      */

if (reason == MQRC_CMD_SERVER_NOT_AVAILABLE)
{
    printf("Please start the command server: <strmqcsv QMgrName>\n")
    MQDISC(&hConn, &compCode, &reason);
    CheckCallResult("MQDISC", compCode, reason);
    exit(98);
}

/*****
/* Check the result from mqExecute call and find the error if it failed. */
*****/
if ( compCode == MQCC_OK )
    printf("Local queue %s successfully created\n", qName);
else
{
    printf("Creation of local queue %s failed: Completion Code = %d
           qName, compCode, reason);
    if (reason == MQRCCF_COMMAND_FAILED)
    {
        /*****
        /* Get the system bag handle out of the mqExecute response bag. */
        /* This bag contains the reason from the command server why the */
        /* command failed. */
        *****/
        mqInquireBag(responseBag, MQHA_BAG_HANDLE, 0, &resultBag, &compCode,
                    &reason);
        CheckCallResult("Get the result bag handle", compCode, reason);

        /*****
        /* Get the completion code and reason code, returned by the command */
        /* server, from the embedded error bag. */
        *****/
        mqInquireInteger(resultBag, MQIASY_COMP_CODE, MQIND_NONE, &mqExecuteCC,
                        &compCode, &reason);
        CheckCallResult("Get the completion code from the result bag",
                        compCode, reason);
        mqInquireInteger(resultBag, MQIASY_REASON, MQIND_NONE, &mqExecuteRC,
                        &compCode, &reason);
        CheckCallResult("Get the reason code from the result bag", compCode,
                        reason);
        printf("Error returned by the command server: Completion code = %d :
              Reason = %d\n", mqExecuteCC, mqExecuteRC);
    }
}
/*****
/* Delete the command bag if successfully created. */
*****/
if (commandBag != MQHB_UNUSABLE_HBAG)
{

```

```

        mqDeleteBag(&commandBag, &compCode, &reason);
        CheckCallResult("Delete the command bag", compCode, reason);
    }

    /*****
    /* Delete the response bag if successfully created.
    /*
    /*****
    if (responseBag != MQHB_UNUSABLE_HBAG)
    {
        mqDeleteBag(&responseBag, &compCode, &reason);
        CheckCallResult("Delete the response bag", compCode, reason);
    }
} /* end of CreateLocalQueue */

```

```

/*****
/*
/* Function: CheckCallResult
/*
/*
/*****
/*
/* Input Parameters: Description of call
/*                      Completion code
/*                      Reason code
/*
/* Output Parameters: None
/*
/* Logic: Display the description of the call, the completion code and the
/*          reason code if the completion code is not successful
/*
/*
/*****
void CheckCallResult(char *callText, MQLONG cc, MQLONG rc)
{
    if (cc != MQCC_OK)
        printf("%s failed: Completion Code = %d :
                Reason = %d\n", callText, cc, rc);
}

```

Displaying events using an event monitor (amqsaiem.c)

```

/*****
/*
/* Program name: AMQSAIEM.C
/*
/*
/* Description: Sample C program to demonstrate a basic event monitor
/*                using the WebSphere MQ Admin Interface (MQAI).
/* Licensed Materials - Property of IBM
/*
/*
/* 63H9336
/* (c) Copyright IBM Corp. 1999, 2005 All Rights Reserved.
/*
/*
/* US Government Users Restricted Rights - Use, duplication or
/* disclosure restricted by GSA ADP Schedule Contract with
/* IBM Corp.
/*****
/*
/* Function:
/* AMQSAIEM is a sample C program that demonstrates how to write a simple
/* event monitor using the mqGetBag call and other MQAI calls.
/*
/*

```

```

/* The name of the event queue to be monitored is passed as a parameter */
/* to the program. This would usually be one of the system event queues:- */
/*     SYSTEM.ADMIN.QMGR.EVENT      Queue Manager events      */
/*     SYSTEM.ADMIN.PERFM.EVENT     Performance events         */
/*     SYSTEM.ADMIN.CHANNEL.EVENT   Channel events             */
/*     SYSTEM.ADMIN.LOGGER.EVENT    Logger events              */
/*                                                                           */
/* To monitor the queue manager event queue or the performance event queue,*/
/* the attributes of the queue manager needs to be changed to enable      */
/* these events. For more information about this, see Part 1 of the        */
/* Programmable System Management book. The queue manager attributes can  */
/* be changed using either MQSC commands or the MQAI interface.          */
/* Channel events are enabled by default.                                  */
/*                                                                           */
/* Program logic                                                            */
/* Connect to the Queue Manager.                                          */
/* Open the requested event queue with a wait interval of 30 seconds.     */
/* Wait for a message, and when it arrives get the message from the queue */
/* and format it into an MQAI bag using the mqGetBag call.                */
/* There are many types of event messages and it is beyond the scope of  */
/* this sample to program for all event messages. Instead the program    */
/* prints out the contents of the formatted bag.                          */
/* Loop around to wait for another message until either there is an error */
/* or the wait interval of 30 seconds is reached.                        */
/*                                                                           */
/*****
/*
/* AMQSAIEM has 2 parameters - the name of the event queue to be monitored */
/*                               - the queue manager name (optional)      */
/*                                                                           */
/*****

/*****
/* Includes                                                                 */
/*****
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <ctype.h>

#include <cmqc.h>          /* MQI          */
#include <cmqcfh.h>       /* PCF          */
#include <cmqbc.h>        /* MQAI         */

/*****
/* Macros                                                                    */
/*****
#if MQAT_DEFAULT == MQAT_WINDOWS_NT
    #define Int64 "I64"
#elif defined(MQ_64_BIT)
    #define Int64 "l"
#else
    #define Int64 "ll"
#endif

/*****
/* Function prototypes                                                       */
/*****
void CheckCallResult(MQCHAR *, MQLONG , MQLONG);
void GetQEvents(MQHCONN, MQCHAR *);

```

```

int PrintBag(MQHBAG);
int PrintBagContents(MQHBAG, int);

/*****
/* Function: main */
*****/
int main(int argc, char *argv[])
{
    MQHCONN hConn; /* handle to connection */
    MQCHAR QMName[MQ_Q_MGR_NAME_LENGTH+1]=""; /* default QM name */
    MQLONG reason; /* reason code */
    MQLONG connReason; /* MQCONN reason code */
    MQLONG compCode; /* completion code */

    /*****
    /* First check the required parameters */
    *****/
    printf("Sample Event Monitor (times out after 30 secs)\n");
    if (argc < 2)
    {
        printf("Required parameter missing - event queue to be monitored\n");
        exit(99);
    }

    /*****
    /* Connect to the queue manager */
    *****/
    if (argc > 2)
        strncpy(QMName, argv[2], (size_t)MQ_Q_MGR_NAME_LENGTH);
    MQCONN(QMName, &hConn, &compCode, &connReason);
    /*****
    /* Report the reason and stop if the connection failed */
    *****/
    if (compCode == MQCC_FAILED)
    {
        CheckCallResult("MQCONN", compCode, connReason);
        exit( (int)connReason);
    }

    /*****
    /* Call the routine to open the event queue and format any event messages */
    /* read from the queue. */
    *****/
    GetQEvents(hConn, argv[1]);

    /*****
    /* Disconnect from the queue manager if not already connected */
    *****/
    if (connReason != MQRC_ALREADY_CONNECTED)
    {
        MQDISC(&hConn, &compCode, &reason);
        CheckCallResult("MQDISC", compCode, reason);
    }

    return 0;
}

/*****
/* */
*****/

```

```

/* Function: CheckCallResult                                     */
/*                                                             */
/*****                                                        */
/*                                                             */
/* Input Parameters:  Description of call                       */
/*                   Completion code                           */
/*                   Reason code                               */
/*                                                             */
/* Output Parameters: None                                     */
/*                                                             */
/* Logic: Display the description of the call, the completion  */
/*        reason code if the completion code is not successful */
/*                                                             */
/*****                                                        */
void CheckCallResult(char *callText, MQLONG cc, MQLONG rc)
{
    if (cc != MQCC_OK)
        printf("%s failed: Completion Code = %d : Reason = %d\n",
            callText, cc, rc);
}

/*****                                                        */
/* Function: GetQEvents                                         */
/*                                                             */
/*****                                                        */
/* Input Parameters:  Handle to the queue manager             */
/*                   Name of the event queue to be monitored  */
/*                                                             */
/* Output Parameters: None                                     */
/*                                                             */
/* Logic:  Open the event queue.                               */
/*         Get a message off the event queue and format the   */
/*         message into a bag.                                */
/*         A real event monitor would need to be programmed  */
/*         to deal with each type of event that it receives  */
/*         from the queue. This is outside the scope of this */
/*         sample, so instead, the contents of the bag are   */
/*         printed.                                           */
/*         The program waits for 30 seconds for an event     */
/*         message and then terminates if no more messages   */
/*         are available.                                     */
/*                                                             */
/*****                                                        */
void GetQEvents(MQHCONN hConn, MQCHAR *qName)
{
    MQLONG openReason;          /* MQOPEN reason code */
    MQLONG reason;             /* reason code */
    MQLONG compCode;           /* completion code */
    MQHOBJ eventQueue;         /* handle to event queue */

    MQHBAG eventBag = MQHB_UNUSABLE_HBAG; /* event bag to receive event msg */
    MQOD  od = {MQOD_DEFAULT};           /* Object Descriptor */
    MQMD  md = {MQMD_DEFAULT};           /* Message Descriptor */
    MQGMO gmo = {MQGMO_DEFAULT};         /* get message options */
    MQLONG bQueueOK = 1;                 /* keep reading msgs while true */

    /*****                                                        */
    /* Create an Event Bag in which to receive the event.         */
    /* Exit the function if the create fails.                       */
    /*****                                                        */

```

```

/*****/
mqCreateBag(MQ_CBO_USER_BAG, &eventBag, &compCode, &reason);
CheckCallResult("Create event bag", compCode, reason);
if (compCode != MQCC_OK)
    return;

/*****/
/* Open the event queue chosen by the user */
/*****/
strncpy(od.ObjectName, qName, (size_t)MQ_Q_NAME_LENGTH);
MQOPEN(hConn, &od, MQOO_INPUT_AS_Q_DEF+MQOO_FAIL_IF_QUIESCING, &eventQueue,
    &compCode, &openReason);
CheckCallResult("Open event queue", compCode, openReason);

/*****/
/* Set the GMO options to control the action of the get message from the */
/* queue. */
/*****/
gmo.WaitInterval = 30000; /* 30 second wait for message */
gmo.Options = MQGMO_WAIT + MQGMO_FAIL_IF_QUIESCING + MQGMO_CONVERT;
gmo.Version = MQGMO_VERSION_2; /* Avoid need to reset Message ID */
gmo.MatchOptions = MQMO_NONE; /* and Correlation ID after every */
/* mqGetBag */

/*****/
/* If open fails, we cannot access the queue and must stop the monitor. */
/*****/
if (compCode != MQCC_OK)
    bQueueOK = 0;

/*****/
/* Main loop to get an event message when it arrives */
/*****/
while (bQueueOK)
{
    printf("\nWaiting for an event\n");

    /*****/
    /* Get the message from the event queue and convert it into the event */
    /* bag. */
    /*****/
    mqGetBag(hConn, eventQueue, &md, &gmo, eventBag, &compCode, &reason);

    /*****/
    /* If get fails, we cannot access the queue and must stop the monitor. */
    /*****/
    if (compCode != MQCC_OK)
    {
        bQueueOK = 0;

        /*****/
        /* If get fails because no message available then we have timed out, */
        /* so report this, otherwise report an error. */
        /*****/
        if (reason == MQRC_NO_MSG_AVAILABLE)
        {
            printf("No more messages\n");
        }
        else
        {
            CheckCallResult("Get bag", compCode, reason);
        }
    }
}

```

```

    }
}

/*****
/* Event message read - Print the contents of the event bag */
*****/
else
{
    if ( PrintBag(eventBag) )
        printf("\nError found while printing bag contents\n");

    } /* end of msg found */
} /* end of main loop */
/*****
/* Close the event queue if successfully opened */
*****/
if (openReason == MQRC_NONE)
{
    MQCLOSE(hConn, &eventQueue, MQCO_NONE, &compCode, &reason);
    CheckCallResult("Close event queue", compCode, reason);
}

/*****
/* Delete the event bag if successfully created. */
*****/
if (eventBag != MQHB_UNUSABLE_HBAG)
{
    mqDeleteBag(&eventBag, &compCode, &reason);
    CheckCallResult("Delete the event bag", compCode, reason);
}

} /* end of GetQEvents */

/*****
/*
/* Function: PrintBag
/*
*****/
/*
/* Input Parameters: Bag Handle
/*
/*
/* Output Parameters: None
/*
/*
/* Returns:          Number of errors found
/*
/*
/* Logic: Calls PrintBagContents to display the contents of the bag.
/*
*****/

int PrintBag(MQHBAG dataBag)
{
    int errors;

    printf("\n");
    errors = PrintBagContents(dataBag, 0);
    printf("\n");

    return errors;
}

```

```

/*****
/*
/* Function: PrintBagContents
/*
/*
/*****
/*
/* Input Parameters:  Bag Handle
/*                    Indentation level of bag
/*
/*
/* Output Parameters: None
/*
/*
/* Returns:           Number of errors found
/*
/*
/* Logic: Count the number of items in the bag
/*      Obtain selector and item type for each item in the bag.
/*      Obtain the value of the item depending on item type and display the
/*      index of the item, the selector and the value.
/*      If the item is an embedded bag handle then call this function again
/*      to print the contents of the embedded bag increasing the
/*      indentation level.
/*
/*
/*****
int PrintBagContents(MQHBAG dataBag, int indent)
{
    /*****
    /* Definitions
    /*****
#define LENGTH 500                /* Max length of string to be read*/
#define INDENT 4                  /* Number of spaces to indent
/* embedded bag display
/*

    /*****
    /* Variables
    /*****
    MQLONG  itemCount;            /* Number of items in the bag
    MQLONG  itemType;            /* Type of the item
    int     i;                   /* Index of item in the bag
    MQCHAR  stringVal[LENGTH+1]; /* Value if item is a string
    MQBYTE  byteStringVal[LENGTH]; /* Value if item is a byte string
    MQLONG  stringLength;        /* Length of string value
    MQLONG  ccSID;               /* CCSID of string value
    MQINT32 iValue;              /* Value if item is an integer
    MQINT64 i64Value;           /* Value if item is a 64-bit
/* integer
/*
    MQLONG  selector;           /* Selector of item
    MQHBAG  bagHandle;          /* Value if item is a bag handle
    MQLONG  reason;             /* reason code
    MQLONG  compCode;           /* completion code
    MQLONG  trimLength;         /* Length of string to be trimmed
    int     errors = 0;         /* Count of errors found
    char    blanks[] = "      "; /* Blank string used to
/* indent display
/*

    /*****
    /* Count the number of items in the bag
    /*****
    mqCountItems(dataBag, MQSEL_ALL_SELECTORS, &itemCount, &compCode, &reason);

    if (compCode != MQCC_OK)

```

```

    errors++;
else
{
    printf("
    printf("
    printf("
}

/*****
/* If no errors found, display each item in the bag */
*****/
if (!errors)
{
    for (i = 0; i < itemCount; i++)
    {

        /*****
        /* First inquire the type of the item for each item in the bag */
        *****/
        mqInquireItemInfo(dataBag, /* Bag handle */
            MQSEL_ANY_SELECTOR, /* Item can have any selector*/
            i, /* Index position in the bag */
            &selector, /* Actual value of selector */
            /* returned by call */
            &itemType, /* Actual type of item */
            /* returned by call */
            &compCode, /* Completion code */
            &reason); /* Reason Code */

        if (compCode != MQCC_OK)
            errors++;

        switch(itemType)
        {
        case MQITEM_INTEGER:
            /*****
            /* Item is an integer. Find its value and display its index, */
            /* selector and value. */
            *****/
            mqInquireInteger(dataBag, /* Bag handle */
                MQSEL_ANY_SELECTOR, /* Allow any selector */
                i, /* Index position in the bag */
                &iValue, /* Returned integer value */
                &compCode, /* Completion code */
                &reason); /* Reason Code */

            if (compCode != MQCC_OK)
                errors++;
            else
                printf("%.4s %-2d %-4d (%d)\n",
                    indent, blanks, i, selector, iValue);
            break

        case MQITEM_INTEGER64:
            /*****
            /* Item is a 64-bit integer. Find its value and display its */
            /* index, selector and value. */
            *****/
            mqInquireInteger64(dataBag, /* Bag handle */
                MQSEL_ANY_SELECTOR, /* Allow any selector */

```

```

        i,          /* Index position in the bag */
        &i64Value,  /* Returned integer value */
        &compCode, /* Completion code */
        &reason);  /* Reason Code */

if (compCode != MQCC_OK)
    errors++;
else
    printf("%.*s %-2d %-4d (%"Int64"d)\n",
           indent, blanks, i, selector, i64Value);
break;

case MQITEM_STRING:
    /******
    /* Item is a string. Obtain the string in a buffer, prepare */
    /* the string for displaying and display the index, selector, */
    /* string and Character Set ID. */
    /******
    mqInquireString(dataBag, /* Bag handle */
                    MQSEL_ANY_SELECTOR, /* Allow any selector */
                    i, /* Index position in the bag */
                    LENGTH, /* Maximum length of buffer */
                    stringVal, /* Buffer to receive string */
                    &stringLength, /* Actual length of string */
                    &ccsid, /* Coded character set id */
                    &compCode, /* Completion code */
                    &reason); /* Reason Code */

    /******
    /* The call can return a warning if the string is too long for */
    /* the output buffer and has been truncated, so only check */
    /* explicitly for call failure. */
    /******
    if (compCode == MQCC_FAILED)
        errors++;
    else
    {
        /******
        /* Remove trailing blanks from the string and terminate with*/
        /* a null. First check that the string should not have been */
        /* longer than the maximum buffer size allowed. */
        /******
        if (stringLength > LENGTH)
            trimLength = LENGTH;
        else
            trimLength = stringLength;
        mqTrim(trimLength, stringVal, stringVal, &compCode, &reason);
        printf("%.*s %-2d %-4d '%s' %d\n",
               indent, blanks, i, selector, stringVal, ccsid);
    }
    break;

case MQITEM_BYTE_STRING:
    /******
    /* Item is a byte string. Obtain the byte string in a buffer, */
    /* prepare the byte string for displaying and display the */
    /* index, selector and string. */
    /******
    mqInquireByteString(dataBag, /* Bag handle */

```

```

        MQSEL_ANY_SELECTOR, /* Allow any selector */
        i,                 /* Index position in the bag */
        LENGTH,           /* Maximum length of buffer */
        byteStringVal,    /* Buffer to receive string */
        &stringLength,    /* Actual length of string */
        &compCode,        /* Completion code */
        &reason);        /* Reason Code

/*****/
/* The call can return a warning if the string is too long for */
/* the output buffer and has been truncated, so only check */
/* explicitly for call failure. */
/*****/
if (compCode == MQCC_FAILED)
    errors++;
else
{
    printf("%.4s %-2d %-4d X'",
           indent, blanks, i, selector);

    for (i = 0 ; i < stringLength ; i++)
        printf("

    printf("\n");
}
break;

case MQITEM_BAG:
/*****/
/* Item is an embedded bag handle, so call the PrintBagContents*/
/* function again to display the contents. */
/*****/
mqInquireBag(dataBag,      /* Bag handle */
             MQSEL_ANY_SELECTOR, /* Allow any selector */
             i,           /* Index position in the bag */
             &bagHandle,  /* Returned embedded bag hdl*/
             &compCode,   /* Completion code */
             &reason);    /* Reason Code */

if (compCode != MQCC_OK)
    errors++;
else
{
    printf("%.4s %-2d %-4d (%d)\n", indent, blanks, i,
           selector, bagHandle);
    if (selector == MQHA_BAG_HANDLE)
        printf("
    else
        printf("
        PrintBagContents(bagHandle, indent+INDENT);
}
break;

default:
    printf("
}
}
}
return errors;
}

```

Inquire channel objects (amqsaicl.c)

```
/* **** */
/*
/* Program name: AMQSAICL.C
/*
/*
/* Description: Sample C program to inquire channel objects
/*              using the WebSphere MQ Administration Interface (MQAI)
/*
/*
/* <N_OCO_COPYRIGHT>
/* Licensed Materials - Property of IBM
/*
/*
/* 63H9336
/* (c) Copyright IBM Corp. 2008, 2019 All Rights Reserved.
/*
/* US Government Users Restricted Rights - Use, duplication or
/* disclosure restricted by GSA ADP Schedule Contract with
/* IBM Corp.
/* <NOC_COPYRIGHT>
/* **** */
/*
/* Function:
/* AMQSAICL is a sample C program that demonstrates how to inquire
/* attributes of the local queue manager using the MQAI interface. In
/* particular, it inquires all channels and their types.
/*
/* - A PCF command is built from items placed into an MQAI administration
/* bag.
/* These are:-
/* - The generic channel name "*"
/* - The attributes to be inquired. In this sample we just want
/* name and type attributes
/*
/* - The mqExecute MQCMD_INQUIRE_CHANNEL call is executed.
/* The call generates the correct PCF structure.
/* The default options to the call are used so that the command is sent
/* to the SYSTEM.ADMIN.COMMAND.QUEUE.
/* The reply from the command server is placed on a temporary dynamic
/* queue.
/* The reply from the MQCMD_INQUIRE_CHANNEL is read from the
/* temporary queue and formatted into the response bag.
/*
/* - The completion code from the mqExecute call is checked and if there
/* is a failure from the command server, then the code returned by the
/* command server is retrieved from the system bag that has been
/* embedded in the response bag to the mqExecute call.
/*
/* Note: The command server must be running.
/*
/* **** */
/* AMQSAICL has 2 parameter - the queue manager name (optional)
/* - output file (optional) default varies
/* **** */
/* **** */
/* Includes
/* **** */
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
```

```

#include <ctype.h>
#if (MQAT_DEFAULT == MQAT_OS400)
#include <recio.h>
#endif

#include <cmqc.h> /* MQI */
#include <cmqcf.c.h> /* PCF */
#include <cmqbc.h> /* MQAI */
#include <cmqxc.h> /* MQCD */

/*****
/* Function prototypes */
/*****
void CheckCallResult(MQCHAR *, MQLONG , MQLONG);

/*****
/* DataTypes */
/*****
#if (MQAT_DEFAULT == MQAT_OS400)
typedef _RFILE OUTFILEHDL;
#else
typedef FILE OUTFILEHDL;
#endif

/*****
/* Constants */
/*****
#if (MQAT_DEFAULT == MQAT_OS400)
const struct
{
    char name[9];
} ChlTypeMap[9] =
{
    "*SDR", /* MQCHT_SENDER */
    "*SVR", /* MQCHT_SERVER */
    "*RCVR", /* MQCHT_RECEIVER */
    "*RQSTR", /* MQCHT_REQUESTER */
    "*ALL", /* MQCHT_ALL */
    "*CLTCN", /* MQCHT_CLNTCONN */
    "*SVRCONN", /* MQCHT_SVRCONN */
    "*CLUSRCVR", /* MQCHT_CLUSRCVR */
    "*CLUSSDR", /* MQCHT_CLUSSDR */
};
#else
const struct
{
    char name[9];
} ChlTypeMap[9] =
{
    "sdr", /* MQCHT_SENDER */
    "svr", /* MQCHT_SERVER */
    "rcvr", /* MQCHT_RECEIVER */
    "rqstr", /* MQCHT_REQUESTER */
    "all", /* MQCHT_ALL */
    "cltconn", /* MQCHT_CLNTCONN */
    "svrcn", /* MQCHT_SVRCONN */
    "clusrcvr", /* MQCHT_CLUSRCVR */
    "clusdr", /* MQCHT_CLUSSDR */
};
#endif

```

```

/*****/
/* Macros */
/*****/
#if (MQAT_DEFAULT == MQAT_OS400)
#define OUTFILE "QTEMP/AMQSAICL(AMQSAICL)"
#define OPENOUTFILE(hdl, fname) \
    (hdl) = _Ropen((fname),"wr, rtncode=Y");
#define CLOSEOUTFILE(hdl) \
    _Rclose((hdl));
#define WRITEOUTFILE(hdl, buf, buflen) \
    _Rwrite((hdl),(buf),(buflen));

#elif (MQAT_DEFAULT == MQAT_UNIX)
#define OUTFILE "/tmp/amqsaicl.txt"
#define OPENOUTFILE(hdl, fname) \
    (hdl) = fopen((fname),"w");
#define CLOSEOUTFILE(hdl) \
    fclose((hdl));
#define WRITEOUTFILE(hdl, buf, buflen) \
    fwrite((buf),(buflen),1,(hdl)); fflush((hdl));

#else
#define OUTFILE "amqsaicl.txt"
#define OPENOUTFILE(fname) \
    fopen((fname),"w");
#define CLOSEOUTFILE(hdl) \
    fclose((hdl));
#define WRITEOUTFILE(hdl, buf, buflen) \
    fwrite((buf),(buflen),1,(hdl)); fflush((hdl));

#endif

#define ChlType2String(t) ChlTypeMap[(t)-1].name

/*****/
/* Function: main */
/*****/
int main(int argc, char *argv[])
{
    /*****/
    /* MQAI variables */
    /*****/
    MQHCONN hConn; /* handle to MQ connection */
    MQCHAR qmName[MQ_Q_MGR_NAME_LENGTH+1]=""; /* default QMgr name */
    MQLONG reason; /* reason code */
    MQLONG connReason; /* MQCONN reason code */
    MQLONG compCode; /* completion code */
    MQHBAG adminBag = MQHB_UNUSABLE_HBAG; /* admin bag for mqExecute */
    MQHBAG responseBag = MQHB_UNUSABLE_HBAG; /* response bag for mqExecute */
    MQHBAG cAttrsBag; /* bag containing chl attributes */
    MQHBAG errorBag; /* bag containing cmd server error */
    MQLONG mqExecuteCC; /* mqExecute completion code */
    MQLONG mqExecuteRC; /* mqExecute reason code */
    MQLONG chlNameLength; /* Actual length of chl name */
    MQLONG chlType; /* Channel type */
    MQLONG i; /* loop counter */
    MQLONG numberOfBags; /* number of bags in response bag */
    MQCHAR chlName[MQ_OBJECT_NAME_LENGTH+1]; /* name of chl extracted from bag */
    MQCHAR OutputBuffer[100]; /* output data buffer */

```

```

OUTFILEHDL *outfp = NULL;                /* output file handle */

/*****
/* Connect to the queue manager */
*****/
if (argc > 1)
    strncpy(qmName, argv[1], (size_t)MQ_Q_MGR_NAME_LENGTH);
MQCONN(qmName, &hConn;, &compCode;, &connReason;);

/*****
/* Report the reason and stop if the connection failed. */
*****/
if (compCode == MQCC_FAILED)
{
    CheckCallResult("Queue Manager connection", compCode, connReason);
    exit( (int)connReason);
}

/*****
/* Open the output file */
*****/
if (argc > 2)
{
    OPENOUTFILE(outfp, argv[2]);
}
else
{
    OPENOUTFILE(outfp, OUTFILE);
}

if(outfp == NULL)
{
    printf("Could not open output file.\n");
    goto MOD_EXIT;
}

/*****
/* Create an admin bag for the mqExecute call */
*****/
mqCreateBag(MQCBO_ADMIN_BAG, &adminBag;, &compCode;, &reason;);
CheckCallResult("Create admin bag", compCode, reason);

/*****
/* Create a response bag for the mqExecute call */
*****/
mqCreateBag(MQCBO_ADMIN_BAG, &responseBag;, &compCode;, &reason;);
CheckCallResult("Create response bag", compCode, reason);

/*****
/* Put the generic channel name into the admin bag */
*****/
mqAddString(adminBag, MQCACH_CHANNEL_NAME, MQBL_NULL_TERMINATED, "*",
            &compCode;, &reason;);
CheckCallResult("Add channel name", compCode, reason);

/*****
/* Put the channel type into the admin bag */
*****/
mqAddInteger(adminBag, MQIACH_CHANNEL_TYPE, MQCHT_ALL, &compCode;, &reason;);
CheckCallResult("Add channel type", compCode, reason);

```

```

/*****
/* Add an inquiry for various attributes */
/*****
mqAddInquiry(adminBag, MQIACH_CHANNEL_TYPE, &compCode;, &reason;);
CheckCallResult("Add inquiry", compCode, reason);

/*****
/* Send the command to find all the channel names and channel types. */
/* The mqExecute call creates the PCF structure required, sends it to */
/* the command server, and receives the reply from the command server into */
/* the response bag. The attributes are contained in system bags that are */
/* embedded in the response bag, one set of attributes per bag. */
/*****
mqExecute(hConn, /* MQ connection handle */
          MQCMD_INQUIRE_CHANNEL, /* Command to be executed */
          MQHB_NONE, /* No options bag */
          adminBag, /* Handle to bag containing commands */
          responseBag, /* Handle to bag to receive the response */
          MQHO_NONE, /* Put msg on SYSTEM.ADMIN.COMMAND.QUEUE */
          MQHO_NONE, /* Create a dynamic q for the response */
          &compCode;, /* Completion code from the mqexecute */
          &reason;); /* Reason code from mqexecute call */

/*****
/* Check the command server is started. If not exit. */
/*****
if (reason == MQRC_CMD_SERVER_NOT_AVAILABLE)
{
    printf("Please start the command server: <strmqcsv QMgrName=" ">\n");
    goto MOD_EXIT;
}

/*****
/* Check the result from mqExecute call. If successful find the channel */
/* types for all the channels. If failed find the error. */
/*****
if ( compCode == MQCC_OK ) /* Successful mqExecute */
{
    /*****
    /* Count the number of system bags embedded in the response bag from the */
    /* mqExecute call. The attributes for each channel are in separate bags. */
    /*****
    mqCountItems(responseBag, MQHA_BAG_HANDLE, &numberOfBags;,
                &compCode;, &reason;);
    CheckCallResult("Count number of bag handles", compCode, reason);

    for ( i=0; i<numberOfbags; i++)
    {
        /*****
        /* Get the next system bag handle out of the mqExecute response bag. */
        /* This bag contains the channel attributes */
        /*****
        mqInquireBag(responseBag, MQHA_BAG_HANDLE, i, &Attrsbag,
                    &compCode, &reason);
        CheckCallResult("Get the result bag handle", compCode, reason);

        /*****
        /* Get the channel name out of the channel attributes bag */
        /*****
        mqInquireString(cAttrsBag, MQCACH_CHANNEL_NAME, 0, MQ_OBJECT_NAME_LENGTH,

```

```

        ch1Name, &ch1NameLength, NULL, &compCode, &reason);
    CheckCallResult("Get channel name", compCode, reason);

    /******
    /* Get the channel type out of the channel attributes bag          */
    /******

mqInquireInteger(cAttrsBag, MQIACH_CHANNEL_TYPE, MQIND_NONE, &ch1Type,
                &compCode, &reason);
    CheckCallResult("Get type", compCode, reason);

    /******
    /* Use mqTrim to prepare the channel name for printing.          */
    /* Print the result.                                             */
    /******
mqTrim(MQ_CHANNEL_NAME_LENGTH, ch1Name, ch1Name, &compCode, &reason);
sprintf(OutputBuffer, "%-20s%-9s", ch1Name, Ch1Type2String(ch1Type));
WRITEOUTFILE(outfp,OutputBuffer,29)
}
}

else /* Failed mqExecute */
{
    printf("Call to get channel attributes failed: Cc = %ld : Rc = %ld\n",
           compCode, reason);
    /******
    /* If the command fails get the system bag handle out of the mqexecute */
    /* response bag.This bag contains the reason from the command server */
    /* why the command failed.                                           */
    /******
    if (reason == MQRCCF_COMMAND_FAILED)
    {
        mqInquireBag(responseBag, MQHA_BAG_HANDLE, 0, &errorBag,
                    &compCode, &reason);
        CheckCallResult("Get the result bag handle", compCode, reason);

        /******
        /* Get the completion code and reason code, returned by the command */
        /* server, from the embedded error bag.                             */
        /******
mqInquireInteger(errorBag, MQIASY_COMP_CODE, MQIND_NONE, &mqExecuteCC,
                &compCode, &reason );
        CheckCallResult("Get the completion code from the result bag",
                        compCode, reason);
mqInquireInteger(errorBag, MQIASY_REASON, MQIND_NONE, &mqExecuteRC,
                &compCode, &reason);
        CheckCallResult("Get the reason code from the result bag",
                        compCode, reason);
        printf("Error returned by the command server: Cc = %ld : Rc = %ld\n",
               mqExecuteCC, mqExecuteRC);
    }
}

MOD_EXIT:
/******
/* Delete the admin bag if successfully created.                      */
/******
if (adminBag != MQHB_UNUSABLE_HBAG)
{
    mqDeleteBag(&adminBag, &compCode, &reason);
}
}

```

```

    CheckCallResult("Delete the admin bag", compCode, reason);
}

/*****
/* Delete the response bag if successfully created.          */
*****/
if (responseBag != MQHB_UNUSABLE_HBAG)
{
    mqDeleteBag(&responseBag, &compCode, &reason);
    CheckCallResult("Delete the response bag", compCode, reason);
}

/*****
/* Disconnect from the queue manager if not already connected  */
*****/
if (connReason != MQRC_ALREADY_CONNECTED)
{
    MQDISC(&hConn, &compCode, &reason);
    CheckCallResult("Disconnect from Queue Manager", compCode, reason);
}

/*****
/* Close the output file if open                               */
*****/
if(outfp != NULL)
    CLOSEOUTFILE(outfp);

return 0;
}

/*****
/*
/* Function: CheckCallResult                                  */
/*
*****/
/*
/* Input Parameters:  Description of call                    */
/*                   Completion code                        */
/*                   Reason code                            */
/*
/* Output Parameters: None                                  */
/*
/* Logic: Display the description of the call, the completion code and the */
/*        reason code if the completion code is not successful */
/*
*****/
void CheckCallResult(char *callText, MQLONG cc, MQLONG rc)
{
    if (cc != MQCC_OK)
        printf("%s failed: Completion Code = %ld : Reason = %ld\n", callText,
            cc, rc);
}

```

Inquiring about queues and printing information (amqsailq.c)

```
*****  
/*                                                                    */  
/* Program name: AMQSAILQ.C                                          */  
/*                                                                    */  
/* Description:  Sample C program to inquire the current depth of the local */  
/*               queues using the WebSphere MQ Administration Interface (MQAI)*/  
/*                                                                    */  
/* Statement:    Licensed Materials - Property of IBM                */  
/*                                                                    */  
/*               84H2000, 5765-B73                                    */  
/*               84H2001, 5639-B42                                    */  
/*               84H2002, 5765-B74                                    */  
/*               84H2003, 5765-B75                                    */  
/*               84H2004, 5639-B43                                    */  
/*                                                                    */  
/*               (C) Copyright IBM Corp. 1999, 2019                 */  
/*                                                                    */  
*****  
/*                                                                    */  
/* Function:                                                       */  
/* AMQSAILQ is a sample C program that demonstrates how to inquire  */  
/* attributes of the local queue manager using the MQAI interface. In */  
/* particular, it inquires the current depths of all the local queues.*/  
/*                                                                    */  
/* - A PCF command is built by placing items into an MQAI administration */  
/* bag.                                                                    */  
/* These are:-                                                       */  
/*   - The generic queue name "*"                                       */  
/*   - The type of queue required. In this sample we want to          */  
/*     inquire local queues.                                             */  
/*   - The attribute to be inquired. In this sample we want the      */  
/*     current depths.                                                  */  
/*                                                                    */  
/* - The mqExecute call is executed with the command MQCMD_INQUIRE_Q. */  
/* The call generates the correct PCF structure.                       */  
/* The default options to the call are used so that the command is sent */  
/* to the SYSTEM.ADMIN.COMMAND.QUEUE.                                  */  
/* The reply from the command server is placed on a temporary dynamic  */  
/* queue.                                                                */  
/* The reply from the MQCMD_INQUIRE_Q command is read from the       */  
/* temporary queue and formatted into the response bag.                */  
/*                                                                    */  
/* - The completion code from the mqExecute call is checked and if there */  
/* is a failure from the command server, then the code returned by    */  
/* command server is retrieved from the system bag that has been      */  
/* embedded in the response bag to the mqExecute call.                 */  
/*                                                                    */  
/* - If the call is successful, the depth of each local queue is placed */  
/* in system bags embedded in the response bag of the mqExecute call. */  
/* The name and depth of each queue is obtained from each of the bags */  
/* and the result displayed on the screen.                              */  
/*                                                                    */  
/* Note: The command server must be running.                           */  
/*                                                                    */  
*****  
/*                                                                    */  
/* AMQSAILQ has 1 parameter - the queue manager name (optional)      */  
/*                                                                    */  
*****
```

```

/*****/
/* Includes */
/*****/
#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <ctype.h>

#include <cmqc.h> /* MQI */
#include <cmqfc.h> /* PCF */
#include <cmqbc.h> /* MQAI */

/*****/
/* Function prototypes */
/*****/
void CheckCallResult(MQCHAR *, MQLONG , MQLONG);

/*****/
/* Function: main */
/*****/
int main(int argc, char *argv[])
{
    /*****/
    /* MQAI variables */
    /*****/
    MQHCONN hConn; /* handle to WebSphere MQ connection */
    MQCHAR qmName[MQ_Q_MGR_NAME_LENGTH+1]=""; /* default QMgr name */
    MQLONG reason; /* reason code */
    MQLONG connReason; /* MQCONN reason code */
    MQLONG compCode; /* completion code */
    MQHBAG adminBag = MQHB_UNUSABLE_HBAG; /* admin bag for mqExecute */
    MQHBAG responseBag = MQHB_UNUSABLE_HBAG; /* response bag for mqExecute */
    MQHBAG qAttrsBag; /* bag containing q attributes */
    MQHBAG errorBag; /* bag containing cmd server error */
    MQLONG mqExecuteCC; /* mqExecute completion code */
    MQLONG mqExecuteRC; /* mqExecute reason code */
    MQLONG qNameLength; /* Actual length of q name */
    MQLONG qDepth; /* depth of queue */
    MQLONG i; /* loop counter */
    MQLONG numberOfBags; /* number of bags in response bag */
    MQCHAR qName[MQ_Q_NAME_LENGTH+1]; /* name of queue extracted from bag*/

    printf("Display current depths of local queues\n\n");

    /*****/
    /* Connect to the queue manager */
    /*****/
    if (argc > 1)
        strncpy(qmName, argv[1], (size_t)MQ_Q_MGR_NAME_LENGTH);
    MQCONN(qmName, &hConn, &compCode, &connReason);

    /*****/
    /* Report the reason and stop if the connection failed. */
    /*****/
    if (compCode == MQCC_FAILED)
    {
        CheckCallResult("Queue Manager connection", compCode, connReason);
        exit( (int)connReason);
    }
}

```

```

}

/*****
/* Create an admin bag for the mqExecute call */
/*****
mqCreateBag(MQCBO_ADMIN_BAG, &adminBag, &compCode, &reason);
CheckCallResult("Create admin bag", compCode, reason);
/*****
/* Create a response bag for the mqExecute call */
/*****
mqCreateBag(MQCBO_ADMIN_BAG, &responseBag, &compCode, &reason);
CheckCallResult("Create response bag", compCode, reason);

/*****
/* Put the generic queue name into the admin bag */
/*****
mqAddString(adminBag, MQCA_Q_NAME, MQBL_NULL_TERMINATED, "*",
            &compCode, &reason);
CheckCallResult("Add q name", compCode, reason);

/*****
/* Put the local queue type into the admin bag */
/*****
mqAddInteger(adminBag, MQIA_Q_TYPE, MQQT_LOCAL, &compCode, &reason);
CheckCallResult("Add q type", compCode, reason);

/*****
/* Add an inquiry for current queue depths */
/*****
mqAddInquiry(adminBag, MQIA_CURRENT_Q_DEPTH, &compCode, &reason);
CheckCallResult("Add inquiry", compCode, reason);

/*****
/* Send the command to find all the local queue names and queue depths. */
/* The mqExecute call creates the PCF structure required, sends it to */
/* the command server, and receives the reply from the command server into */
/* the response bag. The attributes are contained in system bags that are */
/* embedded in the response bag, one set of attributes per bag. */
/*****
mqExecute(hConn, /* WebSphere MQ connection handle */
          MQCMD_INQUIRE_Q, /* Command to be executed */
          MQHB_NONE, /* No options bag */
          adminBag, /* Handle to bag containing commands */
          responseBag, /* Handle to bag to receive the response */
          MQHO_NONE, /* Put msg on SYSTEM.ADMIN.COMMAND.QUEUE */
          MQHO_NONE, /* Create a dynamic q for the response */
          &compCode, /* Completion code from the mqExecute */
          &reason); /* Reason code from mqExecute call */

/*****
/* Check the command server is started. If not exit. */
/*****
if (reason == MQRC_CMD_SERVER_NOT_AVAILABLE)
{
    printf("Please start the command server: <strmqcsv QMgrName>\n");
    MQDISC(&hConn, &compCode, &reason);
    CheckCallResult("Disconnect from Queue Manager", compCode, reason);
    exit(98);
}

```

```

/*****
/* Check the result from mqExecute call. If successful find the current */
/* depths of all the local queues. If failed find the error.          */
/*****
if ( compCode == MQCC_OK )          /* Successful mqExecute */
{
    /*****
    /* Count the number of system bags embedded in the response bag from the */
    /* mqExecute call. The attributes for each queue are in a separate bag. */
    /*****
mqCountItems(responseBag, MQHA_BAG_HANDLE, &numberOfBags, &compCode,
              &reason);
    CheckCallResult("Count number of bag handles", compCode, reason);

    for ( i=0; i<numberOfBags; i++)
    {
        /*****
        /* Get the next system bag handle out of the mqExecute response bag. */
        /* This bag contains the queue attributes                          */
        /*****
mqInquireBag(responseBag, MQHA_BAG_HANDLE, i, &qAttrsBag, &compCode,
              &reason);
        CheckCallResult("Get the result bag handle", compCode, reason);

        /*****
        /* Get the queue name out of the queue attributes bag              */
        /*****
mqInquireString(qAttrsBag, MQCA_Q_NAME, 0, MQ_Q_NAME_LENGTH, qName,
                &qNameLength, NULL, &compCode, &reason);
        CheckCallResult("Get queue name", compCode, reason);

        /*****
        /* Get the depth out of the queue attributes bag                    */
        /*****
mqInquireInteger(qAttrsBag, MQIA_CURRENT_Q_DEPTH, MQIND_NONE, &qDepth,
                 &compCode, &reason);
        CheckCallResult("Get depth", compCode, reason);

        /*****
        /* Use mqTrim to prepare the queue name for printing.              */
        /* Print the result.                                                */
        /*****
mqTrim(MQ_Q_NAME_LENGTH, qName, qName, &compCode, &reason)
printf("%4d %-48s\n", qDepth, qName);
    }
}

else          /* Failed mqExecute */
{
    printf("Call to get queue attributes failed: Completion Code = %d :
           Reason = %d\n", compCode, reason);

    /*****
    /* If the command fails get the system bag handle out of the mqExecute */
    /* response bag. This bag contains the reason from the command server   */
    /* why the command failed.                                             */
    /*****
if (reason == MQRCCF_COMMAND_FAILED)
{

```

```

mqInquireBag(responseBag, MQHA_BAG_HANDLE, 0, &errorBag, &compCode,
              &reason);
CheckCallResult("Get the result bag handle", compCode, reason);

/*****
/* Get the completion code and reason code, returned by the command */
/* server, from the embedded error bag. */
*****/
mqInquireInteger(errorBag, MQIASY_COMP_CODE, MQIND_NONE, &mqExecuteCC,
                 &compCode, &reason );
CheckCallResult("Get the completion code from the result bag",
                compCode, reason);
mqInquireInteger(errorBag, MQIASY_REASON, MQIND_NONE, &mqExecuteRC,
                 &compCode, &reason);
CheckCallResult("Get the reason code from the result bag",
                compCode, reason);
printf("Error returned by the command server: Completion Code = %d :
        Reason = %d\n", mqExecuteCC, mqExecuteRC);
}
}

/*****
/* Delete the admin bag if successfully created. */
*****/
if (adminBag != MQHB_UNUSABLE_HBAG)
{
mqDeleteBag(&adminBag, &compCode, &reason);
CheckCallResult("Delete the admin bag", compCode, reason);
}

/*****
/* Delete the response bag if successfully created. */
*****/
if (responseBag != MQHB_UNUSABLE_HBAG)
{
mqDeleteBag(&responseBag, &compCode, &reason);
CheckCallResult("Delete the response bag", compCode, reason);
}

/*****
/* Disconnect from the queue manager if not already connected */
*****/
if (connReason != MQRC_ALREADY_CONNECTED)
{
MQDISC(&hConn, &compCode, &reason);
CheckCallResult("Disconnect from queue manager", compCode, reason);
}
return 0;
}

*****/
*
* Function: CheckCallResult
*
*****/
*
* Input Parameters: Description of call
*                   Completion code
*                   Reason code
*
*****/

```

```

* Output Parameters: None */
* */
* Logic: Display the description of the call, the completion code and the */
* reason code if the completion code is not successful */
* */
*****/
void CheckCallResult(char *callText, MQLONG cc, MQLONG rc)
{
    if (cc != MQCC_OK)
        printf("%s failed: Completion Code = %d : Reason = %d\n",
            callText, cc, rc);
}

```

Hints and tips for configuring WebSphere MQ

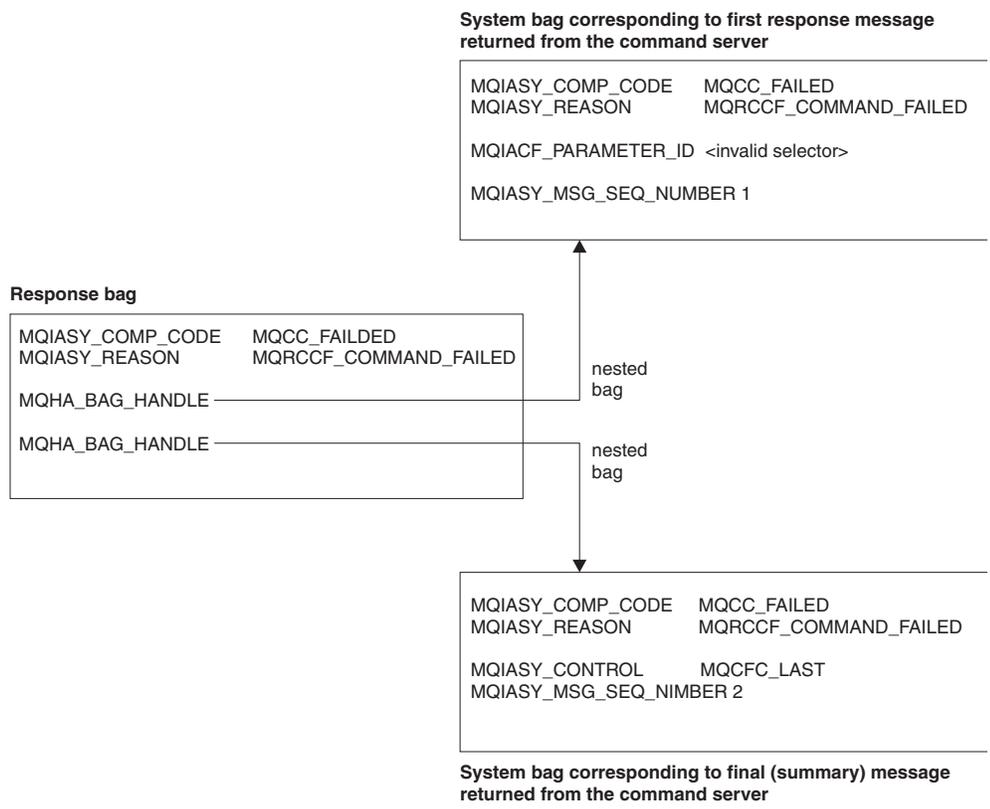
Programming hints and tips when using MQAI.

The MQAI uses PCF messages to send administration commands to the command server rather than dealing directly with the command server itself. Here are some tips for configuring WebSphere MQ using the MQAI:

- Character strings in WebSphere MQ are blank padded to a fixed length. Using C, null-terminated strings can normally be supplied as input parameters to WebSphere MQ programming interfaces.
- To clear the value of a string attribute, set it to a single blank rather than an empty string.
- Consider in advance the attributes that you want to change and inquire on just those attributes.
- Certain attributes cannot be changed, for example a queue name or a channel type. Ensure that you attempt to change only those attributes that can be modified. Refer to the list of required and optional parameters for the specific PCF change object. See  Definitions of the Programmable Command Formats (*WebSphere MQ V7.1 Reference*).
- If an MQAI call fails, some detail of the failure is returned to the response bag. Further detail can then be found in a nested bag that can be accessed by the selector MQHA_BAG_HANDLE. For example, if an mqExecute call fails with a reason code of MQRCCF_COMMAND_FAILED, this information is returned in the response bag. A possible reason for this reason code is that a selector specified was not valid for the type of command message and this detail of information is found in a nested bag that can be accessed by a bag handle.

For more information on MQExecute, see “Sending administration commands to the command server using the mqExecute call” on page 58

The following diagram shows this scenario:



Advanced topics

Information on indexing, data conversion and use of message descriptor

- Indexing
Indexes are used when replacing or removing existing data items from a bag to preserve insertion order. Full details on indexing can be found in Indexing (*WebSphere MQ V7.1 Reference*)
- Data conversion
The strings contained in an MQAI data bag can be in a variety of coded character sets and these can be converted using the mqSetInteger call. Full details on data conversion can be found in Data conversion (*WebSphere MQ V7.1 Reference*)
- Use of the message descriptor
MQAI generates a message descriptor which is set to an initial value when the data bag is created. Full details of the use of the message descriptor can be found in Use of the message descriptor (*WebSphere MQ V7.1 Reference*)

Data bags

A data bag is a means of handling properties or parameters of objects using the MQAI.

Data Bags

- The data bag contains zero or more *data items*. These data items are ordered within the bag as they are placed into the bag. This is called the *insertion order*. Each data item contains a *selector* that identifies the data item and a *value* of that data item that can be either an integer, a 64-bit integer, an integer filter, a string, a string filter, a byte string, a byte string filter, or a handle of another bag. Data items are described in details in “Data item” on page 51
- There are two types of selector; *user selectors* and *system selectors*. These are described in MQAI Selectors (*WebSphere MQ V7.1 Reference*). The selectors are usually unique, but it is possible to have

multiple values for the same selector. In this case, an *index* identifies the particular occurrence of selector that is required. Indexes are described in  Indexing (*WebSphere MQ V7.1 Reference*). A hierarchy of these concepts is shown in Figure 1. The hierarchy has been explained in a previous paragraph.

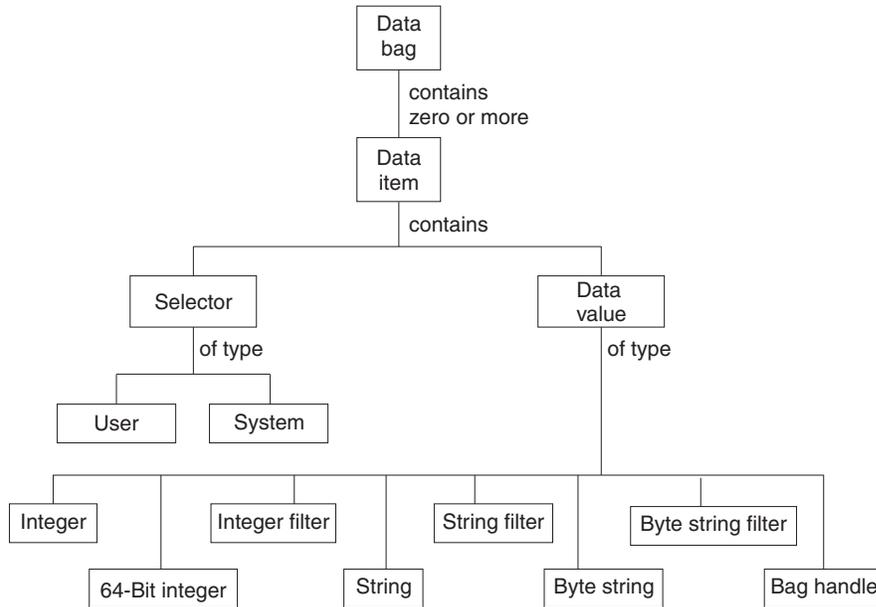


Figure 1. Hierarchy of MQAI concepts

Types of data bag

You can choose the type of data bag that you want to create depending on the task that you wish to perform:

user bag

A simple bag used for user data.

administration bag

A bag created for data used to administer WebSphere MQ objects by sending administration messages to a command server. The administration bag automatically implies certain options as described in “Creating and deleting data bags” on page 50.

command bag

A bag also created for commands for administering WebSphere MQ objects. However, unlike the administration bag, the command bag does not automatically imply certain options although these options are available. For more information about options, see “Creating and deleting data bags” on page 50.

group bag

A bag used to hold a set of grouped data items. Group bags cannot be used for administering WebSphere MQ objects.

In addition, the **system bag** is created by the MQAI when a reply message is returned from the command server and placed into a user's output bag. A system bag cannot be modified by the user.

Using Data Bags The different ways of using data bags are listed in this topic:

Using Data Bags

The different ways of using data bags are shown in the following list:

- You can create and delete data bags “Creating and deleting data bags.”
- You can send data between applications using data bags “Putting and receiving data bags” on page 51.
- You can add data items to data bags “Adding data items to bags” on page 52.
- You can add an inquiry command within a data bag “Adding an inquiry command to a bag” on page 53.
- You can inquire within data bags “Inquiring within data bags” on page 54.
- You can count data items within a data bag “Counting data items” on page 57.
- You can change information within a data bag “Changing information within a bag” on page 54.
- You can clear a data bag “Clearing a bag using the mqClearBag call” on page 55.
- You can truncate a data bag “Truncating a bag using the mqTruncateBag call” on page 55.
- You can convert bags and buffers “Converting bags and buffers” on page 56.

Creating and deleting data bags:

Creating data bags

To use the MQAI, you first create a data bag using the mqCreateBag call. As input to this call, you supply one or more options to control the creation of the bag.

The *Options* parameter of the MQCreateBag call lets you choose whether to create a user bag, a command bag, a group bag, or an administration bag.

To create a user bag, a command bag, or a group bag, you can choose one or more further options to:

- Use the list form when there are two or more adjacent occurrences of the same selector in a bag.
- Reorder the data items as they are added to a PCF message to ensure that the parameters are in their correct order. For more information on data items, see “Data item” on page 51.
- Check the values of user selectors for items that you add to the bag.

Administration bags automatically imply these options.

A data bag is identified by its handle. The bag handle is returned from mqCreateBag and must be supplied on all other calls that use the data bag.

For a full description of the mqCreateBag call, see  mqCreateBag (*WebSphere MQ V7.1 Reference*).

Deleting data bags

Any data bag that is created by the user must also be deleted using the mqDeleteBag call. For example, if a bag is created in the user code, it must also be deleted in the user code.

System bags are created and deleted automatically by the MQAI. For more information about this, see “Sending administration commands to the command server using the mqExecute call” on page 58. User code cannot delete a system bag.

For a full description of the mqDeleteBag call, see  mqDeleteBag (*WebSphere MQ V7.1 Reference*).

Putting and receiving data bags:

Data can also be sent between applications by putting and getting data bags using the mqPutBag and mqGetBag calls. This lets the MQAI handle the buffer rather than the application. The mqPutBag call converts the contents of the specified bag into a PCF message and sends the message to the specified queue and the mqGetBag call removes the message from the specified queue and converts it back into a data bag. Therefore, the mqPutBag call is the equivalent of the mqBagToBuffer call followed by MQPUT, and the mqGetBag is the equivalent of the MQGET call followed by mqBufferToBag.

For more information on sending and receiving PCF messages in a specific queue, see “Sending and receiving PCF messages in a specified queue” on page 9

Note: If you choose to use the mqGetBag call, the PCF details within the message must be correct; if they are not, an appropriate error results and the PCF message is not returned.

Data item:

Data items are used to populate Data bags when they are created. These data items can be user or system items.

These user items contain user data such as attributes of objects that are being administered. System items should be used for more control over the messages generated: for example, the generation of message headers. For more information about system items, see “System items” on page 52.

Types of Data Items

When you have created a data bag, you can populate it with integer or character-string items. You can inquire about all three types of item.

The data item can either be integer or character-string items. Here are the types of data item available within the MQAI:

- Integer
- 64-bit integer
- Integer filter
- Character-string
- String filter
- Byte string
- Byte string filter
- Bag handle

Using Data Items

These are the following ways of using data items:

- “Counting data items” on page 57.
- “Deleting data items” on page 57.
- “Adding data items to bags” on page 52.
- “Filtering and querying data items” on page 53.

System items:

System items can be used for:

- The generation of PCF headers. System items can control the PCF command identifier, control options, message sequence number, and command type.
- Data conversion. System items handle the character-set identifier for the character-string items in the bag.

Like all data items, system items consist of a selector and a value. For information about these selectors and what they are for, see  MQAI Selectors (*WebSphere MQ V7.1 Reference*).

System items are unique. One or more system items can be identified by a system selector. There is only one occurrence of each system selector.

Most system items can be modified (see “Changing information within a bag” on page 54), but the bag-creation options cannot be changed by the user. You cannot delete system items. (See “Deleting data items” on page 57.)

Adding data items to bags:

When a data bag is created, you can populate it with data items. These data items can be user or system items. Detailed information on data items can be found in “Data item” on page 51,

The MQAI lets you add integer items, 64-bit integer items, integer filter items, character-string items, string filter, byte string items, and byte string filter items to bags and this is shown in Figure 2. The items are identified by a selector. Usually one selector identifies one item only, but this is not always the case. If a data item with the specified selector is already present in the bag, an additional instance of that selector is added to the end of the bag.

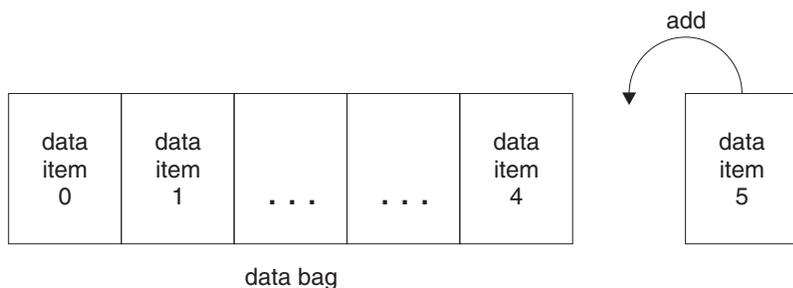


Figure 2. Adding data items

Add data items to a bag using the mqAdd* calls:

- To add integer items, use the mqAddInteger call as described in  mqAddInteger (*WebSphere MQ V7.1 Reference*)
- To add 64-bit integer items, use the mqAddInteger64 call as described in  mqAddInteger64 (*WebSphere MQ V7.1 Reference*)
- To add integer filter items, use the mqAddIntegerFilter call as described in  mqAddIntegerFilter (*WebSphere MQ V7.1 Reference*)
- To add character-string items, use the mqAddString call as described in  mqAddString (*WebSphere MQ V7.1 Reference*)

- To add string filter items, use the `mqAddStringFilter` call as described in  `mqAddStringFilter` (*WebSphere MQ V7.1 Reference*)
- To add byte string items, use the `mqAddByteString` call as described in  `mqAddByteString` (*WebSphere MQ V7.1 Reference*)
- To add byte string filter items, use the `mqAddByteStringFilter` call as described in  `mqAddByteStringFilter` (*WebSphere MQ V7.1 Reference*)

For more information on adding data items to a bag, see “System items” on page 52. .

Adding an inquiry command to a bag:

The `mqAddInquiry` call is used to add an inquiry command to a bag. The call is specifically for administration purposes, so it can be used with administration bags only. It lets you specify the selectors of attributes on which you want to inquire from WebSphere MQ.

For a full description of the `mqAddInquiry` call, see  `mqAddInquiry` (*WebSphere MQ V7.1 Reference*).

Filtering and querying data items:

When using the MQAI to inquire about the attributes of WebSphere MQ objects, you can control the data that is returned to your program in two ways.

- You can *filter* the data that is returned using the `mqAddInteger` and `mqAddString` calls. This approach lets you specify a *Selector* and *ItemValue* pair, for example:

```
mqAddInteger(inputbag, MQIA_Q_TYPE, MQQT_LOCAL)
```

This example specifies that the queue type (*Selector*) must be local (*ItemValue*) and this specification must match the attributes of the object (in this case, a queue) about which you are inquiring.

Other attributes that can be filtered correspond to the PCF Inquire* commands that can be found in “Introduction to Programmable Command Formats” on page 6. For example, to inquire about the attributes of a channel, see the Inquire Channel command in this product documentation. The “Required parameters” and “Optional parameters” of the Inquire Channel command identify the selectors that you can use for filtering.

- You can *query* particular attributes of an object using the `mqAddInquiry` call. This specifies the selector in which you are interested. If you do not specify the selector, all attributes of the object are returned.

Here is an example of filtering and querying the attributes of a queue:

```
/* Request information about all queues */
mqAddString(adminbag, MQCA_Q_NAME, "*")

/* Filter attributes so that local queues only are returned */
mqAddInteger(adminbag, MQIA_Q_TYPE, MQQT_LOCAL)

/* Query the names and current depths of the local queues */
mqAddInquiry(adminbag, MQCA_Q_NAME)
mqAddInquiry(adminbag, MQIA_CURRENT_Q_DEPTH)

/* Send inquiry to the command server and wait for reply */
mqExecute(MQCMD_INQUIRE_Q, ...)
```

For more examples of filtering and querying data items, see “Examples of using the MQAI” on page 19.

Inquiring within data bags:

You can inquire about:

- The value of an integer item using the `mqInquireInteger` call. See  `mqInquireInteger` (*WebSphere MQ V7.1 Reference*).
- The value of a 64-bit integer item using the `mqInquireInteger64` call. See  `mqInquireInteger64` (*WebSphere MQ V7.1 Reference*).
- The value of an integer filter item using the `mqInquireIntegerFilter` call. See  `mqInquireIntegerFilter` (*WebSphere MQ V7.1 Reference*).
- The value of a character-string item using the `mqInquireString` call. See  `mqInquireString` (*WebSphere MQ V7.1 Reference*).
- The value of a string filter item using the `mqInquireStringFilter` call. See  `mqInquireStringFilter` (*WebSphere MQ V7.1 Reference*).
- The value of a byte string item using the `mqInquireByteString` call. See  `mqInquireByteString` (*WebSphere MQ V7.1 Reference*).
- The value of a byte string filter item using the `mqInquireByteStringFilter` call. See  `mqInquireByteStringFilter` (*WebSphere MQ V7.1 Reference*).
- The value of a bag handle using the `mqInquireBag` call. See  `mqInquireBag` (*WebSphere MQ V7.1 Reference*).

You can also inquire about the type (integer, 64-bit integer, integer filter, character string, string filter, byte string, byte string filter or bag handle) of a specific item using the `mqInquireItemInfo` call. See

 `mqInquireItemInfo` (*WebSphere MQ V7.1 Reference*).

Changing information within a bag:

The MQAI lets you change information within a bag using the `mqSet*` calls. You can:

1. Modify data items within a bag. The index allows an individual instance of a parameter to be replaced by identifying the occurrence of the item to be modified (see Figure 3).

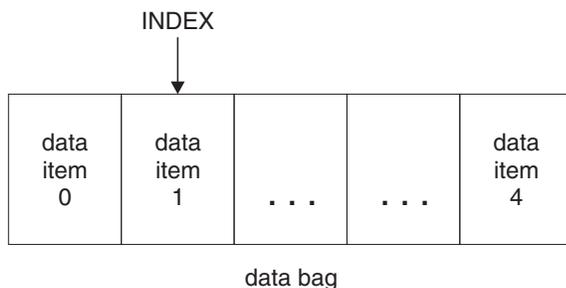


Figure 3. Modifying a single data item

2. Delete all existing occurrences of the specified selector and add a new occurrence to the end of the bag. (See Figure 4 on page 55.) A special index value allows *all* instances of a parameter to be replaced.

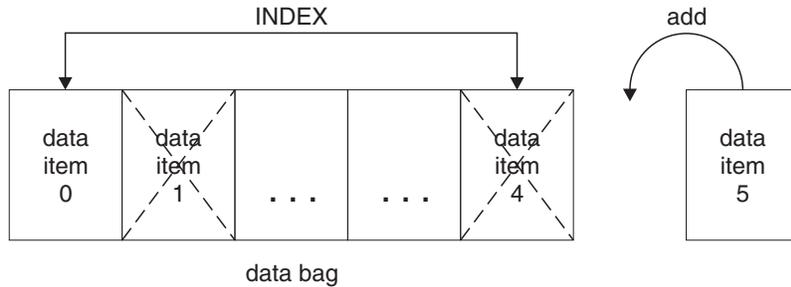


Figure 4. Modifying all data items

Note: The index preserves the insertion order within the bag but can affect the indices of other data items.

The `mqSetInteger` call lets you modify integer items within a bag. The `mqSetInteger64` call lets you modify 64-bit integer items. The `mqSetIntegerFilter` call lets you modify integer filter items. The `mqSetString` call lets you modify character-string items. The `mqSetStringFilter` call lets you modify string filter items. The `mqSetByteString` call lets you modify byte string items. The `mqSetByteStringFilter` call lets you modify byte string filter items. Alternatively, you can use these calls to delete all existing occurrences of the specified selector and add a new occurrence at the end of the bag. The data item can be a user item or a system item.

For a full description of these calls, see:

-  `mqSetInteger` (*WebSphere MQ V7.1 Reference*)
-  `mqSetInteger64` (*WebSphere MQ V7.1 Reference*)
-  `mqSetIntegerFilter` (*WebSphere MQ V7.1 Reference*)
-  `mqSetString` (*WebSphere MQ V7.1 Reference*)
-  `mqSetStringFilter` (*WebSphere MQ V7.1 Reference*)
-  `mqSetByteString` (*WebSphere MQ V7.1 Reference*)
-  `mqSetByteStringFilter` (*WebSphere MQ V7.1 Reference*)

Clearing a bag using the `mqClearBag` call:

The `mqClearBag` call removes all user items from a user bag and resets system items to their initial values. System bags contained within the bag are also deleted.

For a full description of the `mqClearBag` call, see  `mqClearBag` (*WebSphere MQ V7.1 Reference*).

Truncating a bag using the `mqTruncateBag` call:

The `mqTruncateBag` call reduces the number of user items in a user bag by deleting the items from the end of the bag, starting with the most recently added item. For example, it can be used when using the same header information to generate more than one message.

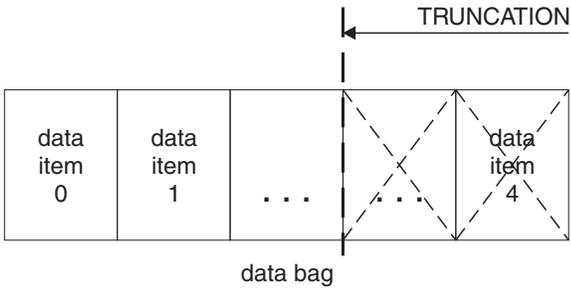


Figure 5. Truncating a bag

For a full description of the `mqTruncateBag` call, see [mqTruncateBag \(WebSphere MQ V7.1 Reference\)](#).

Converting bags and buffers:

To send data between applications, firstly the message data is placed in a bag. Then, the data in the bag is converted into a PCF message using the `mqBagToBuffer` call. The PCF message is sent to the required queue using the `MQPUT` call. This is shown in Figure Figure 6. For a full description of the `mqBagToBuffer` call, see [mqBagToBuffer \(WebSphere MQ V7.1 Reference\)](#).

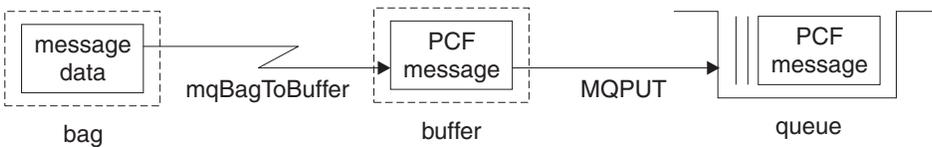


Figure 6. Converting bags to PCF messages

To receive data, the message is received into a buffer using the `MQGET` call. The data in the buffer is then converted into a bag using the `mqBufferToBag` call, providing the buffer contains a valid PCF message. This is shown in Figure Figure 7. For a full description of the `mqBufferToBag` call, see

[mqBufferToBag \(WebSphere MQ V7.1 Reference\)](#).

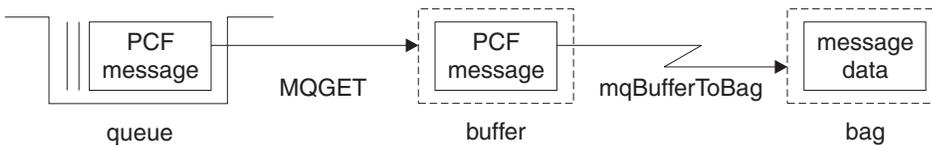


Figure 7. Converting PCF messages to bag form

Counting data items: The `mqCountItems` call counts the number of user items, system items, or both, that are stored in a data bag, and returns this number. For example, `mqCountItems(Bag, 7, ...)`, returns the number of items in the bag with a selector of 7. It can count items by individual selector, by user selectors, by system selectors, or by all selectors.

Note: This call counts the number of data items, not the number of unique selectors in the bag. A selector can occur multiple times, so there might be fewer unique selectors in the bag than data items.

For a full description of the `mqCountItems` call, see  `mqCountItems` (*WebSphere MQ V7.1 Reference*).

Deleting data items:

You can delete items from bags in a number of ways. You can:

- Remove one or more user items from a bag. For detailed information, see “Deleting data items from a bag using the `mqDeleteItem` call.”
- Delete *all* user items from a bag, that is, *clear* a bag. For detailed information see “Clearing a bag using the `mqClearBag` call” on page 55.
- Delete user items from the end of a bag, that is, *truncate* a bag. For detailed information, see “Truncating a bag using the `mqTruncateBag` call” on page 55.

Deleting data items from a bag using the `mqDeleteItem` call:

The `mqDeleteItem` call removes one or more user items from a bag. The index is used to delete either:

1. A single occurrence of the specified selector. (See Figure 8.)

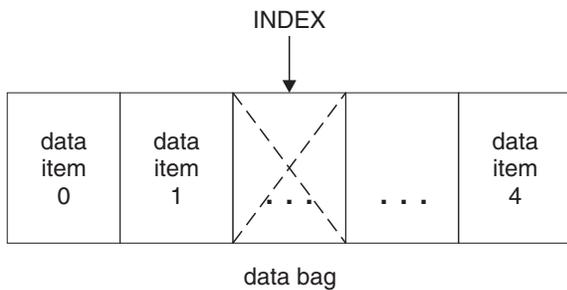


Figure 8. Deleting a single data item

or

2. All occurrences of the specified selector. (See Figure 9.)

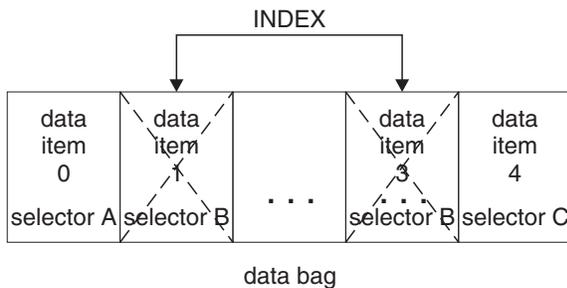


Figure 9. Deleting all data items

Note: The index preserves the insertion order within the bag but can affect the indices of other data items. For example, the `mqDeleteItem` call does not preserve the index values of the data items that follow the deleted item because the indices are reorganized to fill the gap that remains from the deleted item.

For a full description of the `mqDeleteItem` call, see  `mqDeleteItem` (*WebSphere MQ V7.1 Reference*).

Sending administration commands to the command server using the `mqExecute` call

When a data bag has been created and populated, an administrative command message can be sent to the command server of a queue manager using the `mqExecute` call. This handles the exchange with the command server and returns responses in a bag.

After you have created and populated your data bag, you can send an administration command message to the command server of a queue manager. The easiest way to do this is by using the `mqExecute` call. The `mqExecute` call sends an administration command message as a nonpersistent message and waits for any responses. Responses are returned in a response bag. These might contain information about attributes relating to several WebSphere MQ objects or a series of PCF error response messages, for example. Therefore, the response bag could contain a return code only or it could contain *nested bags*.

Response messages are placed into system bags that are created by the system. For example, for inquiries about the names of objects, a system bag is created to hold those object names and the bag is inserted into the user bag. Handles to these bags are then inserted into the response bag and the nested bag can be accessed by the selector `MQHA_BAG_HANDLE`. The system bag stays in storage, if it is not deleted, until the response bag is deleted.

The concept of *nesting* is shown in Figure 10.

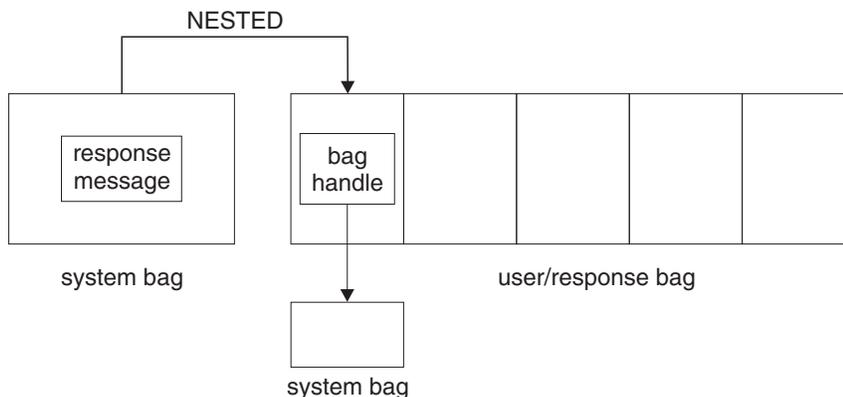


Figure 10. Nesting

As input to the `mqExecute` call, you must supply:

- An MQI connection handle.
- The command to be executed. This should be one of the `MQCMD_*` values.

Note: If this value is not recognized by the MQAI, the value is still accepted. However, if the `mqAddInquiry` call was used to insert values into the bag, this parameter must be an `INQUIRE` command recognized by the MQAI. That is, the parameter should be of the form `MQCMD_INQUIRE_*`.

- Optionally, a handle of the bag containing options that control the processing of the call. This is also where you can specify the maximum time in milliseconds that the MQAI should wait for each reply message.

- A handle of the administration bag that contains details of the administration command to be issued.
- A handle of the response bag that receives the reply messages.

The following are optional:

- An object handle of the queue where the administration command is to be placed.
If no object handle is specified, the administration command is placed on the SYSTEM.ADMIN.COMMAND.QUEUE belonging to the currently connected queue manager. This is the default.
- An object handle of the queue where reply messages are to be placed.
You can choose to place the reply messages on a dynamic queue that is created automatically by the MQAI. The queue created exists for the duration of the call only, and is deleted by the MQAI on exit from the mqExecute call.

For examples uses of the mqExecute call, see  [Example code \(WebSphere MQ V7.1 Reference\)](#)

Administration using the IBM WebSphere MQ Explorer

The IBM WebSphere MQ Explorer allows you to perform local or remote administration of your network from a computer running Windows, or Linux (x86 and x86-64 platforms) only.

IBM WebSphere MQ for Windows, and IBM WebSphere MQ for Linux (x86 and x86-64 platforms) provide an administration interface called the IBM WebSphere MQ Explorer to perform administration tasks as an alternative to using control or MQSC commands.  [Comparing command sets \(WebSphere MQ V7.1 Reference\)](#) shows you what you can do using the IBM WebSphere MQ Explorer.

The IBM WebSphere MQ Explorer allows you to perform local or remote administration of your network from a computer running Windows, or Linux (x86-64 platforms), by pointing the IBM WebSphere MQ Explorer at the queue managers and clusters you are interested in. The platforms and levels of IBM WebSphere MQ that can be administered using the IBM WebSphere MQ Explorer are described in “Remote queue managers” on page 61.

To configure remote IBM WebSphere MQ queue managers so that IBM WebSphere MQ Explorer can administer them, see “Prerequisite software and definitions” on page 62.

It allows you to perform tasks, typically associated with setting up and fine-tuning the working environment for IBM WebSphere MQ, either locally or remotely within a Windows or Linux (x86 and x86-64 platforms) system domain.

On Linux, the IBM WebSphere MQ Explorer might fail to start if you have more than one Eclipse installation. If this happens, start the IBM WebSphere MQ Explorer using a different user ID to the one you use for the other Eclipse installation.

On Linux, to start the IBM WebSphere MQ Explorer successfully, you must be able to write a file to your home directory, and the home directory must exist.

What you can do with the WebSphere MQ Explorer

This is a list of the tasks that you can perform using the WebSphere MQ Explorer.

With the WebSphere MQ Explorer, you can:

- Create and delete a queue manager (on your local machine only).
- Start and stop a queue manager (on your local machine only).
- Define, display, and alter the definitions of WebSphere MQ objects such as queues and channels.
- Browse the messages on a queue.
- Start and stop a channel.
- View status information about a channel, listener, queue, or service objects.
- View queue managers in a cluster.
- Check to see which applications, users, or channels have a particular queue open.
- Create a new queue manager cluster using the *Create New Cluster* wizard.
- Add a queue manager to a cluster using the *Add Queue Manager to Cluster* wizard.
- Manage the authentication information object, used with Secure Sockets Layer (SSL) channel security.
- Create and delete channel initiators, trigger monitors, and listeners.
- Start or stop the command servers, channel initiators, trigger monitors, and listeners.
- Set specific services to start automatically when a queue manager is started.
- Modify the properties of queue managers.
- Change the local default queue manager.
- Invoke the ikeyman GUI to manage secure sockets layer (SSL) certificates, associate certificates with queue managers, and configure and setup certificate stores (on your local machine only).
- Create JMS objects from WebSphere MQ objects, and WebSphere MQ objects from JMS objects.
- Create a JMS Connection Factory for any of the currently supported types.
- Modify the parameters for any service, such as the TCP port number for a listener, or a channel initiator queue name.
- Start or stop the service trace.

You perform administration tasks using a series of *Content Views* and *Property dialogs*.

Content View

A Content View is a panel that can display the following:

- Attributes, and administrative options relating to WebSphere MQ itself.
- Attributes, and administrative options relating to one or more related objects.
- Attributes, and administrative options for a cluster.

Property dialogs

A property dialog is a panel that displays attributes relating to an object in a series of fields, some of which you can edit.

You navigate through the WebSphere MQ Explorer using the *Navigator view*. The Navigator allows you to select the Content View you require.

Remote queue managers

There are two exceptions to the supported queue managers that you can connect to.

From a Windows or Linux (x86 and x86-64 platforms) system, the WebSphere MQ Explorer can connect to all supported queue managers with the following exceptions:

- WebSphere MQ for z/OS queue managers earlier than Version 6.0.
- Currently supported MQSeries® V2 queue managers.

The WebSphere MQ Explorer handles the differences in the capabilities between the different command levels and platforms. However, if it encounters an attribute that it does not recognize, the attribute will not be visible.

If you intend to remotely administer a V6.0 or later queue manager on Windows using the WebSphere MQ Explorer on a WebSphere MQ V5.3 computer, you must install Fix Pack 9 (CSD9) or later on your WebSphere MQ for Windows V5.3 computer.

If you intend to remotely administer a V5.3 queue manager on iSeries using the WebSphere MQ Explorer on a WebSphere MQ V6.0 or later computer, you must install Fix Pack 11 (CSD11) or later on your WebSphere MQ for iSeries V5.3 computer. This fix pack corrects connection problems between the WebSphere MQ Explorer and the iSeries queue manager.

Deciding whether to use the WebSphere MQ Explorer

When deciding whether to use the WebSphere MQ Explorer at your installation, consider the information listed in this topic.

You need to be aware of the following points:

Object names

If you use lowercase names for queue managers and other objects with the WebSphere MQ Explorer, when you work with the objects using MQSC commands, you must enclose the object names in single quotation marks, or WebSphere MQ does not recognize them.

Large queue managers

The WebSphere MQ Explorer works best with small queue managers. If you have a large number of objects on a single queue manager, you might experience delays while the WebSphere MQ Explorer extracts the required information to present in a view.

Clusters

WebSphere MQ clusters can potentially contain hundreds or thousands of queue managers. The WebSphere MQ Explorer presents the queue managers in a cluster using a tree structure. The physical size of a cluster does not affect the speed of the WebSphere MQ Explorer dramatically because the WebSphere MQ Explorer does not connect to the queue managers in the cluster until you select them.

Setting up the WebSphere MQ Explorer

This section outlines the steps you need to take to set up the WebSphere MQ Explorer.

- “Prerequisite software and definitions” on page 62
- “Security” on page 62
- “Showing and hiding queue managers and clusters” on page 66
- “Cluster membership” on page 67
- “Data conversion” on page 68

Prerequisite software and definitions

Ensure that you satisfy the following requirements before trying to use the IBM WebSphere MQ Explorer.

The IBM WebSphere MQ Explorer can connect to remote queue managers using the TCP/IP communication protocol only.

Check that:

1. A command server is running on every remotely administered queue manager.
2. A suitable TCP/IP listener object must be running on every remote queue manager. This object can be the IBM WebSphere MQ listener or, on UNIX and Linux systems, the `inetd` daemon.
3. A server-connection channel, by default named `SYSTEM.ADMIN.SVRCONN`, exists on all remote queue managers.

You can create the channel using the following MQSC command:

```
DEFINE CHANNEL(SYSTEM.ADMIN.SVRCONN) CHLTYPE(SVRCONN)
```

This command creates a basic channel definition. If you want a more sophisticated definition (to set up security, for example), you need additional parameters. For more information, see  `DEFINE CHANNEL` (*WebSphere MQ V7.1 Reference*).

The Client Attachment feature (CAF) is an option of IBM WebSphere MQ Explorer for z/OS that supports the attachment of clients to z/OS. If you require more than five attaches of the IBM WebSphere MQ Explorer you must install the Client Attachment feature (CAF). With the CAF installed, you can set the `MAXINST` attribute on the `SYSTEM.ADMIN.SVRCONN` channel from zero through 999 999 999.

If you do not have the Client Attachment feature (CAF) installed, the `MAXINST` attribute can be set from zero to five only on the `SYSTEM.ADMIN.SVRCONN` channel. A value greater than five is interpreted as zero without the CAF installed.

The Client Attachment Feature can be enabled for free by applying the PTF for PI13429.

4. The system queue, `SYSTEM.MQEXPLORER.REPLY.MODEL`, must exist.

Security

If you are using WebSphere MQ in an environment where it is important for you to control user access to particular objects, you might need to consider the security aspects of using the WebSphere MQ Explorer.

Authorization to use the WebSphere MQ Explorer:

Any user can use the WebSphere MQ Explorer, but certain authorities are required to connect, access, and manage queue managers.

To perform local administrative tasks using the WebSphere MQ Explorer, a user is required to have the necessary authority to perform the administrative tasks. If the user is a member of the `mqm` group, the user has authority to perform all local administrative tasks.

To connect to a remote queue manager and perform remote administrative tasks using the WebSphere MQ Explorer, the user executing the WebSphere MQ Explorer is required to have the following authorities:

- `CONNECT` authority on the target queue manager object
- `INQUIRE` authority on the target queue manager object
- `DISPLAY` authority to the target queue manager object
- `INQUIRE` authority to the queue, `SYSTEM.MQEXPLORER.REPLY.MODEL`
- `DISPLAY` authority to the queue, `SYSTEM.MQEXPLORER.REPLY.MODEL`
- `INPUT (get)` authority to the queue, `SYSTEM.MQEXPLORER.REPLY.MODEL`
- `OUTPUT (put)` authority to the queue, `SYSTEM.ADMIN.COMMAND.QUEUE`

- INQUIRE authority on the queue, SYSTEM.ADMIN.COMMAND.QUEUE
- Authority to perform the action selected

Note: INPUT authority relates to input to the user from a queue (a get operation). OUTPUT authority relates to output from the user to a queue (a put operation).

To connect to a remote queue manager on WebSphere MQ for z/OS and perform remote administrative tasks using the WebSphere MQ Explorer, the following must be provided:

- A RACF® profile for the system queue, SYSTEM.MQEXPLORER.REPLY.MODEL
- A RACF profile for the queues, AMQ.MQEXPLORER.*

In addition, the user executing the WebSphere MQ Explorer is required to have the following authorities:

- RACF UPDATE authority to the system queue, SYSTEM.MQEXPLORER.REPLY.MODEL
- RACF UPDATE authority to the queues, AMQ.MQEXPLORER.*
- CONNECT authority on the target queue manager object
- Authority to perform the action selected
- READ authority to all the hlq.DISPLAY.object profiles in the MQCMD5 class

For information about how to grant authority to WebSphere MQ objects, see “Giving access to a WebSphere MQ object on UNIX or Linux systems and Windows” on page 704.

If a user attempts to perform an operation that they are not authorized to perform, the target queue manager invokes authorization failure procedures and the operation fails.

The default filter in the WebSphere MQ Explorer is to display all WebSphere MQ objects. If there are any WebSphere MQ objects that a user does not have DISPLAY authority to, authorization failures are generated. If authority events are being recorded, restrict the range of objects that are displayed to those objects that the user has DISPLAY authority to.

Security for connecting to remote queue managers:

You must secure the channel between the IBM WebSphere MQ Explorer and each remote queue manager.

The IBM WebSphere MQ Explorer connects to remote queue managers as an MQI client application. This means that each remote queue manager must have a definition of a server-connection channel and a suitable TCP/IP listener. If you do not secure your server connection channel it is possible for a malicious application to connect to the same server connection channel and gain access to the queue manager objects with unlimited authority. In order to secure your server connection channel either specify a non-blank value for the MCAUSER attribute of the channel, use channel authentication records, or use a security exit.

The default value of the MCAUSER attribute is the local user ID. If you specify a non-blank user name as the MCAUSER attribute of the server connection channel, all programs connecting to the queue manager using this channel run with the identity of the named user and have the same level of authority. This does not happen if you use channel authentication records.

Related concepts:

“Channel authentication records” on page 408

Using a security exit with the WebSphere MQ Explorer:

You can specify a default security exit and queue manager specific security exits using the WebSphere MQ Explorer.

You can define a default security exit, which can be used for all new client connections from the WebSphere MQ Explorer. This default exit can be overridden at the time a connection is made. You can also define a security exit for a single queue manager or a set of queue managers, which takes effect when a connection is made. You specify exits using the WebSphere MQ Explorer. For more information, see the WebSphere MQ Help Center.

Using the IBM WebSphere MQ Explorer to connect to a remote queue manager using SSL-enabled MQI channels:

The IBM WebSphere MQ Explorer connects to remote queue managers using an MQI channel. If you want to secure the MQI channel using SSL security, you must establish the channel using a client channel definition table.

For information how to establish an MQI channel using a client channel definition table, see

 Overview of IBM WebSphere MQ MQI clients (*WebSphere MQ V7.1 Product Overview Guide*).

When you have established the channel using a client channel definition table, you can use the IBM WebSphere MQ Explorer to connect to a remote queue manager using SSL-enabled MQI channel, as described in “Tasks on the system that hosts the remote queue manager” and “Tasks on the system that hosts the IBM WebSphere MQ Explorer” on page 65.

Tasks on the system that hosts the remote queue manager

On the system hosting the remote queue manager, perform the following tasks:

1. Define a server connection and client connection pair of channels, and specify the appropriate value for the `SSLCIPH` variable on the server connection on both channels. For more information about the `SSLCIPH` variable, see “Protecting channels with SSL” on page 449
2. Send the channel definition table `AMQCLCHL.TAB`, which is found in the queue manager's `@ipcc` directory, to the system hosting the IBM WebSphere MQ Explorer.
3. Start a TCP/IP listener on a designated port.
4. Place both the CA and personal SSL certificates into the SSL directory of the queue manager:
 - `/var/mqm/qmgrs/+QMNAME+/SSL` for UNIX and Linux systems
 - `C:\Program Files\WebSphere MQ\qmgrs\+QMNAME+\SSL` for Windows systemsWhere `+QMNAME+` is a token representing the name of the queue manager.
5. Create a key database file of type CMS named `key.kdb`. Stash the password in a file either by checking the option in the iKeyman GUI, or by using the `-stash` option with the `runmqckm`, or `runmqakm` commands.
6. Add the CA certificates to the key database created in the previous step.
7. Import the personal certificate for the queue manager into the key database.

For more detailed information about working with the Secure Sockets Layer on Windows systems, see “Working with SSL or TLS on UNIX, Linux and Windows systems” on page 631.

Tasks on the system that hosts the IBM WebSphere MQ Explorer

On the system hosting the IBM WebSphere MQ Explorer, perform the following tasks:

1. Create a key database file of type JKS named `key.jks`. Set a password for this key database file.
The IBM WebSphere MQ Explorer uses Java™ keystore files (JKS) for SSL security, and so the keystore file being created for configuring SSL for the IBM WebSphere MQ Explorer must match this.
2. Add the CA certificates to the key database created in the previous step.
3. Import the personal certificate for the queue manager into the key database.
4. On Windows and Linux systems, start MQ Explorer by using the system menu, the `MQExplorer` executable file, or the `strmqcfcfg` command.
5. From the IBM WebSphere MQ Explorer toolbar, click **Window** -> **Preferences**, then expand **WebSphere MQ Explorer** and click **SSL Client Certificate Stores**. Enter the name of, and password for, the JKS file created in step 1 of “Tasks on the system that hosts the IBM WebSphere MQ Explorer,” in both the Trusted Certificate Store and the Personal Certificate Store, then click **OK**.
6. Close the **Preferences** window, and right-click **Queue Managers**. Click **Show/Hide Queue Managers**, and then click **Add** on the **Show/Hide Queue Managers** screen.
7. Type the name of the queue manager, and select the **Connect directly** option. Click next.
8. Select **Use client channel definition table (CCDT)** and specify the location of the channel table file that you transferred from the remote queue manager in step 2 in “Tasks on the system that hosts the remote queue manager” on page 64 on the system hosting the remote queue manager.
9. Click **Finish**. You can now access the remote queue manager from the IBM WebSphere MQ Explorer.

Connecting through another queue manager:

The WebSphere MQ Explorer allows you to connect to a queue manager through an intermediate queue manager, to which the WebSphere MQ Explorer is already connected.

In this case, the WebSphere MQ Explorer puts PCF command messages to the intermediate queue manager, specifying the following:

- The *ObjectQMgrName* parameter in the object descriptor (MQOD) as the name of the target queue manager. For more information on queue name resolution, see the  *Name resolution (WebSphere MQ V7.1 Programming Guide)*.
- The *UserIdentifier* parameter in the message descriptor (MQMD) as the local `userId`.

If the connection is then used to connect to the target queue manager via an intermediate queue manager, the `userId` is flowed in the *UserIdentifier* parameter of the message descriptor (MQMD) again. In order for the MCA listener on the target queue manager to accept this message, either the `MCAUSER` attribute must be set, or the `userId` must already exist with put authority.

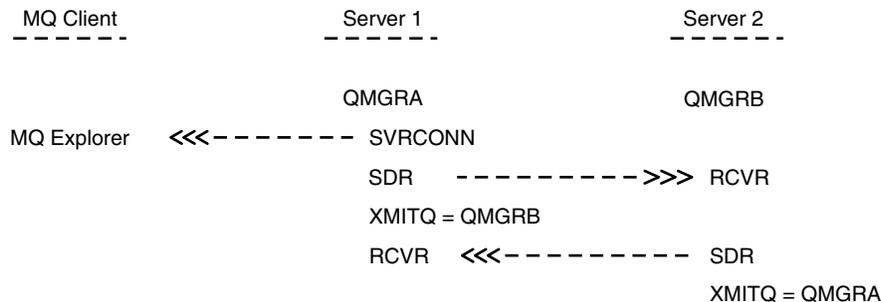
The command server on the target queue manager puts messages to the transmission queue specifying the `userId` in the *UserIdentifier* parameter in the message descriptor (MQMD). For this put to succeed the `userId` must already exist on the target queue manager with put authority.

The following example shows you how to connect a queue manager, through an intermediate queue manager, to the WebSphere MQ Explorer.

Establish a remote administration connection to a queue manager. Verify that the:

- Queue manager on the server is active and has a server-connection channel (SVRCONN) defined.
- Listener is active.
- Command server is active.

- SYSTEM.MQ EXPLORER.REPLY.MODEL queue has been created and that you have sufficient authority.
- Queue manager listeners, command servers, and sender channels are started.



In this example:

- WebSphere MQ Explorer is connected to queue manager QMGRA (running on Server1) using a client connection.
- Queue manager QMGRB on Server2 can be now connected to WebSphere MQ Explorer through an intermediate queue manager (QMGRA)
- When connecting to QMGRB with WebSphere MQ Explorer, select QMGRA as the intermediate queue manager

In this situation, there is no direct connection to QMGRB from WebSphere MQ Explorer; the connection to QMGRB is through QMGRA.

Queue manager QMGRB on Server2 is connected to QMGRA on Server1 using sender-receiver channels. The channel between QMGRA and QMGRB must be set up in such a way that remote administration is possible; see “Preparing channels and transmission queues for remote administration” on page 113.

Showing and hiding queue managers and clusters

The WebSphere MQ Explorer can display more than one queue manager at a time. From the Show/Hide Queue Manager panel (selectable from the menu for the Queue Managers tree node), you can choose whether you display information about another (remote) machine. Local queue managers are detected automatically.

To show a remote queue manager:

1. Right-click the **Queue Managers** tree node, then select *Show/Hide Queue Managers...*
2. Click **Add**. The Show/Hide Queue Managers panel is displayed.
3. Enter the name of the remote queue manager and the host name or IP address in the fields provided.
The host name or IP address is used to establish a client connection to the remote queue manager using either its default server connection channel, SYSTEM.ADMIN.SVRCONN, or a user-defined server connection channel.
4. Click **Finish**.

The Show/Hide Queue Managers panel also displays a list of all visible queue managers. You can use this panel to hide queue managers from the navigation view.

If the WebSphere MQ Explorer displays a queue manager that is a member of a cluster, the cluster is detected, and displayed automatically.

To export the list of remote queue managers from this panel:

1. Close the Show/Hide Queue Managers panel.
2. Right-click the top **IBM WebSphere MQ** tree node in the Navigation pane of the WebSphere MQ Explorer, then select **Export MQ Explorer Settings**
3. Click **MQ Explorer > MQ Explorer Settings**
4. Select **Connection Information > Remote queue managers**.
5. Select a file to store the exported settings in.
6. Finally, click **Finish** to export the remote queue manager connection information to the specified file.

To import a list of remote queue managers:

1. Right-click the top **IBM WebSphere MQ** tree node in the Navigation pane of the WebSphere MQ Explorer, then select **Import MQ Explorer Settings**
2. Click **MQ Explorer > MQ Explorer Settings**
3. Click **Browse**, and navigate to the path of the file that contains the remote queue manager connection information.
4. Click **Open**. If the file contains a list of remote queue managers, the **Connection Information > Remote queue managers** box is selected.
5. Finally, click **Finish** to import the remote queue manager connection information into the WebSphere MQ Explorer.

Cluster membership

WebSphere MQ Explorer requires information about queue managers that are members of a cluster.

If a queue manager is a member of a cluster, then the cluster tree node will be populated automatically.

If queue managers become members of clusters while the WebSphere MQ Explorer is running, then you must maintain the WebSphere MQ Explorer with up-to-date administration data about clusters so that it can communicate effectively with them and display correct cluster information when requested. In order to do this, the WebSphere MQ Explorer needs the following information:

- The name of a repository queue manager
- The connection name of the repository queue manager if it is on a remote queue manager

With this information, the WebSphere MQ Explorer can:

- Use the repository queue manager to obtain a list of queue managers in the cluster.
- Administer the queue managers that are members of the cluster and are on supported platforms and command levels.

Administration is not possible if:

- The chosen repository becomes unavailable. The WebSphere MQ Explorer does not automatically switch to an alternative repository.
- The chosen repository cannot be contacted over TCP/IP.
- The chosen repository is running on a queue manager that is running on a platform and command level not supported by the WebSphere MQ Explorer.

The cluster members that can be administered can be local, or they can be remote if they can be contacted using TCP/IP. The WebSphere MQ Explorer connects to local queue managers that are members of a cluster directly, without using a client connection.

Data conversion

The WebSphere MQ Explorer works in CCSID 1208 (UTF-8). This enables the WebSphere MQ Explorer to display the data from remote queue managers correctly. Whether connecting to a queue manager directly, or by using an intermediate queue manager, the WebSphere MQ Explorer requires all incoming messages to be converted to CCSID 1208 (UTF-8).

An error message is issued if you try to establish a connection between the WebSphere MQ Explorer and a queue manager with a CCSID that the WebSphere MQ Explorer does not recognize.

Supported conversions are described in  Code page conversion (*WebSphere MQ V7.1 Reference*).

Security on Windows

The Prepare WebSphere MQ wizard creates a special user account so that the Windows service can be shared by processes that need to use it.

A Windows service is shared between client processes for a IBM WebSphere MQ installation. One service is created for each installation. Each service is named `MQ_InstallationName`, and has a display name of `IBM WebSphere MQ(InstallationName)`. Before Version 7.1, with only one installation on a server the single, Windows service was named `MQSeriesServices` with the display name `IBM MQSeries`.

Because each service must be shared between non-interactive and interactive logon sessions, you must launch each under a special user account. You can use one special user account for all the services, or create different special user accounts. Each special user account must have the user right to “Logon as a service”, for more information see “User rights required for a IBM WebSphere MQ Windows Service” on page 69

When you install IBM WebSphere MQ and run the Prepare IBM WebSphere MQ wizard for the first time, it creates a local user account for the service called `MUSR_MQADMIN` with the required settings and permissions, including “Logon as a service”.

For subsequent installations, the Prepare IBM WebSphere MQ wizard creates a user account named `MUSR_MQADMINx`, where `x` is the next available number representing a user ID that does not exist. The password for `MUSR_MQADMINx` is randomly generated when the account is created, and used to configure the logon environment for the service. The generated password does not expire.

This IBM WebSphere MQ account is not affected by any account policies that are set up on the system to require that account passwords are changed after a certain period.

The password is not known outside this one-time processing and is stored by the Windows operating system in a secure part of the registry.

Related reference:

 Windows: `odqLogon as a service` required (*WebSphere MQ V7.1 Installing Guide*)

Using Active directory (Windows only)

In some network configurations, where user accounts are defined on domain controllers that are using Active Directory, the local user account IBM WebSphere MQ is running under might not have the authority it requires to query the group membership of other domain user accounts. The Prepare IBM WebSphere MQ Wizard identifies whether this is the case by carrying out tests and asking the user questions about the network configuration.

If the local user account IBM WebSphere MQ is running under does not have the required authority, the Prepare IBM WebSphere MQ Wizard prompts the user for the account details of a domain user account with particular user rights. For the user rights that the domain user account requires see “User rights required for a IBM WebSphere MQ Windows Service” on page 69. Once the user has entered valid

account details for the domain user account into the Prepare IBM WebSphere MQ Wizard, it configures a IBM WebSphere MQ Windows service to run under the new account. The account details are held in the secure part of the Registry and cannot be read by users.

When the service is running, a IBM WebSphere MQ Windows service is launched and remains running for as long as the service is running. A IBM WebSphere MQ administrator who logs on to the server after the Windows service is launched can use the IBM WebSphere MQ Explorer to administer queue managers on the server. This connects the IBM WebSphere MQ Explorer to the existing Windows service process. These two actions need different levels of permission before they can work:

- The launch process requires a launch permission.
- The IBM WebSphere MQ administrator requires Access permission.

User rights required for a IBM WebSphere MQ Windows Service

The table in this topic lists the user rights required for the local and domain user account under which the Windows service for a IBM WebSphere MQ installation runs.

Log on as batch job	Enables a IBM WebSphere MQ Windows service to run under this user account.
Log on as service	Enables users to set the IBM WebSphere MQ Windows service to log on using the configured account.
Shut down the system	Allows the IBM WebSphere MQ Windows service to restart the server if configured to do so when recovery of a service fails.
Increase quotas	Required for operating system CreateProcessAsUser call.
Act as part of the operating system	Required for operating system LogonUser call.
Bypass traverse checking	Required for operating system LogonUser call.
Replace a process level token	Required for operating system LogonUser call.

Note: Debug programs rights might be needed in environments running ASP and IIS applications. Your domain user account must have these Windows user rights set as effective user rights as listed in the Local Security Policy application. If they are not, set them using either the Local Security Policy application locally on the server, or by using the Domain Security Application domain wide.

Changing the user name associated with the IBM WebSphere MQ Service

You might need to change the user name associated with the IBM WebSphere MQ Service from MUSR_MQADMIN to something else. (For example, you might need to do this if your queue manager is associated with Db2, which does not accept user names of more than 8 characters.)

Procedure

1. Create a new user account (for example **NEW_NAME**)
2. Use the Prepare IBM WebSphere MQ Wizard to enter the account details of the new user account or use the Computer Management panel to alter the 'Log On' details for the installation specific IBM WebSphere MQ Service.

Changing the password of the IBM WebSphere MQ Windows service user account

About this task

To change the password of the IBM WebSphere MQ Windows service local user account, perform the following steps:

Procedure

1. Identify the user the service is running under.
2. Stop the IBM WebSphere MQ service from the Computer Management panel.
3. Change the required password in the same way that you would change the password of an individual.
4. Go to the properties for the IBM WebSphere MQ service from the Computer Management panel.
5. Select the **Log On Page**.
6. Confirm that the account name specified matches the user for which the password was modified.
7. Type the password into the **Password** and **Confirm password** fields and click **OK**.

IBM WebSphere MQ Windows service for an installation running under a domain user account:

About this task

If the IBM WebSphere MQ Windows service for an installation is running under a domain user account, you can also change the password for the account as follows:

Procedure

1. Change the password for the domain account on the domain controller. You might need to ask your domain administrator to do this for you.
2. Follow the steps to modify the **Log On** page for the IBM WebSphere MQ service.

The user account that IBM WebSphere MQ Windows service runs under executes any MQSC commands that are issued by user interface applications, or performed automatically on system startup, shutdown, or service recovery. This user account must therefore have IBM WebSphere MQ administration rights. By default it is added to the local **mqm** group on the server. If this membership is removed, the IBM WebSphere MQ Windows service does not work. For more information about user rights, see “User rights required for a IBM WebSphere MQ Windows Service” on page 69

If a security problem arises with the user account that the IBM WebSphere MQ Windows service runs under, error messages and descriptions appear in the system event log.

Related concepts:

“Using Active directory (Windows only)” on page 68

IBM WebSphere MQ coordinating with Db2 as the resource manager

If you start your queue managers from the IBM WebSphere MQ Explorer, or are using IBM WebSphere MQ V7, and are having problems when coordinating DB2®, check your queue manager error logs.

Check your queue manager error logs for an error like the following:

```
23/09/2008 15:43:54 - Process(5508.1) User(MUSR_MQADMIN) Program(amqzma0.exe)
Host(HOST_1) Installation(Installation1)
VMRF(7.1.0.0) QMgr(A.B.C)
AMQ7604: The XA resource manager 'DB2 MQBankDB database' was not available when
called for xa_open. The queue manager is continuing without this resource
manager.
```

Explanation: The user ID (default name is MUSR_MQADMIN) which runs the IBM WebSphere MQ Service process amqsvc.exe is still running with an access token which does not contain group membership information for the group DB2USERS.

Solve: After you have ensured that the IBM WebSphere MQ Service user ID is a member of DB2USERS, use the following sequence of commands:

- stop the service.
- stop any other processes running under the same user ID.
- restart these processes.

Rebooting the machine would ensure the previous steps, but is not necessary.

Extending the WebSphere MQ Explorer

WebSphere MQ for Windows, and WebSphere MQ for Linux (x86 and x86-64 platforms) provide an administration interface called the WebSphere MQ Explorer to perform administration tasks as an alternative to using control or MQSC commands.

This information applies to WebSphere MQ for Windows, and WebSphere MQ for Linux (x86 and x86-64 platforms) only.

The WebSphere MQ Explorer presents information in a style consistent with that of the Eclipse framework and the other plug-in applications that Eclipse supports.

Through extending the WebSphere MQ Explorer, system administrators have the ability to customize the WebSphere MQ Explorer to improve the way they administer WebSphere MQ.

For more information, see *Extending the WebSphere MQ Explorer* in the WebSphere MQ Explorer product documentation.

Using the WebSphere MQ Taskbar application (Windows only)

The WebSphere MQ Taskbar application displays an icon in the Windows system tray on the server. The icon provides you with the current status of WebSphere MQ and a menu from which you can perform some simple actions.

On Windows, the WebSphere MQ icon is in the system tray on the server and is overlaid with a color-coded status symbol, which can have one of the following meanings:

Green Working correctly; no alerts at present

Blue Indeterminate; WebSphere MQ is starting up or shutting down

Yellow

Alert; one or more services are failing or have already failed

To display the menu, right-click the WebSphere MQ icon. From the menu you can perform the following actions:

- Click **Open** to open the WebSphere MQ Alert Monitor
- Click **Exit** to exit the WebSphere MQ Taskbar application
- Click **WebSphere MQ Explorer** to start the WebSphere MQ Explorer
- Click **Stop WebSphere MQ** to stop WebSphere MQ
- Click **About WebSphere MQ** to display information about the WebSphere MQ Alert Monitor

The WebSphere MQ alert monitor application (Windows only)

The WebSphere MQ alert monitor is an error detection tool that identifies and records problems with WebSphere MQ on a local machine.

The alert monitor displays information about the current status of the local installation of a WebSphere MQ server. It also monitors the Windows Advanced Configuration and Power Interface (ACPI) and ensures the ACPI settings are enforced.

From the WebSphere MQ alert monitor, you can:

- Access the WebSphere MQ Explorer directly
- View information relating to all outstanding alerts
- Shut down the WebSphere MQ service on the local machine
- Route alert messages over the network to a configurable user account, or to a Windows workstation or server

Administering local WebSphere MQ objects

This section tells you how to administer local WebSphere MQ objects to support application programs that use the Message Queue Interface (MQI). In this context, local administration means creating, displaying, changing, copying, and deleting WebSphere MQ objects.

In addition to the approaches detailed in this section you can use the WebSphere MQ Explorer to administer local WebSphere MQ objects; see “Administration using the IBM WebSphere MQ Explorer” on page 59.

This section contains the following information:

-  Application programs using the MQI (*WebSphere MQ V7.1 Programming Guide*)
- “Performing local administration tasks using MQSC commands” on page 76
- “Working with queue managers” on page 85
- “Working with local queues” on page 87
- “Working with alias queues” on page 93
- “Working with model queues” on page 95
- “Working with services” on page 102
- “Managing objects for triggering” on page 108

Starting and stopping a queue manager

Use this topic as an introduction to stopping and starting a queue manager.

Starting a queue manager

To start a queue manager, use the **strmqm** command as follows:

```
strmqm saturn.queue.manager
```

On WebSphere MQ for Windows and WebSphere MQ for Linux (x86 and x86-64 platforms) systems, you can start a queue manager as follows:

1. Open the WebSphere MQ Explorer.
2. Select the queue manager from the Navigator View.
3. Click **Start**. The queue manager starts.

If the queue manager start-up takes more than a few seconds WebSphere MQ issues information messages intermittently detailing the start-up progress.

The **strmqm** command does not return control until the queue manager has started and is ready to accept connection requests.

Starting a queue manager automatically

In WebSphere MQ for Windows you can start a queue manager automatically when the system starts using the WebSphere MQ Explorer. For more information, see “Administration using the IBM WebSphere MQ Explorer” on page 59.

Stopping a queue manager

Use the **endmqm** command to stop a queue manager.

Note: You must use the **endmqm** command from the installation associated with the queue manager that you are working with. You can find out which installation a queue manager is associated with using the **dspmqs -o** installation command.

For example, to stop a queue manager called QMB, enter the following command:

```
endmqm QMB
```

On WebSphere MQ for Windows and WebSphere MQ for Linux (x86 and x86-64 platforms) systems, you can stop a queue manager as follows:

1. Open the WebSphere MQ Explorer.
2. Select the queue manager from the Navigator View.
3. Click **Stop...** The End Queue Manager panel is displayed.
4. Select Controlled, or Immediate.
5. Click **OK**. The queue manager stops.

Quiesced shutdown

By default, the **endmqm** command performs a quiesced shutdown of the specified queue manager. This might take a while to complete. A quiesced shutdown waits until all connected applications have disconnected.

Use this type of shutdown to notify applications to stop. If you issue:

```
endmqm -c QMB
```

you are not told when all applications have stopped. (An **endmqm -c QMB** command is equivalent to an **endmqm QMB** command.)

However, if you issue:

```
endmqm -w QMB
```

the command waits until all applications have stopped and the queue manager has ended.

Immediate shutdown

For an immediate shutdown any current MQI calls are allowed to complete, but any new calls fail. This type of shutdown does not wait for applications to disconnect from the queue manager.

For an immediate shutdown, type:

```
endmqm -i QMB
```

Preemptive shutdown

Note: Do not use this method unless all other attempts to stop the queue manager using the **endmqm** command have failed. This method can have unpredictable consequences for connected applications.

If an immediate shutdown does not work, you must resort to a *preemptive* shutdown, specifying the **-p** flag. For example:

```
endmqm -p QMB
```

This stops the queue manager immediately. If this method still does not work, see “Stopping a queue manager manually” for an alternative solution.

For a detailed description of the **endmqm** command and its options, see  **endmqm** (*WebSphere MQ V7.1 Reference*).

If you have problems shutting down a queue manager

Problems in shutting down a queue manager are often caused by applications. For example, when applications:

- Do not check MQI return codes properly
- Do not request notification of a quiesce
- Terminate without disconnecting from the queue manager (by issuing an **MQDISC** call)

If a problem occurs when you stop the queue manager, you can break out of the **endmqm** command using Ctrl-C. You can then issue another **endmqm** command, but this time with a flag that specifies the type of shutdown that you require.

Stopping a queue manager manually

If the standard methods for stopping queue managers fail, try the methods described here.

The standard way of stopping queue managers is by using the **endmqm** command. To stop a queue manager manually, use one of the procedures described in this section. For details of how to perform operations on queue managers using control commands, see  **Creating and managing queue managers**.

Stopping queue managers in WebSphere MQ for Windows

How to end the processes and the WebSphere MQ service, to stop queue managers in WebSphere MQ for Windows.

To stop a queue manager running under WebSphere MQ for Windows:

1. List the names (IDs) of the processes that are running, by using the Windows Task Manager.
2. End the processes by using Windows Task Manager, or the **taskkill** command, in the following order (if they are running):

AMQZMUC0	Critical process manager
AMQZXMA0	Execution controller
AMQZFUMA	OAM process
AMQZLAA0	LQM agents
AMQZLSA0	LQM agents
AMQZMUF0	Utility Manager
AMQZMGR0	Process controller
AMQZMUR0	Restartable process manager
AMQFQPUB	Publish Subscribe process
AMQFCXBA	Broker worker process
AMQRMPPA	Process pooling process
AMQCRSTA	Non-threaded responder job process
AMQCRS6B	LU62 receiver channel and client connection
AMQRRMFA	The repository process (for clusters)
AMQZDMAA	Deferred message processor
AMQPCSEA	The command server
RUNMQTRM	Invoke a trigger monitor for a server
RUNMQDLQ	Invoke dead-letter queue handler
RUNMQCHI	The channel initiator process
RUNMQLSR	The channel listener process
AMQXSSVN	Shared memory servers
AMQZTRCN	Trace

3. Stop the WebSphere MQ service from **Administration tools > Services** on the Windows Control Panel.
4. If you have tried all methods and the queue manager has not stopped, reboot your system.

The Windows Task Manager and the **tasklist** command give limited information about tasks. For more information to help to determine which processes relate to a particular queue manager, consider using a tool such as *Process Explorer* (procxp.exe), available for download from the Microsoft website at

 <http://www.microsoft.com>.

Stopping queue managers in WebSphere MQ for UNIX and Linux systems

How to end the processes and the WebSphere MQ service, to stop queue managers in WebSphere MQ for UNIX and Linux. You can try the methods described here if the standard methods for stopping and removing queue managers fail.

To stop a queue manager running under WebSphere MQ for UNIX and Linux systems:

1. Find the process IDs of the queue manager programs that are still running by using the **ps** command. For example, if the queue manager is called QMNAME, use the following command:

```
ps -ef | grep QMNAME
```
2. End any queue manager processes that are still running. Use the **kill** command, specifying the process IDs discovered by using the **ps** command.
 End the processes in the following order:

amqzmuc0	Critical process manager
amqzma0	Execution controller
amqzfuma	OAM process
amqzlaa0	LQM agents
amqzlsa0	LQM agents
amqzmuf0	Utility Manager
amqzmur0	Restartable process manager
amqzmgr0	Process controller
amqfqpub	Publish Subscribe process
amqfcxba	Broker worker process
amqrmppa	Process pooling process
amqcrsta	Non-threaded responder job process
amqcrs6b	LU62 receiver channel and client connection
amqrrmfa	The repository process (for clusters)
amqzdmaa	Deferred message processor
amqpcsea	The command server
runmqtrm	Invoke a trigger monitor for a server
runmqdlq	Invoke dead-letter queue handler
runmqchi	The channel initiator process
runmqlsr	The channel listener process

Note: You can use the `kill -9` command to end processes that fail to stop.

If you stop the queue manager manually, FFSTs might be taken, and FDC files placed in `/var/mqm/errors`. Do not regard this as a defect in the queue manager.

The queue manager will restart normally, even after you have stopped it using this method.

Performing local administration tasks using MQSC commands

This section introduces you to MQSC commands and tells you how to use them for some common tasks.

If you use IBM WebSphere MQ for Windows or IBM WebSphere MQ for Linux (x86 and x86-64 platforms), you can also perform the operations described in this section using the IBM WebSphere MQ Explorer. See “Administration using the IBM WebSphere MQ Explorer” on page 59 for more information.

You can use MQSC commands to manage queue manager objects, including the queue manager itself, queues, process definitions, channels, client connection channels, listeners, services, namelists, clusters, and authentication information objects. This section deals with queue managers, queues, and process definitions; for information about administering channel, client connection channel, and listener objects, see  *Objects (WebSphere MQ V7.1 Product Overview Guide)*. For information about all the MQSC commands for managing queue manager objects, see “Script (MQSC) Commands” on page 77.

You issue MQSC commands to a queue manager using the `runmqsc` command. (For details of this command, see  *runmqsc (WebSphere MQ V7.1 Reference)*.) You can do this interactively, issuing commands from a keyboard, or you can redirect the standard input device (`stdin`) to run a sequence of commands from an ASCII text file. In both cases, the format of the commands is the same. (For information about running the commands from a text file, see “Running MQSC commands from text files” on page 81.)

You can run the `runmqsc` command in three ways, depending on the flags set on the command:

- Verify a command without running it, where the MQSC commands are verified on a local queue manager, but are not run.

- Run a command on a local queue manager, where the MQSC commands are run on a local queue manager.
- Run a command on a remote queue manager, where the MQSC commands are run on a remote queue manager.

You can also run the command followed by a question mark to display the syntax.

Object attributes specified in MQSC commands are shown in this section in uppercase (for example, RQMNAME), although they are not case-sensitive. MQSC command attribute names are limited to eight characters. MQSC commands are available on other platforms, including IBM i and z/OS.

MQSC commands are summarized in the collection of topics in the  Comparing command sets (*WebSphere MQ V7.1 Reference*) section. “Script (MQSC) Commands” contains a description of each MQSC command and its syntax.

Script (MQSC) Commands

MQSC commands provide a uniform method of issuing human-readable commands on WebSphere MQ platforms. For information about *programmable command format* (PCF) commands, see “Introduction to Programmable Command Formats” on page 6.

The general format of the commands is shown in  The MQSC commands (*WebSphere MQ V7.1 Reference*).

You should observe the following rules when using MQSC commands:

- Each command starts with a primary parameter (a verb), and this is followed by a secondary parameter (a noun). This is then followed by the name or generic name of the object (in parentheses) if there is one, which there is on most commands. Following that, parameters can usually occur in any order; if a parameter has a corresponding value, the value must occur directly after the parameter to which it relates.

Note: On z/OS, the secondary parameter does not have to be second.

- Keywords, parentheses, and values can be separated by any number of blanks and commas. A comma shown in the syntax diagrams can always be replaced by one or more blanks. There must be at least one blank immediately preceding each parameter (after the primary parameter) except on z/OS.
- Any number of blanks can occur at the beginning or end of the command, and between parameters, punctuation, and values. For example, the following command is valid:

```
ALTER QLOCAL ('Account' )          TRIGDPTH ( 1)
```

Blanks within a pair of quotation marks are significant.

- Additional commas can appear anywhere where blanks are allowed and are treated as if they were blanks (unless, of course, they are inside strings enclosed by quotation marks).
- Repeated parameters are not allowed. Repeating a parameter with its ‘NO’ version, as in REPLACE NOREPLACE, is also not allowed.
- Strings that contain blanks, lowercase characters or special characters other than:
 - Period (.)
 - Forward slash (/)
 - Underscore (_)
 - Percent sign (%)

must be enclosed in single quotation marks, unless they are:

- Issued from the WebSphere MQ for z/OS operations and control panels

- Generic values ending with an asterisk (on IBM i these must be enclosed in single quotation marks)
- A single asterisk (for example, TRACE(*)) (on IBM i these must be enclosed in single quotation marks)
- A range specification containing a colon (for example, CLASS(01:03))

If the string itself contains a single quotation mark, the single quotation mark is represented by two single quotation marks. Lowercase characters not contained within quotation marks are folded to uppercase.

- On platforms other than z/OS, a string containing no characters (that is, two single quotation marks with no space in between) is interpreted as a blank space enclosed in single quotation marks, that is, interpreted in the same way as (' '). The exception to this is if the attribute being used is one of the following:
 - TOPICSTR
 - SUB
 - USERDATA
 - SELECTOR

then two single quotation marks with no space are interpreted as a zero-length string.

On z/OS, if you want a blank space enclosed in single quotation marks, you must enter it as such (' '). A string containing no characters ('') is the same as entering ().

- In v7.0, any trailing blanks in those string attributes which are based on MQCHARV types, such as SELECTOR, sub user data, are treated as significant which means that 'abc ' does not equal 'abc'.
- A left parenthesis followed by a right parenthesis, with no significant information in between, for example
NAME ()

is not valid except where specifically noted.

- Keywords are not case sensitive – ALTER, alter, and ALTER are all acceptable. Anything that is not contained within quotation marks is folded to uppercase.
- Synonyms are defined for some parameters. For example, DEF is always a synonym for DEFINE, so DEF QLOCAL is valid. Synonyms are not, however, just minimum strings; DEFI is not a valid synonym for DEFINE.

Note: There is no synonym for the DELETE parameter. This is to avoid accidental deletion of objects when using DEF, the synonym for DEFINE.

For an overview of using MQSC commands for administering IBM WebSphere MQ, see “Performing local administration tasks using MQSC commands” on page 76.

MQSC commands use certain special characters to have certain meanings. For more information about these special characters and how to use them, see  Characters with special meanings (*WebSphere MQ V7.1 Reference*).

To find out how you can build scripts using MQSC commands, see  Building command scripts (*WebSphere MQ V7.1 Reference*).

For information about how to use MQSC commands on z/OS, see  Using commands in z/OS (*WebSphere MQ V7.1 Reference*).

For the full list of MQSC commands, see  The MQSC commands (*WebSphere MQ V7.1 Reference*).

Related reference:

 Building command scripts (*WebSphere MQ V7.1 Reference*)

WebSphere MQ object names

How to use object names in MQSC commands.

In examples, we use some long names for objects. This is to help you identify the type of object you are dealing with.

When you issue MQSC commands, you need specify only the local name of the queue. In our examples, we use queue names such as:

```
ORANGE.LOCAL.QUEUE
```

The LOCAL.QUEUE part of the name is to illustrate that this queue is a local queue. It is *not* required for the names of local queues in general.

We also use the name saturn.queue.manager as a queue manager name. The queue.manager part of the name is to illustrate that this object is a queue manager. It is *not* required for the names of queue managers in general.

Case-sensitivity in MQSC commands

MQSC commands, including their attributes, can be written in uppercase or lowercase. Object names in MQSC commands are folded to uppercase (that is, QUEUE and queue are not differentiated), unless the names are enclosed within single quotation marks. If quotation marks are not used, the object is

processed with a name in uppercase. See the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for more information.

The **runmqsc** command invocation, in common with all WebSphere MQ control commands, is case sensitive in some WebSphere MQ environments. See  Using control commands (*WebSphere MQ V7.1 Reference*) for more information.

Standard input and output

The *standard input device*, also referred to as `stdin`, is the device from which input to the system is taken. Typically this is the keyboard, but you can specify that input is to come from a serial port or a disk file, for example. The *standard output device*, also referred to as `stdout`, is the device to which output from the system is sent. Typically this is a display, but you can redirect output to a serial port or a file.

On operating-system commands and WebSphere MQ control commands, the `<` operator redirects input. If this operator is followed by a file name, input is taken from the file. Similarly, the `>` operator redirects output; if this operator is followed by a file name, output is directed to that file.

Using MQSC commands interactively

You can use MQSC commands interactively by using a command window or shell.

To use MQSC commands interactively, open a command window or shell and enter:

```
runmqsc
```

In this command, a queue manager name has not been specified, so the MQSC commands are processed by the default queue manager. If you want to use a different queue manager, specify the queue manager name on the **runmqsc** command. For example, to run MQSC commands on queue manager `jupiter.queue.manager`, use the command:

```
runmqsc jupiter.queue.manager
```

After this, all the MQSC commands you type in are processed by this queue manager, assuming that it is on the same node and is already running.

Now you can type in any MQSC commands, as required. For example, try this one:

```
DEFINE QLOCAL (ORANGE.LOCAL.QUEUE)
```

For commands that have too many parameters to fit on one line, use continuation characters to indicate that a command is continued on the following line:

- A minus sign (-) indicates that the command is to be continued from the start of the following line.
- A plus sign (+) indicates that the command is to be continued from the first nonblank character on the following line.

Command input terminates with the final character of a nonblank line that is not a continuation character. You can also terminate command input explicitly by entering a semicolon (;). (This is especially useful if you accidentally enter a continuation character at the end of the final line of command input.)

Feedback from MQSC commands

When you issue MQSC commands, the queue manager returns operator messages that confirm your actions or tell you about the errors you have made. For example:

```
AMQ8006: WebSphere MQ queue created.
```

This message confirms that a queue has been created.

```
AMQ8405: Syntax error detected at or near end of command segment below:-
```

```
AMQ8426: Valid MQSC commands are:
```

```
ALTER
CLEAR
DEFINE
DELETE
DISPLAY
END
PING
REFRESH
RESET
RESOLVE
RESUME
START
STOP
SUSPEND
4 : end
```

This message indicates that you have made a syntax error.

These messages are sent to the standard output device. If you have not entered the command correctly, refer to the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for the correct syntax.

Ending interactive input of MQSC commands

To stop working with MQSC commands, enter the END command.

Alternatively, you can use the EOF character for your operating system.

Running MQSC commands from text files

Running MQSC commands interactively is suitable for quick tests, but if you have very long commands, or are using a particular sequence of commands repeatedly, consider redirecting stdin from a text file.

“Standard input and output” on page 79 contains information about stdin and stdout. To redirect stdin from a text file, first create a text file containing the MQSC commands using your usual text editor. When you use the **runmqsc** command, use the redirection operators. For example, the following command runs a sequence of commands contained in the text file `myprog.in`:

```
runmqsc < myprog.in
```

Similarly, you can also redirect the output to a file. A file containing the MQSC commands for input is called an *MQSC command file*. The output file containing replies from the queue manager is called the *output file*.

To redirect both stdin and stdout on the **runmqsc** command, use this form of the command:

```
runmqsc < myprog.in > myprog.out
```

This command invokes the MQSC commands contained in the MQSC command file `myprog.in`. Because we have not specified a queue manager name, the MQSC commands run against the default queue manager. The output is sent to the text file `myprog.out`. Figure 11 shows an extract from the MQSC command file `myprog.in` and Figure 12 on page 82 shows the corresponding extract of the output in `myprog.out`.

To redirect stdin and stdout on the **runmqsc** command, for a queue manager (`saturn.queue.manager`) that is not the default, use this form of the command:

```
runmqsc saturn.queue.manager < myprog.in > myprog.out
```

MQSC command files

MQSC commands are written in human-readable form, that is, in ASCII text. Figure 11 is an extract from an MQSC command file showing an MQSC command (`DEFINE QLOCAL`) with its attributes. The

 [WebSphere MQ Script \(MQSC\) Command Reference \(WebSphere MQ V7.1 Reference\)](#) contains a description of each MQSC command and its syntax.

```
.
.
.
DEFINE QLOCAL(ORANGE.LOCAL.QUEUE) REPLACE +
  DESCR(' ') +
  PUT(ENABLED) +
  DEFPRTY(0) +
  DEFPSIST(NO) +
  GET(ENABLED) +
  MAXDEPTH(5000) +
  MAXMSGL(1024) +
  DEFSOPT(SHARED) +
  NOHARDENBO +
  USAGE(NORMAL) +
  NOTRIGGER;
.
.
.
```

Figure 11. Extract from an MQSC command file

For portability among WebSphere MQ environments, limit the line length in MQSC command files to 72 characters. The plus sign indicates that the command is continued on the next line.

MQSC command reports

The `runmqsc` command returns a *report*, which is sent to stdout. The report contains:

- A header identifying MQSC commands as the source of the report:

```
Starting MQSC for queue manager jupiter.queue.manager.
```

Where `jupiter.queue.manager` is the name of the queue manager.

- An optional numbered listing of the MQSC commands issued. By default, the text of the input is echoed to the output. Within this output, each command is prefixed by a sequence number, as shown in Figure 12. However, you can use the `-e` flag on the `runmqsc` command to suppress the output.
- A syntax error message for any commands found to be in error.
- An *operator message* indicating the outcome of running each command. For example, the operator message for the successful completion of a `DEFINE QLOCAL` command is:

```
AMQ8006: WebSphere MQ queue created.
```
- Other messages resulting from general errors when running the script file.
- A brief statistical summary of the report indicating the number of commands read, the number of commands with syntax errors, and the number of commands that could not be processed.

Note: The queue manager attempts to process only those commands that have no syntax errors.

```
Starting MQSC for queue manager jupiter.queue.manager.
.
.
12:    DEFINE QLOCAL('ORANGE.LOCAL.QUEUE') REPLACE +
:      DESCR(' ') +
:      PUT(ENABLED) +
:      DEFPRTY(0) +
:      DEFPSIST(NO) +
:      GET(ENABLED) +
:      MAXDEPTH(5000) +
:      MAXMSGL(1024) +
:      DEFSOPT(SHARED) +
:      NOHARDENBO +
:      USAGE(NORMAL) +
:      NOTRIGGER;
AMQ8006: WebSphere MQ queue created.
:
.
.
```

Figure 12. Extract from an MQSC command report file

Running the supplied MQSC command files

The following MQSC command files are supplied with WebSphere MQ:

amqscos0.tst

Definitions of objects used by sample programs.

amqscic0.tst

Definitions of queues for CICS® transactions.

In WebSphere MQ for Windows, these files are located in the directory `MQ_INSTALLATION_PATH\tools\mqsc\samples`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

On UNIX and Linux systems these files are located in the directory `MQ_INSTALLATION_PATH/samp`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

The command that runs them is:

```
runmqsc < amqscos0.tst >test.out
```

Using runmqsc to verify commands

You can use the **runmqsc** command to verify MQSC commands on a local queue manager without actually running them. To do this, set the `-v` flag in the **runmqsc** command, for example:

```
runmqsc -v < myprog.in > myprog.out
```

When you invoke **runmqsc** against an MQSC command file, the queue manager verifies each command and returns a report without actually running the MQSC commands. This allows you to check the syntax of the commands in your command file. This is particularly important if you are:

- Running a large number of commands from a command file.
- Using an MQSC command file many times over.

The returned report is similar to that shown in Figure 12 on page 82.

You cannot use this method to verify MQSC commands remotely. For example, if you attempt this command:

```
runmqsc -w 30 -v jupiter.queue.manager < myprog.in > myprog.out
```

the `-w` flag, which you use to indicate that the queue manager is remote, is ignored, and the command is run locally in verification mode. 30 is the number of seconds that WebSphere MQ waits for replies from the remote queue manager.

Running MQSC commands from batch files

If you have very long commands, or are using a particular sequence of commands repeatedly, consider redirecting `stdin` from a batch file.

To redirect `stdin` from a batch file, first create a batch file containing the MQSC commands using your usual text editor. When you use the **runmqsc** command, use the redirection operators. The following example:

1. Creates a test queue manager, TESTQM
2. Creates a matching CLNTCONN and listener set to use TCP/IP port 1600
3. Creates a test queue, TESTQ
4. Puts a message on the queue, using the `amqsputc` sample program

```

export MYTEMPQM=TESTQM
export MYPOR=1600
export MQCHLLIB=/var/mqm/qmgrs/$MQTEMPQM/@ipcc

crtmqm $MYTEMPQM
strmqm $MYTEMPQM
runmqtsr -m $MYTEMPQM -t TCP -p $MYPOR &

runmqsc $MYTEMPQM << EOF
  DEFINE CHANNEL(NTLM) CHLTYPE(SVRCONN) TRPTYPE(TCP)
  DEFINE CHANNEL(NTLM) CHLTYPE(CLNTCONN) QMNAME('$MYTEMPQM') CONNAME('hostname($MYPOR)')
  ALTER CHANNEL(NTLM) CHLTYPE(CLNTCONN)
  DEFINE QLOCAL(TESTQ)
EOF

amqspcsc TESTQ $MYTEMPQM << EOF
hello world
EOF

endmqm -i $MYTEMPQM

```

Figure 13. Example script for running MQSC commands from a batch file

Resolving problems with MQSC commands

If you cannot get MQSC commands to run, use the information in this topic to see if any of these common problems apply to you. It is not always obvious what the problem is when you read the error that a command generates.

If you cannot get MQSC commands to run, use the following information to see if any of these common problems apply to you. It is not always obvious what the problem is when you read the error generated.

When you use the **runmqsc** command, remember the following:

- Use the < operator to redirect input from a file. If you omit this operator, the queue manager interprets the file name as a queue manager name, and issues the following error message:
AMQ8118: WebSphere MQ queue manager does not exist.
- If you redirect output to a file, use the > redirection operator. By default, the file is put in the current working directory at the time **runmqsc** is invoked. Specify a fully-qualified file name to send your output to a specific file and directory.
- Check that you have created the queue manager that is going to run the commands, by using the following command to display all queue managers:
dspmq
- The queue manager must be running. If it is not, start it; (see  Starting a queue manager (*WebSphere MQ V7.1 Installing Guide*)). You get an error message if you try to start a queue manager that is already running.
- Specify a queue manager name on the **runmqsc** command if you have not defined a default queue manager, or you get this error:
AMQ8146: WebSphere MQ queue manager not available.
- You cannot specify an MQSC command as a parameter of the **runmqsc** command. For example, this is not valid:
runmqsc DEFINE QLOCAL(FRED)
- You cannot enter MQSC commands before you issue the **runmqsc** command.
- You cannot run control commands from **runmqsc**. For example, you cannot issue the **strmqm** command to start a queue manager while you are running MQSC commands interactively. If you do this, you receive error messages similar to the following:

```
runmqsc
.
.
Starting MQSC for queue manager jupiter.queue.manager.

1 : strmqm saturn.queue.manager
AMQ8405: Syntax error detected at or near end of cmd segment below:-s

AMQ8426: Valid MQSC commands are:
ALTER
CLEAR
DEFINE
DELETE
DISPLAY
END
PING
REFRESH
RESET
RESOLVE
RESUME
START
STOP
SUSPEND
2 : end
```

Working with queue managers

Examples of MQSC commands that you can use to display or alter queue manager attributes.

Displaying queue manager attributes

To display the attributes of the queue manager specified on the **runmqsc** command, use the following MQSC command:

```
DISPLAY QMGR
```

Typical output from this command is shown in Figure 14 on page 86

```

DISPLAY QMGR
  1 : DISPLAY QMGR
AMQ8408: Display Queue Manager details.
QMNAME(QM1)
ACCTINT(1800)
ACCTQ(OFF)
ACTVCONO (DISABLED)
ALTDATE(2012-05-27)
AUTHOREV(DISABLED)
CHAD(DISABLED)
CHADEXIT( )
CLWLDATA( )
CLWLEN(100)
CLWLUSEQ(LOCAL)
CMDLEVEL(710)
CONFIGEV(DISABLED)
CRTIME(16.14.01)
DEFXMITQ( )
DISTL(YES)
IPADDRV(IPV4)
LOGGEREV(DISABLED)
MAXHANDS(256)
MAXPROPL(NOLIMIT)
MAXUMSGS(10000)
MONCHL(OFF)
PARENT( )
PLATFORM(WINDOWSNT)
PSNPMSG(DISCARD)
PSSYNCP(IFPER)
PSMODE(ENABLED)
REPOS( )
ROUTEREC(MSG)
SCMDSERV(QMGR)
SSLCRYP( )
SSLFIPS(NO)
MQ\Data\qmgrs\QM1\ssl\key)
SSLKEYC(0)
STATCHL(OFF)
STATMQI(OFF)
STRSTPEV(ENABLED)
TREELIFE(1800)
ACCTCONO(DISABLED)
ACCTMQI(OFF)
ACTIVREC(MSG)
ACTVTRC (OFF)
ALTTIME(16.14.01)
CCSID(850)
CHADEV(DISABLED)
CHLEV(DISABLED)
CLWLEXIT( )
CLWLMRUC(999999999)
CMDEV(DISABLED)
COMMANDQ(SYSTEM.ADMIN.COMMAND.QUEUE)
CRDATE(2011-05-27)
DEADQ()
DESCR( )
INHIBTEV(DISABLED)
LOCALEV(DISABLED)
MARKINT(5000)
MAXMSGL(4194304)
MAXPRTY(9)
MONACLS(QMGR)
MONQ(OFF)
PERFMEV(DISABLED)
PSRTYCNT(5)
PSNPRES(NORMAL)
QMID(QM1_2011-05-27_16.14.01)
REMOTEEV(DISABLED)
REPOSNL( )
SCHINIT(QMGR)
SSLCRLNL( )
SSLEV(DISABLED)
SSLKEYR(C:\Program Files\IBM\WebSphere
STATACLS(QMGR)
STATINT(1800)
STATQ(OFF)
SYNCPT
TRIGINT(999999999)

```

Figure 14. Typical output from a DISPLAY QMGR command

The ALL parameter is the default on the DISPLAY QMGR command. It displays all the queue manager attributes. In particular, the output tells you the default queue manager name, the dead-letter queue name, and the command queue name.

You can confirm that these queues exist by entering the command:

```
DISPLAY QUEUE (SYSTEM.*)
```

This displays a list of queues that match the stem SYSTEM.*. The parentheses are required.

Altering queue manager attributes

To alter the attributes of the queue manager specified on the **runmqsc** command, use the MQSC command ALTER QMGR, specifying the attributes and values that you want to change. For example, use the following commands to alter the attributes of jupiter.queue.manager:

```
runmqsc jupiter.queue.manager
ALTER QMGR DEADQ (ANOTHERDLQ) INHIBTEV (ENABLED)
```

The ALTER QMGR command changes the dead-letter queue used, and enables inhibit events.

Related reference:

 Attributes for the queue manager (*WebSphere MQ V7.1 Reference*)

Working with local queues

This section contains examples of some MQSC commands that you can use to manage local, model, and alias queues.

See the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for detailed information about these commands.

Defining a local queue

For an application, the local queue manager is the queue manager to which the application is connected. Queues managed by the local queue manager are said to be local to that queue manager.

Use the MQSC command `DEFINE QLOCAL` to create a local queue. You can also use the default defined in the default local queue definition, or you can modify the queue characteristics from those of the default local queue.

Note: The default local queue is named `SYSTEM.DEFAULT.LOCAL.QUEUE` and it was created on system installation.

For example, the `DEFINE QLOCAL` command that follows defines a queue called `ORANGE.LOCAL.QUEUE` with these characteristics:

- It is enabled for gets, enabled for puts, and operates on a priority order basis.
- It is an *normal* queue; it is not an initiation queue or transmission queue, and it does not generate trigger messages.
- The maximum queue depth is 5000 messages; the maximum message length is 4194304 bytes.

```
DEFINE QLOCAL (ORANGE.LOCAL.QUEUE) +
  DESCR('Queue for messages from other systems') +
  PUT (ENABLED) +
  GET (ENABLED) +
  NOTRIGGER +
  MSGDLVSQ (PRIORITY) +
  MAXDEPTH (5000) +
  MAXMSGL (4194304) +
  USAGE (NORMAL);
```

Note:

1. With the exception of the value for the description, all the attribute values shown are the default values. We have shown them here for purposes of illustration. You can omit them if you are sure that the defaults are what you want or have not been changed. See also “Displaying default object attributes” on page 89.
2. `USAGE (NORMAL)` indicates that this queue is not a transmission queue.
3. If you already have a local queue on the same queue manager with the name `ORANGE.LOCAL.QUEUE`, this command fails. Use the `REPLACE` attribute if you want to overwrite the existing definition of a queue, but see also “Changing local queue attributes” on page 90.

Defining a dead-letter queue:

Each queue manager must have a local queue to be used as a dead-letter queue so that messages that cannot be delivered to their correct destination can be stored for later retrieval. You must tell the queue manager about the dead-letter queue.

To tell the queue manager about the dead-letter queue, specify a dead-letter queue name on the `crtmqm` command (`crtmqm -u DEAD.LETTER.QUEUE`, for example), or by using the `DEADQ` attribute on the `ALTER QMGR` command to specify one later. You must define the dead-letter queue before using it.

A sample dead-letter queue called `SYSTEM.DEAD.LETTER.QUEUE` is available with the product. This queue is automatically created when you create the queue manager. You can modify this definition if required, and rename it.

A dead-letter queue has no special requirements except that:

- It must be a local queue
- Its `MAXMSGL` (maximum message length) attribute must enable the queue to accommodate the largest messages that the queue manager has to handle **plus** the size of the dead-letter header (`MQDLH`)

WebSphere MQ provides a dead-letter queue handler that allows you to specify how messages found on a dead-letter queue are to be processed or removed. For further information, see  Handling undelivered messages with the WebSphere MQ dead-letter queue handler (*WebSphere MQ V7.1 Installing Guide*).

Resolving problems with undelivered messages:

Use the advice given here to help you to resolve problems when messages do not arrive successfully.

- **Scenario:** Messages do not arrive on a queue when you are expecting them.
- **Explanation:** Messages that cannot be delivered for some reason are placed on the dead-letter queue.
- **Solution:** You can check whether the queue contains any messages by issuing an `MQSC DISPLAY QUEUE` command.

If the queue contains messages, you can use the provided browse sample application (`amqsbcg`) to browse messages on the queue using the `MQGET` call. The sample application steps through all the messages on a named queue for a named queue manager, displaying both the message descriptor and the message context fields for all the messages on the named queue.

You must decide how to dispose of any messages found on the dead-letter queue, depending on the reasons for the messages being put on the queue. Problems might occur if you do not associate a dead-letter queue with each queue manager.

For more information about dead-letter queues and handling undelivered messages, see  Handling undelivered messages with the WebSphere MQ dead-letter queue handler (*WebSphere MQ V7.1 Installing Guide*).

Displaying default object attributes

You can use the DISPLAY QUEUE command to display attributes that were taken from the default object when a WebSphere MQ object was defined.

When you define a WebSphere MQ object, it takes any attributes that you do not specify from the default object. For example, when you define a local queue, the queue inherits any attributes that you omit in the definition from the default local queue, which is called SYSTEM.DEFAULT.LOCAL.QUEUE. To see exactly what these attributes are, use the following command:

```
DISPLAY QUEUE (SYSTEM.DEFAULT.LOCAL.QUEUE)
```

The syntax of this command is different from that of the corresponding DEFINE command. On the DISPLAY command you can give just the queue name, whereas on the DEFINE command you have to specify the type of the queue, that is, QLOCAL, QALIAS, QMODEL, or QREMOTE.

You can selectively display attributes by specifying them individually. For example:

```
DISPLAY QUEUE (ORANGE.LOCAL.QUEUE) +
    MAXDEPTH +
    MAXMSGL +
    CURDEPTH;
```

This command displays the three specified attributes as follows:

```
AMQ8409: Display Queue details.
    QUEUE (ORANGE.LOCAL.QUEUE)          TYPE (QLOCAL)
    CURDEPTH (0)                        MAXDEPTH (5000)
    MAXMSGL (4194304)
```

CURDEPTH is the current queue depth, that is, the number of messages on the queue. This is a useful attribute to display, because by monitoring the queue depth you can ensure that the queue does not become full.

Copying a local queue definition

You can copy a queue definition using the LIKE attribute on the DEFINE command.

For example:

```
DEFINE QLOCAL (MAGENTA.QUEUE) +
    LIKE (ORANGE.LOCAL.QUEUE)
```

This command creates a queue with the same attributes as our original queue ORANGE.LOCAL.QUEUE, rather than those of the system default local queue. Enter the name of the queue to be copied *exactly* as it was entered when you created the queue. If the name contains lower case characters, enclose the name in single quotation marks.

You can also use this form of the DEFINE command to copy a queue definition, but substitute one or more changes to the attributes of the original. For example:

```
DEFINE QLOCAL (THIRD.QUEUE) +
    LIKE (ORANGE.LOCAL.QUEUE) +
    MAXMSGL (1024);
```

This command copies the attributes of the queue ORANGE.LOCAL.QUEUE to the queue THIRD.QUEUE, but specifies that the maximum message length on the new queue is to be 1024 bytes, rather than 4194304.

Note:

1. When you use the LIKE attribute on a DEFINE command, you are copying the queue attributes only. You are not copying the messages on the queue.

2. If you define a local queue, without specifying LIKE, it is the same as DEFINE LIKE(SYSTEM.DEFAULT.LOCAL.QUEUE).

Changing local queue attributes

You can change queue attributes in two ways, using either the ALTER QLOCAL command or the DEFINE QLOCAL command with the REPLACE attribute.

In “Defining a local queue” on page 87, the queue called ORANGE.LOCAL.QUEUE was defined. Suppose, for example, that you want to decrease the maximum message length on this queue to 10,000 bytes.

- Using the ALTER command:

```
ALTER QLOCAL (ORANGE.LOCAL.QUEUE) MAXMSGL(10000)
```

This command changes a single attribute, that of the maximum message length; all the other attributes remain the same.

- Using the DEFINE command with the REPLACE option, for example:

```
DEFINE QLOCAL (ORANGE.LOCAL.QUEUE) MAXMSGL(10000) REPLACE
```

This command changes not only the maximum message length, but also all the other attributes, which are given their default values. The queue is now put enabled whereas previously it was put inhibited. Put enabled is the default, as specified by the queue SYSTEM.DEFAULT.LOCAL.QUEUE.

If you *decrease* the maximum message length on an existing queue, existing messages are not affected. Any new messages, however, must meet the new criteria.

Clearing a local queue

You can use the CLEAR command to clear a local queue.

To delete all the messages from a local queue called MAGENTA.QUEUE, use the following command:

```
CLEAR QLOCAL (MAGENTA.QUEUE)
```

Note: There is no prompt that enables you to change your mind; once you press the Enter key the messages are lost.

You cannot clear a queue if:

- There are uncommitted messages that have been put on the queue under sync point.
- An application currently has the queue open.

Deleting a local queue

You can use the MQSC command DELETE QLOCAL to delete a local queue.

A queue cannot be deleted if it has uncommitted messages on it. However, if the queue has one or more committed messages and no uncommitted messages, it can be deleted only if you specify the PURGE option. For example:

```
DELETE QLOCAL (PINK.QUEUE) PURGE
```

Specifying NOPURGE instead of PURGE ensures that the queue is not deleted if it contains any committed messages.

Browsing queues

WebSphere MQ provides a sample queue browser that you can use to look at the contents of the messages on a queue. The browser is supplied in both source and executable formats.

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

In WebSphere MQ for Windows, the file names and paths for the sample queue browser are as follows:

Source

MQ_INSTALLATION_PATH\tools\c\samples\

Executable

MQ_INSTALLATION_PATH\tools\c\samples\bin\amqsbcg.exe

In WebSphere MQ for UNIX and Linux, the file names and paths are as follows:

Source

MQ_INSTALLATION_PATH/samp/amqsbcg0.c

Executable

MQ_INSTALLATION_PATH/samp/bin/amqsbcg

The sample requires two input parameters, the queue name and the queue manager name. For example:
amqsbcg SYSTEM.ADMIN.QMGREVENT.tpp01 saturn.queue.manager

Typical results from this command are shown in Figure 15 on page 92.

Some utilities, such as tar, cannot cope with files greater than 2 GB. Before enabling large file support, check your operating system documentation for information on restrictions on utilities you use.

For information about planning the amount of storage you need for queues, visit the IBM WebSphere MQ website for platform-specific performance reports:

<http://www.ibm.com/software/integration/ts/mqseries/>

Working with alias queues

You can define an alias queue to refer indirectly to another queue or topic.

Attention: Distribution lists do not support the use of alias queues that point to topic objects. From Version 7.1.0, Fix Pack 8, if an alias queue points to a topic object in a distribution list, WebSphere MQ returns MQRC_ALIAS_BASE_Q_TYPE_ERROR.

The queue to which an alias queue refers can be any of the following:

- A local queue (see “Defining a local queue” on page 87).
- A local definition of a remote queue (see “Creating a local definition of a remote queue” on page 118).
- A topic.

An alias queue is not a real queue, but a definition that resolves to a real (or target) queue at run time. The alias queue definition specifies the target queue. When an application makes an **MQOPEN** call to an alias queue, the queue manager resolves the alias to the target queue name.

An alias queue cannot resolve to another locally defined alias queue. However, an alias queue can resolve to alias queues that are defined elsewhere in clusters of which the local queue manager is a member. See

 Name resolution (*WebSphere MQ V7.1 Programming Guide*) for further information.

Alias queues are useful for:

- Giving different applications different levels of access authorities to the target queue.
- Allowing different applications to work with the same queue in different ways. (Perhaps you want to assign different default priorities or different default persistence values.)
- Simplifying maintenance, migration, and workload balancing. (Perhaps you want to change the target queue name without having to change your application, which continues to use the alias.)

For example, assume that an application has been developed to put messages on a queue called MY.ALIAS.QUEUE. It specifies the name of this queue when it makes an **MQOPEN** request and, indirectly, if it puts a message on this queue. The application is not aware that the queue is an alias queue. For each MQI call using this alias, the queue manager resolves the real queue name, which could be either a local queue or a remote queue defined at this queue manager.

By changing the value of the TARGQ attribute, you can redirect MQI calls to another queue, possibly on another queue manager. This is useful for maintenance, migration, and load-balancing.

Defining an alias queue

The following command creates an alias queue:

```
DEFINE QALIAS (MY.ALIAS.QUEUE) TARGET (YELLOW.QUEUE)
```

This command redirects MQI calls that specify MY.ALIAS.QUEUE to the queue YELLOW.QUEUE. The command does not create the target queue; the MQI calls fail if the queue YELLOW.QUEUE does not exist at run time.

If you change the alias definition, you can redirect the MQI calls to another queue. For example:

```
ALTER QALIAS (MY.ALIAS.QUEUE) TARGET (MAGENTA.QUEUE)
```

This command redirects MQI calls to another queue, MAGENTA.QUEUE.

You can also use alias queues to make a single queue (the target queue) appear to have different attributes for different applications. You do this by defining two aliases, one for each application. Suppose there are two applications:

- Application ALPHA can put messages on YELLOW.QUEUE, but is not allowed to get messages from it.
- Application BETA can get messages from YELLOW.QUEUE, but is not allowed to put messages on it.

The following command defines an alias that is put enabled and get disabled for application ALPHA:

```
DEFINE QALIAS (ALPHAS.ALIAS.QUEUE) +  
    TARGET (YELLOW.QUEUE) +  
    PUT (ENABLED) +  
    GET (DISABLED)
```

The following command defines an alias that is put disabled and get enabled for application BETA:

```
DEFINE QALIAS (BETAS.ALIAS.QUEUE) +  
    TARGET (YELLOW.QUEUE) +  
    PUT (DISABLED) +  
    GET (ENABLED)
```

ALPHA uses the queue name ALPHAS.ALIAS.QUEUE in its MQI calls; BETA uses the queue name BETAS.ALIAS.QUEUE. They both access the same queue, but in different ways.

You can use the LIKE and REPLACE attributes when you define queue aliases, in the same way that you use these attributes with local queues.

Using other commands with alias queues

You can use the appropriate MQSC commands to display or alter alias queue attributes, or to delete the alias queue object. For example:

Use the following command to display the alias queue's attributes:

```
DISPLAY QUEUE (ALPHAS.ALIAS.QUEUE)
```

Use the following command to alter the base queue name, to which the alias resolves, where the force option forces the change even if the queue is open:

```
ALTER QALIAS (ALPHAS.ALIAS.QUEUE) TARGQ(ORANGE.LOCAL.QUEUE) FORCE
```

Use the following command to delete this queue alias:

```
DELETE QALIAS (ALPHAS.ALIAS.QUEUE)
```

You cannot delete an alias queue if an application currently has the queue open. See the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for more information about this and other alias queue commands.

Working with model queues

A queue manager creates a *dynamic queue* if it receives an MQI call from an application specifying a queue name that has been defined as a model queue. The name of the new dynamic queue is generated by the queue manager when the queue is created. A *model queue* is a template that specifies the attributes of any dynamic queues created from it. Model queues provide a convenient method for applications to create queues as required.

Defining a model queue

You define a model queue with a set of attributes in the same way that you define a local queue. Model queues and local queues have the same set of attributes, except that on model queues you can specify whether the dynamic queues created are temporary or permanent. (Permanent queues are maintained across queue manager restarts, temporary ones are not.) For example:

```
DEFINE QMODEL (GREEN.MODEL.QUEUE) +
    DESCR('Queue for messages from application X') +
    PUT (DISABLED) +
    GET (ENABLED) +
    NOTRIGGER +
    MSGDLVSQ (FIFO) +
    MAXDEPTH (1000) +
    MAXMSGL (2000) +
    USAGE (NORMAL) +
    DEFTYPE (PERMDYN)
```

This command creates a model queue definition. From the DEFTYPE attribute, you can see that the actual queues created from this template are permanent dynamic queues. Any attributes not specified are automatically copied from the SYSYSTEM.DEFAULT.MODEL.QUEUE default queue.

You can use the LIKE and REPLACE attributes when you define model queues, in the same way that you use them with local queues.

Using other commands with model queues

You can use the appropriate MQSC commands to display or alter a model queue's attributes, or to delete the model queue object. For example:

Use the following command to display the model queue's attributes:

```
DISPLAY QUEUE (GREEN.MODEL.QUEUE)
```

Use the following command to alter the model to enable puts on any dynamic queue created from this model:

```
ALTER QMODEL (BLUE.MODEL.QUEUE) PUT(ENABLED)
```

Use the following command to delete this model queue:

```
DELETE QMODEL (RED.MODEL.QUEUE)
```

Working with administrative topics

Use MQSC commands to manage administrative topics.

See  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for detailed information about these commands.

Related concepts:

 Administrative topic objects (*WebSphere MQ V7.1 Installing Guide*)

“Defining an administrative topic”

“Displaying administrative topic object attributes” on page 97

“Changing administrative topic attributes” on page 97

“Copying an administrative topic definition” on page 98

“Deleting an administrative topic definition” on page 98

Defining an administrative topic

Use the MQSC command **DEFINE TOPIC** to create an administrative topic. When defining an administrative topic you can optionally set each topic attribute.

Any attribute of the topic that is not explicitly set is inherited from the default administrative topic, SYSTEM.DEFAULT.TOPIC, that was created when the system installation was installed.

For example, the **DEFINE TOPIC** command that follows, defines a topic called **ORANGE.TOPIC** with these characteristics:

- Resolves to the topic string ORANGE. For information about how topic strings can be used, see  Combining topic strings (*WebSphere MQ V7.1 Installing Guide*).
- Any attribute that is set to ASPARENT uses the attribute as defined by the parent topic of this topic. This action is repeated up the topic tree as far as the root topic, SYSTEM.BASE.TOPIC is found. For more information about topic trees, see  Topic trees (*WebSphere MQ V7.1 Installing Guide*).

```
DEFINE TOPIC (ORANGE.TOPIC) +  
  TOPICSTR (ORANGE) +  
  DEFPRTY (ASPARENT) +  
  NPMSGDLV (ASPARENT)
```

Note:

- Except for the value of the topic string, all the attribute values shown are the default values. They are shown here only as an illustration. You can omit them if you are sure that the defaults are what you want or have not been changed. See also “Displaying administrative topic object attributes” on page 97.
- If you already have an administrative topic on the same queue manager with the name ORANGE.TOPIC, this command fails. Use the REPLACE attribute if you want to overwrite the existing definition of a topic, but see also “Changing administrative topic attributes” on page 97

Displaying administrative topic object attributes

Use the MQSC command **DISPLAY TOPIC** to display an administrative topic object.

To display all topics, use:

```
DISPLAY TOPIC(ORANGE.TOPIC)
```

You can selectively display attributes by specifying them individually. For example:

```
DISPLAY TOPIC(ORANGE.TOPIC) +  
    TOPICSTR +  
    DEFPRTY +  
    NPMSGDLV
```

This command displays the three specified attributes as follows:

```
AMQ8633: Display topic details.  
    TOPIC(ORANGE.TOPIC)                TYPE(LOCAL)  
    TOPICSTR(ORANGE)                   DEFPRTY(ASPARENT)  
    NPMSGDLV(ASPARENT)
```

To display the topic ASPARENT values as they are used at Runtime use  **DISPLAY TPSTATUS** (*WebSphere MQ V7.1 Reference*). For example, use:

```
DISPLAY TPSTATUS(ORANGE) DEFPRTY NPMSGDLV
```

The command displays the following details:

```
AMQ8754: Display topic status details.  
    TOPICSTR(ORANGE)                   DEFPRTY(0)  
    NPMSGDLV(ALLAVAIL)
```

When you define an administrative topic, it takes any attributes that you do not specify explicitly from the default administrative topic, which is called `SYSTEM.DEFAULT.TOPIC`. To see what these default attributes are, use the following command:

```
DISPLAY TOPIC (SYSTEM.DEFAULT.TOPIC)
```

Changing administrative topic attributes

You can change topic attributes in two ways, using either the **ALTER TOPIC** command or the **DEFINE TOPIC** command with the **REPLACE** attribute.

If, for example, you want to change the default priority of messages delivered to a topic called `ORANGE.TOPIC`, to be 5, use either of the following commands.

- Using the **ALTER** command:

```
ALTER TOPIC(ORANGE.TOPIC) DEFPRTY(5)
```

This command changes a single attribute, that of the default priority of message delivered to this topic to 5; all other attributes remain the same.

- Using the **DEFINE** command:

```
DEFINE TOPIC(ORANGE.TOPIC) DEFPRTY(5) REPLACE
```

This command changes the default priority of messages delivered to this topic. All the other attributes are given their default values.

If you alter the priority of messages sent to this topic, existing messages are not affected. Any new message, however, use the specified priority if not provided by the publishing application.

Copying an administrative topic definition

You can copy a topic definition using the **LIKE** attribute on the **DEFINE** command.

For example:

```
DEFINE TOPIC (MAGENTA.TOPIC) +  
    LIKE (ORANGE.TOPIC)
```

This command creates a topic, `MAGENTA.TOPIC`, with the same attributes as the original topic, `ORANGE.TOPIC`, rather than those of the system default administrative topic. Enter the name of the topic to be copied exactly as it was entered when you created the topic. If the name contains lowercase characters, enclose the name in single quotation marks.

You can also use this form of the **DEFINE** command to copy a topic definition, but make changes to the attributes of the original. For example:

```
DEFINE TOPIC(BLUE.TOPIC) +  
    TOPICSTR(BLUE) +  
    LIKE(ORANGE.TOPIC)
```

You can also copy the attributes of the topic `BLUE.TOPIC` to the topic `GREEN.TOPIC` and specify that when publications cannot be delivered to their correct subscriber queue they are not placed onto the dead-letter queue. For example:

```
DEFINE TOPIC(GREEN.TOPIC) +  
    TOPICSTR(GREEN) +  
    LIKE(BLUE.TOPIC) +  
    USEDQ(NO)
```

Deleting an administrative topic definition

You can use the MQSC command **DELETE TOPIC** to delete an administrative topic.

```
DELETE TOPIC(ORANGE.TOPIC)
```

Applications will no longer be able to open the topic for publication or make new subscriptions using the object name, `ORANGE.TOPIC`. Publishing applications that have the topic open are able to continue publishing the resolved topic string. Any subscriptions already made to this topic continue receiving publications after the topic has been deleted.

Applications that are not referencing this topic object but are using the resolved topic string that this topic object represented, 'ORANGE' in this example, continue to work. In this case they inherit the properties from a topic object higher in the topic tree. For more information about topic trees, see

 Topic trees (*WebSphere MQ V7.1 Installing Guide*).

Working with subscriptions

Use MQSC commands to manage subscriptions.

Subscriptions can be one of three types, defined in the **SUBTYPE** attribute:

ADMIN

Administratively defined by a user.

PROXY

An internally created subscription for routing publications between queue managers.

API

Created programmatically, for example, using the MQI MQSUB call.

See the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for detailed information about these commands.

Related concepts:

- “Defining an administrative subscription”
- “Displaying attributes of subscriptions”
- “Changing local subscription attributes” on page 100
- “Copying a local subscription definition” on page 101
- “Deleting a subscription” on page 101

Defining an administrative subscription

Use the MQSC command **DEFINE SUB** to create an administrative subscription. You can also use the default defined in the default local subscription definition. Or, you can modify the subscription characteristics from those of the default local subscription, `SYSTEM.DEFAULT.SUB` that was created when the system was installed.

For example, the **DEFINE SUB** command that follows defines a subscription called `ORANGE` with these characteristics:

- Durable subscription, meaning that it persists over queue manager restart, with unlimited expiry.
- Receive publications made on the `ORANGE` topic string, with the message priorities as set by the publishing applications.
- Publications delivered for this subscription are sent to the local queue `SUBQ`, this queue must be defined before the definition of the subscription.

```
DEFINE SUB (ORANGE) +
    TOPICSTR (ORANGE) +
    DESTCLAS (PROVIDED) +
    DEST (SUBQ) +
    EXPIRY (UNLIMITED) +
    PUBPRTY (AS PUB)
```

Note:

- The subscription and topic string name do not have to match.
- Except for the values of the description and topic string, all the attribute values shown are the default values. They are shown here only as an illustration. You can omit them if you are sure that the defaults are what you want or have not been changed. See also “Displaying attributes of subscriptions.”
- If you already have a local subscription on the same queue manager with the name `TEST`, this command fails. Use the **REPLACE** attribute if you want to overwrite the existing definition of a queue, but see also “Changing local subscription attributes” on page 100.
- If the queue `SUBQ` does not exist, this command fails.

Displaying attributes of subscriptions

You can use the **DISPLAY SUB** command to display configured attributes of any subscription known to the queue manager.

For example, use:

```
DISPLAY SUB (ORANGE)
```

You can selectively display attributes by specifying them individually. For example:

```
DISPLAY SUB (ORANGE) +
    SUBID +
    TOPICSTR +
    DURABLE
```

This command displays the three specified attributes as follows:

```
AMQ8096: WebSphere MQ subscription inquired.  
SUBID(414D51204141412020202020202020EE921E4E20002A03)  
SUB(ORANGE) TOPICSTR(ORANGE)  
DURABLE(YES)
```

TOPICSTR is the resolved topic string on which this subscriber is operating. When a subscription is defined to use a topic object the topic string from that object is used as a prefix to the topic string provided when making the subscription. SUBID is a unique identifier assigned by the queue manager when a subscription is created. This is a useful attribute to display because some subscription names might be long or in a different character sets for which it might become impractical.

An alternate method for displaying subscriptions is to use the SUBID:

```
DISPLAY SUB +  
SUBID(414D51204141412020202020202020EE921E4E20002A03) +  
TOPICSTR +  
DURABLE
```

This command gives the same output as before:

```
AMQ8096: WebSphere MQ subscription inquired.  
SUBID(414D51204141412020202020202020EE921E4E20002A03)  
SUB(ORANGE) TOPICSTR(ORANGE)  
DURABLE(YES)
```

Proxy subscriptions on a queue manager are not displayed by default. To display them specify a **SUBTYPE** of PROXY or ALL.

You can use the  **DISPLAY SBSTATUS** (*WebSphere MQ V7.1 Reference*) command to display the Runtime attributes. For example, use the command:

```
DISPLAY SBSTATUS(ORANGE) NUMMSGS
```

The following output is displayed:

```
AMQ8099: WebSphere MQ subscription status inquired.  
SUB(ORANGE)  
SUBID(414D51204141412020202020202020EE921E4E20002A03)  
NUMMSGS(0)
```

When you define an administrative subscription, it takes any attributes that you do not specify explicitly from the default subscription, which is called SYSTEM.DEFAULT.SUB. To see what these default attributes are, use the following command:

```
DISPLAY SUB (SYSTEM.DEFAULT.SUB)
```

Changing local subscription attributes

You can change subscription attributes in two ways, using either the **ALTER SUB** command or the **DEFINE SUB** command with the **REPLACE** attribute.

If, for example, you want to change the priority of messages delivered to a subscription called ORANGE to be 5, use either of the following commands:

- Using the ALTER command:

```
ALTER SUB(ORANGE) PUBPRTY(5)
```

This command changes a single attribute, that of the priority of messages delivered to this subscription to 5; all other attributes remain the same.

- Using the DEFINE command:

```
DEFINE SUB (ORANGE) PUBPRTY(5) REPLACE
```

This command changes not only the priority of messages delivered to this subscription, but all the other attributes which are given their default values.

If you alter the priority of messages sent to this subscription, existing messages are not affected. Any new messages, however, are of the specified priority.

Copying a local subscription definition

You can copy a subscription definition using the **LIKE** attribute on the **DEFINE** command.

For example:

```
DEFINE SUB (BLUE) +  
    LIKE (ORANGE)
```

You can also copy the attributes of the sub **REAL** to the sub **THIRD.SUB**, and specify that the **correlID** of delivered publications is **THIRD**, rather than the publishers **correlID**. For example:

```
DEFINE SUB(THIRD.SUB) +  
    LIKE(BLUE) +  
    DESTCORL(ORANGE)
```

Deleting a subscription

You can use the MQSC command **DELETE SUB** to delete a local subscription.

```
DELETE SUB(ORANGE)
```

You can also delete a subscription using the **SUBID**:

```
DELETE SUB SUBID(414D5120414141202020202020202020EE921E4E20002A03)
```

Checking messages on a subscription

About this task

When a subscription is defined it is associated with a queue. Published messages matching this subscription are put to this queue.

Note that the following **runmqsc** commands show only those subscriptions that received messages.

To check for messages currently queued for a subscription perform the following steps:

Procedure

1. To check for messages queued for a subscription type **DISPLAY SBSTATUS(<sub_name>) NUMMSGs**, see “Displaying attributes of subscriptions” on page 99.
2. If the **NUMMSGs** value is greater than zero identify the queue associated with the subscription by typing **DISPLAY SUB(<sub_name>)DEST**.
3. Using the name of the queue returned you can view the messages by following the technique described in “Browsing queues” on page 91.

Working with services

Service objects are a means by which additional processes can be managed as part of a queue manager. With services, you can define programs that are started and stopped when the queue manager starts and ends. IBM WebSphere MQ services are always started under the user ID of the user who started the queue manager.

Service objects can be either of the following types:

Server A server is a service object that has the parameter `SERVTYPE` specified as `SERVER`. A server service object is the definition of a program that is executed when a specified queue manager is started. Server service objects define programs that typically run for a long time. For example, a server service object can be used to execute a trigger monitor process, such as `runmqtrm`.

Only one instance of a server service object can run concurrently. The status of running server service objects can be monitored using the MQSC command, `DISPLAY SVSTATUS`.

Command

A command is a service object that has the parameter `SERVTYPE` specified as `COMMAND`. Command service objects are similar to server service objects, however multiple instances of a command service object can run concurrently, and their status cannot be monitored using the MQSC command `DISPLAY SVSTATUS`.

If the MQSC command, `STOP SERVICE`, is executed no check is made to determine whether the program started by the MQSC command, `START SERVICE`, is still active before executing the stop program.

Defining a service object

You define a service object with various attributes.

The attributes are as follows:

SERVTYPE

Defines the type of the service object. Possible values are as follows:

SERVER

A server service object.

Only one instance of a server service object can be executed at a time. The status of server service objects can be monitored using the MQSC command, `DISPLAY SVSTATUS`.

COMMAND

A command service object.

Multiple instances of a command service object can be executed concurrently. The status of a command service objects cannot be monitored.

STARTCMD

The program that is executed to start the service. A fully qualified path to the program must be specified.

STARTARG

Arguments passed to the start program.

STDERR

Specifies the path to a file to which the standard error (`stderr`) of the service program should be redirected.

STDOUT

Specifies the path to a file to which the standard output (`stdout`) of the service program should be redirected.

STOPCMD

The program that is executed to stop the service. A fully qualified path to the program must be specified.

STOPARG

Arguments passed to the stop program.

CONTROL

Specifies how the service is to be started and stopped:

MANUAL

The service is not to be started automatically or stopped automatically. It is controlled by use of the START SERVICE and STOP SERVICE commands. This is the default value.

QMGR

The service being defined is to be started and stopped at the same time as the queue manager is started and stopped.

STARTONLY

The service is to be started at the same time as the queue manager is started, but is not requested to stop when the queue manager is stopped.

Managing services

By using the CONTROL parameter, an instance of a service object can be either started and stopped automatically by the queue manager, or started and stopped using the MQSC commands START SERVICE and STOP SERVICE.

When an instance of a service object is started, a message is written to the queue manager error log containing the name of the service object and the process ID of the started process. An example log entry for a server service object starting follows:

```
02/15/2005 11:54:24 AM - Process(10363.1) User(mqm) Program(amqzmgr0)
Host(HOST_1) Installation(Installation1)
VRMF(7.1.0.0) QMgr(A.B.C)
AMQ5028: The Server 'S1' has started. ProcessId(13031).
```

EXPLANATION:

The Server process has started.

ACTION:

None.

An example log entry for a command service object starting follows:

```
02/15/2005 11:53:55 AM - Process(10363.1) User(mqm) Program(amqzmgr0)
Host(HOST_1) Installation(Installation1)
VRMF(7.1.0.0) QMgr(A.B.C)
AMQ5030: The Command 'C1' has started. ProcessId(13030).
```

EXPLANATION:

The Command has started.

ACTION:

None.

When an instance server service stops, a message is written to the queue manager error logs containing the name of the service and the process ID of the ending process. An example log entry for a server service object stopping follows:

```
02/15/2005 11:54:54 AM - Process(10363.1) User(mqm) Program(amqzmgr0)
Host(HOST_1) Installation(Installation1)
VRMF(7.1.0.0) QMgr(A.B.C)
AMQ5029: The Server 'S1' has ended. ProcessId(13031).
```

EXPLANATION:
 The Server process has ended.
 ACTION:
 None.

Additional environment variables

When a service is started, the environment in which the service process is started is inherited from the environment of the queue manager. It is possible to define additional environment variables to be set in the environment of the service process by adding the variables you want to define to one of the `service.env` environment override files.

Note:

There are two possible files to which you can add environment variables:

- The machine scope `service.env` file, which is located in `/var/mqm` on UNIX and Linux systems, or in the data directory selected during installation on Windows systems.
- The queue manager scope `service.env` file, which is located in the queue manager data directory. For example, the location of the environment override file for a queue manager named `QMNAME` is:
 - On UNIX and Linux systems, `/var/mqm/qmgrs/QMNAME/service.env`
 - On Windows systems, `C:\Program Files\IBM\WebSphere MQ\qmgrs\QMNAME\service.env`

Both files are processed, if available, with definitions in the queue manager scope file taking precedence over those definitions in the machine scope file.

Any environment variable can be specified in `service.env`. For example, if the IBM WebSphere MQ service runs a number of commands, it might be useful to set the `PATH` user variable in the `service.env` file. The values that you set the variable to can't be environment variables; for example `CLASSPATH=%CLASSPATH%` is incorrect. Similarly, on Linux `PATH=$PATH:/opt/mqm/bin` would give unexpected results.

`CLASSPATH` must be capitalized, and the class path statement can contain only literals. Some services (Telemetry for example) set their own class path. The `CLASSPATH` defined in `service.env` is added to it.

The format of the variables defined in the file, `service.env` is a list of name and value variable pairs. Each variable must be defined on a new line, and each variable is taken as it is explicitly defined, including white space. An example of the file, `service.env` follows:

```
#####
**                                                                 **
** <N_OCO_COPYRIGHT>                                             **
** Licensed Materials - Property of IBM                          **
**                                                                 **
** 63H9336                                                         **
** (C) Copyright IBM Corporation 2005, 2019                     **
**                                                                 **
** <NOC_COPYRIGHT>                                              **
**                                                                 **
#####
** Module Name: service.env                                     **
** Type       : WebSphere MQ service environment file           **
** Function   : Define additional environment variables to be set **
**            : for SERVICE programs.                          **
** Usage     : <VARIABLE>=<VALUE>                              **
**                                                                 **
#####
```

```
MYLOC=/opt/myloc/bin
MYTMP=/tmp
TRACEDIR=/tmp/trace
MYINITQ=ACCOUNTS.INITIATION.QUEUE
```

Replaceable inserts on service definitions

In the definition of a service object, it is possible to substitute tokens. Tokens that are substituted are automatically replaced with their expanded text when the service program is executed. Substitute tokens can be taken from the following list of common tokens, or from any variables that are defined in the file, `service.env`.

The following are common tokens that can be used to substitute tokens in the definition of a service object:

MQ_INSTALL_PATH

The location where WebSphere MQ is installed.

MQ_DATA_PATH

The location of the WebSphere MQ data directory:

- On UNIX and Linux systems, the WebSphere MQ data directory location is `/var/mqm/`
- On Windows systems, the location of the WebSphere MQ data directory is the data directory selected during the installation of WebSphere MQ

QMNAME

The current queue manager name.

MQ_SERVICE_NAME

The name of the service.

MQ_SERVER_PID

This token can only be used by the `STOPARG` and `STOPCMD` arguments.

For server service objects this token is replaced with the process id of the process started by the `STARTCMD` and `STARTARG` arguments. Otherwise, this token is replaced with 0.

MQ_Q_MGR_DATA_PATH

The location of the queue manager data directory.

MQ_Q_MGR_DATA_NAME

The transformed name of the queue manager. For more information on name transformation, see

 [Understanding WebSphere MQ file names](#) (*WebSphere MQ V7.1 Product Overview Guide*).

To use replaceable inserts, insert the token within + characters into any of the `STARTCMD`, `STARTARG`, `STOPCMD`, `STOPARG`, `STDOUT` or `STDERR` strings. For examples of this, see “Examples on using service objects.”

Examples on using service objects

The services in this section are written with UNIX style path separator characters, except where otherwise stated.

Using a server service object:

This example shows how to define, use, and alter, a server service object to start a trigger monitor.

1. A server service object is defined, using the following MQSC command:

```
DEFINE SERVICE(S1) +
  CONTROL(QMGR) +
  SERVTYPE(SERVER) +
  STARTCMD('+MQ_INSTALL_PATH+bin/runmqtrm') +
  STARTARG('-m +QMNAME+ -q ACCOUNTS.INITIATION.QUEUE') +
  STOPCMD('+MQ_INSTALL_PATH+bin/amqsstop') +
  STOPARG('-m +QMNAME+ -p +MQ_SERVER_PID+')
```

Where:

+MQ_INSTALL_PATH+ is a token representing the installation directory.

+QMNAME+ is a token representing the name of the queue manager.

ACCOUNTS.INITIATION.QUEUE is the initiation queue.

amqsstop is a sample program provided with WebSphere MQ which requests the queue manager to break all connections for the process id. amqsstop generates PCF commands, therefore the command server must be running.

+MQ_SERVER_PID+ is a token representing the process id passed to the stop program.

See "Replaceable inserts on service definitions" on page 105 for a list of the common tokens.

2. An instance of the server service object will execute when the queue manager is next started. However, we will start an instance of the server service object immediately with the following MQSC command:

```
START SERVICE(S1)
```

3. The status of the server service process is displayed, using the following MQSC command:

```
DISPLAY SVSTATUS(S1)
```

4. This example now shows how to alter the server service object and have the updates picked up by manually restarting the server service process. The server service object is altered so that the initiation queue is specified as JUPITER.INITIATION.QUEUE. The following MQSC command is used:

```
ALTER SERVICE(S1) +
  STARTARG('-m +QMNAME+ -q JUPITER.INITIATION.QUEUE')
```

Note: A running service will not pick up any updates to its service definition until it is restarted.

5. The server service process is restarted so that the alteration is picked up, using the following MQSC commands:

```
STOP SERVICE(S1)
```

Followed by:

```
START SERVICE(S1)
```

The server service process is restarted and picks up the alterations made in 4.

Note: The MQSC command, STOP SERVICE, can only be used if a STOPCMD argument is specified in the service definition.

Using a command service object:

This example shows how to define a command service object to start a program that writes entries to the operating system's system log when a queue manager is started or stopped.

1. The command service object is defined, using the following MQSC command:

```
DEFINE SERVICE(S2) +
  CONTROL(QMGR) +
  SERVTYPE(COMMAND) +
  STARTCMD('/usr/bin/logger') +
  STARTARG('Queue manager +QMNAME+ starting') +
  STOPCMD('/usr/bin/logger') +
  STOPARG('Queue manager +QMNAME+ stopping')
```

Where:

logger is the UNIX and Linux system supplied command to write to the system log.

+QMNAME+ is a token representing the name of the queue manager.

Using a command service object when a queue manager ends only:

This example shows how to define a command service object to start a program that writes entries to the operating system's system log when a queue manager is stopped only.

1. The command service object is defined, using the following MQSC command:

```
DEFINE SERVICE(S3) +
  CONTROL(QMGR) +
  SERVTYPE(COMMAND) +
  STOPCMD('/usr/bin/logger') +
  STOPARG('Queue manager +QMNAME+ stopping')
```

Where:

logger is a sample program provided with WebSphere MQ that can write entries to the operating system's system log.

+QMNAME+ is a token representing the name of the queue manager.

More on passing arguments:

This example shows how to define a server service object to start a program called runserv when a queue manager is started.

This example is written with Windows style path separator characters.

One of the arguments that is to be passed to the starting program is a string containing a space. This argument needs to be passed as a single string. To achieve this, double quotation marks are used as shown in the following command to define the command service object:

1. The server service object is defined, using the following MQSC command:

```
DEFINE SERVICE(S1) SERVTYPE(SERVER) CONTROL(QMGR) +
  STARTCMD('C:\Program Files\Tools\runsrv.exe') +
  STARTARG('-m +QMNAME+ -d "C:\Program Files\Tools\'') +
  STDOUT('C:\Program Files\Tools\+MQ_SERVICE_NAME+.out')

DEFINE SERVICE(S4) +
  CONTROL(QMGR) +
  SERVTYPE(SERVER) +
```

```
STARTCMD('C:\Program Files\Tools\runserv.exe') +
STARTARG('-m +QMNAME+ -d "C:\Program Files\Tools\ "') +
STDOUT('C:\Program Files\Tools\+MQ_SERVICE_NAME+.out')
```

Where:

+QMNAME+ is a token representing the name of the queue manager.

"C:\Program Files\Tools\" is a string containing a space, which will be passed as a single string.

Autostarting a Service:

This example shows how to define a server service object that can be used to automatically start the Trigger Monitor when the queue manager starts.

1. The server service object is defined, using the following MQSC command:

```
DEFINE SERVICE(TRIG_MON_START) +
CONTROL(QMGR) +
SERVTYPE(SERVER) +
STARTCMD('runmqtrm') +
STARTARG('-m +QMNAME+ -q +IQNAME+')
```

Where:

+QMNAME+ is a token representing the name of the queue manager.

+IQNAME+ is an environment variable defined by the user in one of the service.env files representing the name of the initiation queue.

Managing objects for triggering

WebSphere MQ enables you to start an application automatically when certain conditions on a queue are met. For example, you might want to start an application when the number of messages on a queue reaches a specified number. This facility is called *triggering*. You have to define the objects that support triggering.

Triggering described in detail in  Starting WebSphere MQ applications using triggers (*WebSphere MQ V7.1 Programming Guide*).

Defining an application queue for triggering

An application queue is a local queue that is used by applications for messaging, through the MQI. Triggering requires a number of queue attributes to be defined on the application queue.

Triggering itself is enabled by the *Trigger* attribute (TRIGGER in MQSC commands). In this example, a trigger event is to be generated when there are 100 messages of priority 5 or greater on the local queue MOTOR.INSURANCE.QUEUE, as follows:

```
DEFINE QLOCAL (MOTOR.INSURANCE.QUEUE) +
PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) +
MAXMSGL (2000) +
DEFPSIST (YES) +
INITQ (MOTOR.INS.INIT.QUEUE) +
TRIGGER +
TRIGTYPE (DEPTH) +
TRIGDPTH (100)+
TRIGMPRI (5)
```

where:

QLOCAL (MOTOR.INSURANCE.QUEUE)

Is the name of the application queue being defined.

PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)

Is the name of the process definition that defines the application to be started by a trigger monitor program.

MAXMSGL (2000)

Is the maximum length of messages on the queue.

DEFPSIST (YES)

Specifies that messages on this queue are persistent by default.

INITQ (MOTOR.INS.INIT.QUEUE)

Is the name of the initiation queue on which the queue manager is to put the trigger message.

TRIGGER

Is the trigger attribute value.

TRIGTYPE (DEPTH)

Specifies that a trigger event is generated when the number of messages of the required priority (TRIGMPRI) reaches the number specified in TRIGDPTH.

TRIGDPTH (100)

Is the number of messages required to generate a trigger event.

TRIGMPRI (5)

Is the priority of messages that are to be counted by the queue manager in deciding whether to generate a trigger event. Only messages with priority 5 or higher are counted.

Defining an initiation queue

When a trigger event occurs, the queue manager puts a trigger message on the initiation queue specified in the application queue definition. Initiation queues have no special settings, but you can use the following definition of the local queue MOTOR.INS.INIT.QUEUE for guidance:

```
DEFINE QLOCAL(MOTOR.INS.INIT.QUEUE) +
  GET (ENABLED) +
  NOSHARE +
  NOTRIGGER +
  MAXMSGL (2000) +
  MAXDEPTH (1000)
```

Defining a process

Use the DEFINE PROCESS command to create a process definition. A process definition defines the application to be used to process messages from the application queue. The application queue definition names the process to be used and thereby associates the application queue with the application to be used to process its messages. This is done through the PROCESS attribute on the application queue MOTOR.INSURANCE.QUEUE. The following MQSC command defines the required process, MOTOR.INSURANCE.QUOTE.PROCESS, identified in this example:

```
DEFINE PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) +
  DESCR ('Insurance request message processing') +
  APPLTYPE (UNIX) +
  APPLICID ('/u/admin/test/IRMP01') +
  USERDATA ('open, close, 235')
```

Where:

MOTOR.INSURANCE.QUOTE.PROCESS

Is the name of the process definition.

DESCR ('Insurance request message processing')

Describes the application program to which this definition relates. This text is displayed when

you use the DISPLAY PROCESS command. This can help you to identify what the process does. If you use spaces in the string, you must enclose the string in single quotation marks.

APPLTYPE (UNIX)

Is the type of application to be started.

APPLICID ('/u/admin/test/IRMP01')

Is the name of the application executable file, specified as a fully qualified file name. In Windows systems, a typical APPLICID value would be c:\appl\test\irmp01.exe.

USERDATA ('open, close, 235')

Is user-defined data, which can be used by the application.

Displaying attributes of a process definition

Use the DISPLAY PROCESS command to examine the results of your definition. For example:

```
DISPLAY PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)
```

```
24 : DISPLAY PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS) ALL
AMQ8407: Display Process details.
DESCR ('Insurance request message processing')
APPLICID ('/u/admin/test/IRMP01')
USERDATA (open, close, 235)
PROCESS (MOTOR.INSURANCE.QUOTE.PROCESS)
APPLTYPE (UNIX)
```

You can also use the MQSC command ALTER PROCESS to alter an existing process definition, and the DELETE PROCESS command to delete a process definition.

Administering remote WebSphere MQ objects

This section tells you how to administer WebSphere MQ objects on a remote queue manager using MQSC commands, and how to use remote queue objects to control the destination of messages and reply messages.

This section describes:

- “Channels, clusters, and remote queuing”
- “Remote administration from a local queue manager” on page 112
- “Creating a local definition of a remote queue” on page 118
- “Using remote queue definitions as aliases” on page 120
- “Data conversion” on page 121

Channels, clusters, and remote queuing

A queue manager communicates with another queue manager by sending a message and, if required, receiving back a response. The receiving queue manager could be:

- On the same machine
- On another machine in the same location (or even on the other side of the world)
- Running on the same platform as the local queue manager
- Running on another platform supported by WebSphere MQ

These messages might originate from:

- User-written application programs that transfer data from one node to another

- User-written administration applications that use PCF commands or the MQAI
- The WebSphere MQ Explorer.
- Queue managers sending:
 - Instrumentation event messages to another queue manager
 - MQSC commands issued from a **runmqsc** command in indirect mode (where the commands are run on another queue manager)

Before a message can be sent to a remote queue manager, the local queue manager needs a mechanism to detect the arrival of messages and transport them consisting of:

- At least one channel
- A transmission queue
- A channel initiator

For a remote queue manager to receive a message, a listener is required.

A channel is a one-way communication link between two queue managers and can carry messages destined for any number of queues at the remote queue manager.

Each end of the channel has a separate definition. For example, if one end is a sender or a server, the other end must be a receiver or a requester. A simple channel consists of a *sender channel definition* at the local queue manager end and a *receiver channel definition* at the remote queue manager end. The two definitions must have the same name and together constitute a single message channel.

If you want the remote queue manager to respond to messages sent by the local queue manager, set up a second channel to send responses back to the local queue manager.

Use the MQSC command `DEFINE CHANNEL` to define channels. In this section, the examples relating to channels use the default channel attributes unless otherwise specified.

There is a message channel agent (MCA) at each end of a channel, controlling the sending and receiving of messages. The MCA takes messages from the transmission queue and puts them on the communication link between the queue managers.

A transmission queue is a specialized local queue that temporarily holds messages before the MCA picks them up and sends them to the remote queue manager. You specify the name of the transmission queue on a *remote queue definition*.

You can allow an MCA to transfer messages using multiple threads. This process is known as *pipelining*. Pipelining enables the MCA to transfer messages more efficiently, improving channel performance. See

 [Attributes of channels \(WebSphere MQ V7.1 Installing Guide\)](#) for details of how to configure a channel to use pipelining.

“Preparing channels and transmission queues for remote administration” on page 113 tells you how to use these definitions to set up remote administration.

For more information about setting up distributed queuing in general, see  [Distributed queuing components \(WebSphere MQ V7.1 Product Overview Guide\)](#).

Remote administration using clusters

In a WebSphere MQ network using distributed queuing, every queue manager is independent. If one queue manager needs to send messages to another queue manager, it must define a transmission queue, a channel to the remote queue manager, and a remote queue definition for every queue to which it wants to send messages.

A *cluster* is a group of queue managers set up in such a way that the queue managers can communicate directly with one another over a single network without complex transmission queue, channel, and queue definitions. Clusters can be set up easily, and typically contain queue managers that are logically related in some way and need to share data or applications. Even the smallest cluster reduces system administration costs.

Establishing a network of queue managers in a cluster involves fewer definitions than establishing a traditional distributed queuing environment. With fewer definitions to make, you can set up or change your network more quickly and easily, and reduce the risk of making an error in your definitions.

To set up a cluster, you need one cluster sender (CLUSDR) and one cluster receiver (CLUSRCVR) definition for each queue manager. You do not need any transmission queue definitions or remote queue definitions. The principles of remote administration are the same when used within a cluster, but the definitions themselves are greatly simplified.

For more information about clusters, their attributes, and how to set them up, see  Queue manager clusters (*WebSphere MQ V7.1 Product Overview Guide*).

Remote administration from a local queue manager

This section tells you how to administer a remote queue manager from a local queue manager using MQSC and PCF commands.

Preparing the queues and channels is essentially the same for both MQSC and PCF commands. In this section, the examples show MQSC commands, because they are easier to understand. For more information about writing administration programs using PCF commands, see “Using Programmable Command Formats” on page 7.

You send MQSC commands to a remote queue manager either interactively or from a text file containing the commands. The remote queue manager might be on the same machine or, more typically, on a different machine. You can remotely administer queue managers in other WebSphere MQ environments, including UNIX and Linux systems, Windows systems, IBM i, and z/OS.

To implement remote administration, you must create specific objects. Unless you have specialized requirements, the default values (for example, for maximum message length) are sufficient.

Preparing queue managers for remote administration

How to use MQSC commands to prepare queue managers for remote administration.

Figure 16 on page 113 shows the configuration of queue managers and channels that you need for remote administration using the **runmqsc** command. The object `source.queue.manager` is the source queue manager from which you can issue MQSC commands and to which the results of these commands (operator messages) are returned. The object `target.queue.manager` is the name of the target queue manager, which processes the commands and generates any operator messages.

Note: If you are using **runmqsc** with the `-w` option, `source.queue.manager` *must* be the default queue manager. For further information on creating a queue manager, see  **crtmqm** (*WebSphere MQ V7.1 Reference*).

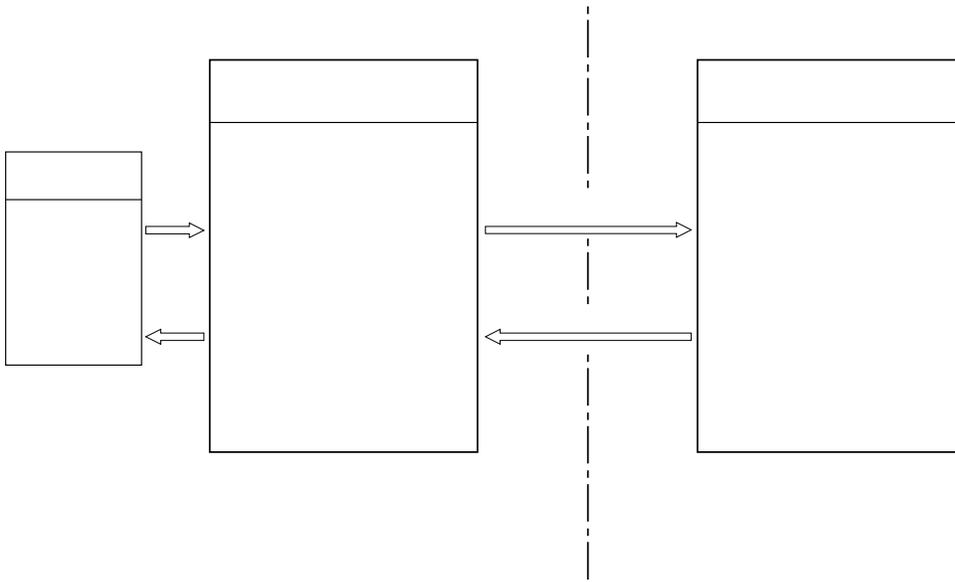


Figure 16. Remote administration using MQSC commands

On both systems, if you have not already done so:

- Create the queue manager and the default objects, using the **crtmqm** command.
- Start the queue manager, using the **strmqm** command.

On the target queue manager:

- The command queue, `SYSTEM.ADMIN.COMMAND.QUEUE`, must be present. This queue is created by default when a queue manager is created.

You have to run these commands locally or over a network facility such as Telnet.

Preparing channels and transmission queues for remote administration

How to use MQSC commands to prepare channels and transmission queues for remote administration.

To run MQSC commands remotely, set up two channels, one for each direction, and their associated transmission queues. This example assumes that you are using TCP/IP as the transport type and that you know the TCP/IP address involved.

The channel `source.to.target` is for sending MQSC commands from the source queue manager to the target queue manager. Its sender is at `source.queue.manager` and its receiver is at `target.queue.manager`. The channel `target.to.source` is for returning the output from commands and any operator messages that are generated to the source queue manager. You must also define a transmission queue for each channel. This queue is a local queue that is given the name of the receiving queue manager. The `XMITQ` name must match the remote queue manager name in order for remote administration to work, unless you are using a queue manager alias. Figure 17 on page 114 summarizes this configuration.

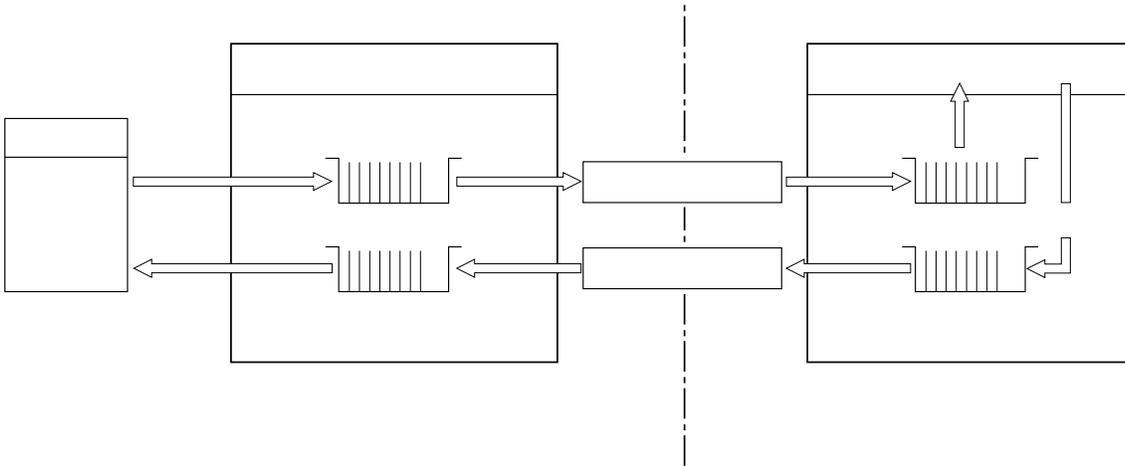


Figure 17. Setting up channels and queues for remote administration

See  Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*) for more information about setting up channels.

Defining channels, listeners, and transmission queues:

On the source queue manager (`source.queue.manager`), issue the following MQSC commands to define the channels, listener, and the transmission queue:

1. Define the sender channel at the source queue manager:


```
DEFINE CHANNEL ('source.to.target') +
  CHLTYPE(SDR) +
  CONNAME (RHX5498) +
  XMITQ ('target.queue.manager') +
  TRPTYPE(TCP)
```
2. Define the receiver channel at the source queue manager:


```
DEFINE CHANNEL ('target.to.source') +
  CHLTYPE(RCVR) +
  TRPTYPE(TCP)
```
3. Define the listener on the source queue manager:


```
DEFINE LISTENER ('source.queue.manager') +
  TRPTYPE (TCP)
```
4. Define the transmission queue on the source queue manager:


```
DEFINE QLOCAL ('target.queue.manager') +
  USAGE (XMITQ)
```

Issue the following commands on the target queue manager (`target.queue.manager`), to create the channels, listener, and the transmission queue:

1. Define the sender channel on the target queue manager:


```
DEFINE CHANNEL ('target.to.source') +
  CHLTYPE(SDR) +
  CONNAME (RHX7721) +
  XMITQ ('source.queue.manager') +
  TRPTYPE(TCP)
```
2. Define the receiver channel on the target queue manager:

```
DEFINE CHANNEL ('source.to.target') +
  CHLTYPE(RCVR) +
  TRPTYPE(TCP)
```

3. Define the listener on the target queue manager:

```
DEFINE LISTENER ('target.queue.manager') +
  TRPTYPE (TCP)
```

4. Define the transmission queue on the target queue manager:

```
DEFINE QLOCAL ('source.queue.manager') +
  USAGE (XMITQ)
```

Note: The TCP/IP connection names specified for the CONNNAME attribute in the sender channel definitions are for illustration only. This is the network name of the machine at the *other* end of the connection. Use the values appropriate for your network.

Starting the listeners and channels:

How to use MQSC commands to start listeners and channels.

Start both listeners by using the following MQSC commands:

1. Start the listener on the source queue manager, `source.queue.manager`, by issuing the following MQSC command:

```
START LISTENER ('source.queue.manager')
```

2. Start the listener on the target queue manager, `target.queue.manager`, by issuing the following MQSC command:

```
START LISTENER ('target.queue.manager')
```

Start both sender channels by using the following MQSC commands:

1. Start the sender channel on the source queue manager, `source.queue.manager`, by issuing the following MQSC command:

```
START CHANNEL ('source.to.target')
```

2. Start the sender channel on the target queue manager, `target.queue.manager`, by issuing the following MQSC command:

```
START CHANNEL ('target.to.source')
```

Automatic definition of channels:

You enable automatic definition of receiver and server-connection definitions by updating the queue manager object using the MQSC command, `ALTER QMGR` (or the PCF command `Change Queue Manager`).

If WebSphere MQ receives an inbound attach request and cannot find an appropriate receiver or server-connection channel, it creates a channel automatically. Automatic definitions are based on two default definitions supplied with WebSphere MQ: `SYSTEM.AUTO.RECEIVER` and `SYSTEM.AUTO.SVRCONN`.

For more information about creating channel definitions automatically, see  [Preparing channels \(WebSphere MQ V7.1 Installing Guide\)](#). For information about automatically defining channels for clusters, see  [Auto-definition of cluster channels \(WebSphere MQ V7.1 Installing Guide\)](#).

Managing the command server for remote administration

How to start, stop, and display the status of the command server. A command server is mandatory for all administration involving PCF commands, the MQAI, and also for remote administration.

Each queue manager can have a command server associated with it. A command server processes any incoming commands from remote queue managers, or PCF commands from applications. It presents the commands to the queue manager for processing and returns a completion code or operator message depending on the origin of the command.

Note: For remote administration, ensure that the target queue manager is running. Otherwise, the messages containing commands cannot leave the queue manager from which they are issued. Instead, these messages are queued in the local transmission queue that serves the remote queue manager. Avoid this situation.

There are separate control commands for starting and stopping the command server. Providing the command server is running, users of WebSphere MQ for Windows or WebSphere MQ for Linux (x86 and x86-64 platforms) can perform the operations described in the following sections using the WebSphere MQ Explorer. For more information, see “Administration using the IBM WebSphere MQ Explorer” on page 59.

Starting the command server

Depending on the value of the queue manager attribute, *SCMDSERV*, the command server is either started automatically when the queue manager starts, or must be started manually. The value of the queue manager attribute can be altered using the MQSC command ALTER QMGR specifying the parameter *SCMDSERV*. By default, the command server is started automatically.

If *SCMDSERV* is set to *MANUAL*, start the command server using the command:

```
strmqcsv saturn.queue.manager
```

where *saturn.queue.manager* is the queue manager for which the command server is being started.

Displaying the status of the command server

For remote administration, ensure that the command server on the target queue manager is running. If it is not running, remote commands cannot be processed. Any messages containing commands are queued in the target queue manager's command queue.

To display the status of the command server for a queue manager, issue the following MQSC command:

```
DISPLAY QMSTATUS CMDSERV
```

Stopping a command server

To end the command server started by the previous example use the following command:

```
endmqcsv saturn.queue.manager
```

You can stop the command server in two ways:

- For a controlled stop, use the **endmqcsv** command with the **-c** flag, which is the default.
- For an immediate stop, use the **endmqcsv** command with the **-i** flag.

Note: Stopping a queue manager also ends the command server associated with it.

Issuing MQSC commands on a remote queue manager

You can use a particular form of the **runmqsc** command to run MQSC commands on a remote queue manager.

The command server *must* be running on the target queue manager, if it is going to process MQSC commands remotely. (This is not necessary on the source queue manager). For information on how to start the command server on a queue manager, see “Managing the command server for remote administration” on page 116.

On the source queue manager, you can then run MQSC commands interactively in indirect mode by typing:

```
runmqsc -w 30 target.queue.manager
```

This form of the **runmqsc** command, with the **-w** flag, runs the MQSC commands in indirect mode, where commands are put (in a modified form) on the command server input queue and executed in order.

When you type in an MQSC command, it is redirected to the remote queue manager, in this case, `target.queue.manager`. The timeout is set to 30 seconds; if a reply is not received within 30 seconds, the following message is generated on the local (source) queue manager:

```
AMQ8416: MQSC timed out waiting for a response from the command server.
```

When you stop issuing MQSC commands, the local queue manager displays any timed-out responses that have arrived and discards any further responses.

The source queue manager defaults to the default local queue manager. If you specify the **-m** *LocalQmgrName* option in the **runmqsc** command, you can direct the commands to be issued by way of any local queue manager.

In indirect mode, you can also run an MQSC command file on a remote queue manager. For example:

```
runmqsc -w 60 target.queue.manager < mycomds.in > report.out
```

where `mycomds.in` is a file containing MQSC commands and `report.out` is the report file.

Suggested method for issuing commands remotely

When you are issuing commands on a remote queue manager, consider using the following approach:

1. Put the MQSC commands to be run on the remote system in a command file.
2. Verify your MQSC commands locally, by specifying the **-v** flag on the **runmqsc** command.
You cannot use **runmqsc** to verify MQSC commands on another queue manager.
3. Check that the command file runs locally without error.
4. Run the command file on the remote system.

If you have problems using MQSC commands remotely

If you have difficulty in running MQSC commands remotely, make sure that you have:

- Started the command server on the target queue manager.
- Defined a valid transmission queue.
- Defined the two ends of the message channels for both:
 - The channel along which the commands are being sent.
 - The channel along which the replies are to be returned.
- Specified the correct connection name (CONNNAME) in the channel definition.

- Started the listeners before you started the message channels.
- Checked that the disconnect interval has not expired, for example, if a channel started but then shut down after some time. This is especially important if you start the channels manually.
- Sent requests from a source queue manager that do not make sense to the target queue manager (for example, requests that include parameters that are not supported on the remote queue manager).

See also “Resolving problems with MQSC commands” on page 84.

Working with queue managers on z/OS:

You can issue MQSC commands to a z/OS queue manager from a queue manager on the platforms described in this guide. However, to do this, you must modify the **runmqsc** command and the channel definitions at the sender.

In particular, you add the **-x** flag to the **runmqsc** command on the source node to specify that the target queue manager is running under z/OS:

```
runmqsc -w 30 -x target.queue.manager
```

Creating a local definition of a remote queue

A local definition of a remote queue is a definition on a local queue manager that refers to a queue on a remote queue manager.

You do not have to define a remote queue from a local position, but the advantage of doing so is that applications can refer to the remote queue by its locally-defined name instead of having to specify a name that is qualified by the ID of the queue manager on which the remote queue is located.

Understanding how local definitions of remote queues work

An application connects to a local queue manager and then issues an **MQOPEN** call. In the open call, the queue name specified is that of a remote queue definition on the local queue manager. The remote queue definition supplies the names of the target queue, the target queue manager, and optionally, a transmission queue. To put a message on the remote queue, the application issues an **MQPUT** call, specifying the handle returned from the **MQOPEN** call. The queue manager uses the remote queue name and the remote queue manager name in a transmission header at the start of the message. This information is used to route the message to its correct destination in the network.

As administrator, you can control the destination of the message by altering the remote queue definition.

The following example shows how an application puts a message on a queue owned by a remote queue manager. The application connects to a queue manager, for example, `saturn.queue.manager`. The target queue is owned by another queue manager.

On the **MQOPEN** call, the application specifies these fields:

Field value	Description
<i>ObjectName</i> CYAN.REMOTE.QUEUE	Specifies the local name of the remote queue object. This defines the target queue and the target queue manager.
<i>ObjectType</i> (Queue)	Identifies this object as a queue.
<i>ObjectQmgrName</i> Blank or saturn.queue.manager	This field is optional. If blank, the name of the local queue manager is assumed. (This is the queue manager on which the remote queue definition exists.)

After this, the application issues an **MQPUT** call to put a message onto this queue.

On the local queue manager, you can create a local definition of a remote queue using the following MQSC commands:

```
DEFINE QREMOTE (CYAN.REMOTE.QUEUE) +
  DESCR ('Queue for auto insurance requests from the branches') +
  RNAME (AUTOMOBILE.INSURANCE.QUOTE.QUEUE) +
  RQMNAME (jupiter.queue.manager) +
  XMITQ (INQUOTE.XMIT.QUEUE)
```

where:

QREMOTE (CYAN.REMOTE.QUEUE)

Specifies the local name of the remote queue object. This is the name that applications connected to this queue manager must specify in the **MQOPEN** call to open the queue **AUTOMOBILE.INSURANCE.QUOTE.QUEUE** on the remote queue manager **jupiter.queue.manager**.

DESCR ('Queue for auto insurance requests from the branches')

Provides additional text that describes the use of the queue.

RNAME (AUTOMOBILE.INSURANCE.QUOTE.QUEUE)

Specifies the name of the target queue on the remote queue manager. This is the real target queue for messages sent by applications that specify the queue name **CYAN.REMOTE.QUEUE**. The queue **AUTOMOBILE.INSURANCE.QUOTE.QUEUE** must be defined as a local queue on the remote queue manager.

RQMNAME (jupiter.queue.manager)

Specifies the name of the remote queue manager that owns the target queue **AUTOMOBILE.INSURANCE.QUOTE.QUEUE**.

XMITQ (INQUOTE.XMIT.QUEUE)

Specifies the name of the transmission queue. This is optional; if the name of a transmission queue is not specified, a queue with the same name as the remote queue manager is used.

In either case, the appropriate transmission queue must be defined as a local queue with a *Usage* attribute specifying that it is a transmission queue (**USAGE(XMITQ)** in MQSC commands).

An alternative way of putting messages on a remote queue

Using a local definition of a remote queue is not the only way of putting messages on a remote queue. Applications can specify the full queue name, including the remote queue manager name, as part of the **MQOPEN** call. In this case, you do not need a local definition of a remote queue. However, this means that applications must either know, or have access to, the name of the remote queue manager at run time.

Using other commands with remote queues

You can use MQSC commands to display or alter the attributes of a remote queue object, or you can delete the remote queue object. For example:

- To display the remote queue's attributes:
`DISPLAY QUEUE (CYAN.REMOTE.QUEUE)`
- To change the remote queue to enable puts. This does not affect the target queue, only applications that specify this remote queue:
`ALTER QREMOTE (CYAN.REMOTE.QUEUE) PUT(ENABLED)`
- To delete this remote queue. This does not affect the target queue, only its local definition:
`DELETE QREMOTE (CYAN.REMOTE.QUEUE)`

Note: When you delete a remote queue, you delete only the local representation of the remote queue. You do not delete the remote queue itself or any messages on it.

Defining a transmission queue

A transmission queue is a local queue that is used when a queue manager forwards messages to a remote queue manager through a message channel.

The channel provides a one-way link to the remote queue manager. Messages are queued at the transmission queue until the channel can accept them. When you define a channel, you must specify a transmission queue name at the sending end of the message channel.

The MQSC command attribute USAGE defines whether a queue is a transmission queue or a normal queue.

Default transmission queues

When a queue manager sends messages to a remote queue manager, it identifies the transmission queue using the following sequence:

1. The transmission queue named on the XMITQ attribute of the local definition of a remote queue.
2. A transmission queue with the same name as the target queue manager. (This value is the default value on XMITQ of the local definition of a remote queue.)
3. The transmission queue named on the DEFEXMITQ attribute of the local queue manager.

For example, the following MQSC command creates a default transmission queue on source.queue.manager for messages going to target.queue.manager:

```
DEFINE QLOCAL ('target.queue.manager') +
        DESCR ('Default transmission queue for target qm') +
        USAGE (XMITQ)
```

Applications can put messages directly on a transmission queue, or indirectly through a remote queue definition. See also “Creating a local definition of a remote queue” on page 118.

Using remote queue definitions as aliases

In addition to locating a queue on another queue manager, you can also use a local definition of a remote queue for Queue manager aliases and reply-to queue aliases. Both types of alias are resolved through the local definition of a remote queue. You must set up the appropriate channels for the message to arrive at its destination.

Queue manager aliases

An alias is the process by which the name of the target queue manager, as specified in a message, is modified by a queue manager on the message route. Queue manager aliases are important because you can use them to control the destination of messages within a network of queue managers.

You do this by altering the remote queue definition on the queue manager at the point of control. The sending application is not aware that the queue manager name specified is an alias.

For more information about queue manager aliases, see  [What are aliases? \(WebSphere MQ V7.1 Product Overview Guide\)](#).

Reply-to queue aliases

Optionally, an application can specify the name of a reply-to queue when it puts a *request message* on a queue.

If the application that processes the message extracts the name of the reply-to queue, it knows where to send the *reply message*, if required.

A reply-to queue alias is the process by which a reply-to queue, as specified in a request message, is altered by a queue manager on the message route. The sending application is not aware that the reply-to queue name specified is an alias.

A reply-to queue alias lets you alter the name of the reply-to queue and optionally its queue manager. This in turn lets you control which route is used for reply messages.

For more information about request messages, reply messages, and reply-to queues, see  [Types of message \(WebSphere MQ V7.1 Programming Guide\)](#) and  [Reply-to queue and queue manager \(WebSphere MQ V7.1 Programming Guide\)](#).

For more information about reply-to queue aliases, see  [Reply-to queue aliases and clusters \(WebSphere MQ V7.1 Installing Guide\)](#).

Data conversion

Message data in WebSphere MQ defined formats (also known as *built-in formats*) can be converted by the queue manager from one coded character set to another, provided that both character sets relate to a single language or a group of similar languages.

For example, conversion between coded character sets with identifiers (CCSIDs) 850 and 500 is supported, because both apply to Western European languages.

For EBCDIC newline (NL) character conversions to ASCII, see  [All queue managers \(WebSphere MQ V7.1 Installing Guide\)](#).

Supported conversions are defined in the  [Data conversion \(WebSphere MQ V7.1 Reference\)](#).

When a queue manager cannot convert messages in built-in formats

The queue manager cannot automatically convert messages in built-in formats if their CCSIDs represent different national-language groups. For example, conversion between CCSID 850 and CCSID 1025 (which is an EBCDIC coded character set for languages using Cyrillic script) is not supported because many of the characters in one coded character set cannot be represented in the other. If you have a network of queue managers working in different national languages, and data conversion among some of the coded character sets is not supported, you can enable a default conversion. Default data conversion is described in “Default data conversion” on page 122.

File ccsid.tbl

The file `ccsid.tbl` is used for the following purposes:

- In WebSphere MQ for Windows it records all the supported code sets.
- On AIX and HP-UX platforms, the supported code sets are held internally by the operating system.
- For all other UNIX and Linux platforms, the supported code sets are held in conversion tables provided by WebSphere MQ.
- It specifies any additional code sets. To specify additional code sets, you need to edit `ccsid.tbl` (guidance on how to do this is provided in the file).
- It specifies any default data conversion.

You can update the information recorded in `ccsid.tbl`; you might want to do this if, for example, a future release of your operating system supports additional coded character sets.

In WebSphere MQ for Windows, `ccsid.tbl` is located in directory `C:\Program Files\IBM\WebSphere MQ\conv\table` by default.

In WebSphere MQ for UNIX and Linux systems, `ccsid.tbl` is located in directory `/var/mqm/conv/table`.

Default data conversion

If you set up channels between two machines on which data conversion is not normally supported, you must enable default data conversion for the channels to work.

To enable default data conversion, edit the `ccsid.tbl` file to specify a default EBCDIC CCSID and a default ASCII CCSID. Instructions on how to do this are included in the file. You must do this on all machines that will be connected using the channels. Restart the queue manager for the change to take effect.

The default data-conversion process is as follows:

- If conversion between the source and target CCSIDs is not supported, but the CCSIDs of the source and target environments are either both EBCDIC or both ASCII, the character data is passed to the target application without conversion.
- If one CCSID represents an ASCII coded character set, and the other represents an EBCDIC coded character set, WebSphere MQ converts the data using the default data-conversion CCSIDs defined in `ccsid.tbl`.

Note: Try to restrict the characters being converted to those that have the same code values in the coded character set specified for the message and in the default coded character set. If you use only the set of characters that is valid for WebSphere MQ object names (as defined in  *Naming IBM WebSphere MQ objects (WebSphere MQ V7.1 Product Overview Guide)*) you will, in general, satisfy this requirement. Exceptions occur with EBCDIC CCSIDs 290, 930, 1279, and 5026 used in Japan, where the lowercase characters have different codes from those used in other EBCDIC CCSIDs.

Converting messages in user-defined formats

The queue manager cannot convert messages in user-defined formats from one coded character set to another. If you need to convert data in a user-defined format, you must supply a data-conversion exit for each such format. Do not use default CCSIDs to convert character data in user-defined formats. For more information about converting data in user-defined formats and about writing data conversion exits, see the  *Writing data-conversion exits (WebSphere MQ V7.1 Programming Guide)*.

Changing the queue manager CCSID

When you have used the `CCSID` attribute of the `ALTER QMGR` command to change the CCSID of the queue manager, stop and restart the queue manager to ensure that all running applications, including the command server and channel programs, are stopped and restarted.

This is necessary because any applications that are running when the queue manager CCSID is changed continue to use the existing CCSID.

Administering IBM WebSphere MQ Telemetry

IBM WebSphere MQ Telemetry is administered using IBM WebSphere MQ Explorer or at a command line. Use the explorer to configure telemetry channels, control the telemetry service, and monitor the MQTT clients that are connected to IBM WebSphere MQ. Configure the security of IBM WebSphere MQ Telemetry using JAAS, SSL and the IBM WebSphere MQ object authority manager.

Administering using IBM WebSphere MQ Explorer

Use the explorer to configure telemetry channels, control the telemetry service, and monitor the MQTT clients that are connected to IBM WebSphere MQ. Configure the security of IBM WebSphere MQ Telemetry using JAAS, SSL and the IBM WebSphere MQ object authority manager.

Administering using the command line

IBM WebSphere MQ Telemetry can be completely administered at the command line using the IBM WebSphere MQ  MQSC (*WebSphere MQ V7.1 Reference*) commands.

The IBM WebSphere MQ Telemetry documentation also has sample scripts that demonstrate the basic usage of the MQ Telemetry Transport v3 Client application.

Read and understand the samples in  MQ Telemetry Transport sample programs (*WebSphere MQ V7.1 Programming Guide*) in the  Developing applications for IBM WebSphere MQ Telemetry (*WebSphere MQ V7.1 Programming Guide*) section before using them.

Related concepts:

“Configure distributed queuing to send messages to MQTT clients” on page 127

“MQTT client identification, authorization, and authentication” on page 129

“Telemetry channel authentication using SSL” on page 137

“Publication privacy on telemetry channels” on page 139

“SSL configuration of MQTT clients and telemetry channels” on page 140

“Telemetry channel JAAS configuration” on page 145

“WebSphere MQ Telemetry daemon for devices concepts” on page 147

Related tasks:

“Configuring a queue manager for telemetry on Linux and AIX”

“Configuring a queue manager for telemetry on Windows” on page 125

Related information:

 IBM WebSphere MQ Telemetry (*WebSphere MQ V7.1 Product Overview Guide*)

Configuring a queue manager for telemetry on Linux and AIX

Follow these manual steps to configure a queue manager to run IBM WebSphere MQ Telemetry. You can run an automated procedure to set up a simpler configuration using the IBM WebSphere MQ Telemetry support for IBM WebSphere MQ Explorer.

Before you begin

1. See  Installing IBM WebSphere MQ Telemetry for information on how to install IBM WebSphere MQ Version 7.1, and the IBM WebSphere MQ Telemetry feature.
2. Create and start a queue manager. The queue manager is referred to as *qMgr* in this task.

About this task

The IBM WebSphere MQ Telemetry support for IBM WebSphere MQ Explorer includes a wizard, and a sample command procedure `sampleMQM`. They set up an initial configuration using the guest user ID; see

 Verifying the installation of IBM WebSphere MQ Telemetry by using IBM WebSphere MQ Explorer (*WebSphere MQ V7.1 Installing Guide*) and  MQ Telemetry Transport sample programs.

Follow the steps in this task to configure IBM WebSphere MQ Telemetry manually using different authorization schemes.

Procedure

1. Open a command window at the telemetry samples directory.

The telemetry samples directory is `/opt/mqm/mqxr/samples`.

2. Create the telemetry transmission queue.

```
echo "DEFINE QLOCAL('SYSTEM.MQTT.TRANSMIT.QUEUE') USAGE(XMITQ)
MAXDEPTH(100000)" | runmqsc qMgr
```

When the telemetry service is first started, it creates `SYSTEM.MQTT.TRANSMIT.QUEUE`.

It is created manually in this task, because `SYSTEM.MQTT.TRANSMIT.QUEUE` must exist before the telemetry service is started, to authorize access to it.

3. Set the default transmission queue for `qMgr`

```
echo "ALTER QMGR DEFQXMITQ('SYSTEM.MQTT.TRANSMIT.QUEUE')" | runmqsc qMgr
```

When the telemetry service is first started, it does not alter the queue manager to make `SYSTEM.MQTT.TRANSMIT.QUEUE` the default transmission queue.

To make `SYSTEM.MQTT.TRANSMIT.QUEUE` the default transmission queue alter the default transmission queue property. Alter the property using the IBM WebSphere MQ Explorer or with the command in Figure 19 on page 126.

Altering the default transmission queue might interfere with your existing configuration. The reason for altering the default transmission queue to `SYSTEM.MQTT.TRANSMIT.QUEUE` is to make sending messages directly to MQTT clients easier. Without altering the default transmission queue you must add a remote queue definition for every client that receives IBM WebSphere MQ messages; see “Sending a message to a client directly” on page 128.

4. Follow a procedure in “Authorizing MQTT clients to access WebSphere MQ objects” on page 131 to create one or more user IDs. The user IDs have the authority to publish, subscribe, and send publications to MQTT clients.

5. Install the telemetry service

```
cat installMQXRService_unix.mqsc | runmqsc qMgr
```

6. Start the service

```
echo "START SERVICE(SYSTEM.MQXR.SERVICE)" | runmqsc qMgr
```

The telemetry service is started automatically when the queue manager is started.

It is started manually in this task, because the queue manager is already running.

7. Using IBM WebSphere MQ Explorer, configure telemetry channels to accept connections from MQTT clients.

8. Verify the configuration by running the sample client.

For the sample client to work with your telemetry channel, the channel must authorize the client to publish, subscribe, and receive publications. The sample client connects to the telemetry channel on port 1883 by default.

Example

Figure 18 shows the `runmqsc` command to create the `SYSTEM.MQXR.SERVICE` manually on Linux.

```
DEF SERVICE(SYSTEM.MQXR.SERVICE) +
  CONTROL(QMGR) +
  DESCR('Manages clients using MQXR protocols such as MQTT') +
  SERVTYPE(SERVER) +
  STARTCMD('+MQ_INSTALL_PATH+/mqxr/bin/runMQXRService.sh') +
  STARTARG('-m +QMNAME+ -d "+MQ_Q_MGR_DATA_PATH+" -g "+MQ_DATA_PATH+"') +
  STOPCMD('+MQ_INSTALL_PATH+/mqxr/bin/endMQXRService.sh') +
  STOPARG('-m +QMNAME+') +
  STDOUT('+MQ_Q_MGR_DATA_PATH+/mqxr.stdout') +
  STDERR('+MQ_Q_MGR_DATA_PATH+/mqxr.stderr')
```

Figure 18. `installMQXRService_unix.mqsc`

Configuring a queue manager for telemetry on Windows

Follow these manual steps to configure a queue manager to run IBM WebSphere MQ Telemetry. You can run an automated procedure to set up a simpler configuration using the IBM WebSphere MQ Telemetry support for IBM WebSphere MQ Explorer.

Before you begin

1. See  Installing IBM WebSphere MQ Telemetry for information on how to install IBM WebSphere MQ Version 7.1, and the IBM WebSphere MQ Telemetry feature.
2. Create and start a queue manager. The queue manager is referred to as *qMgr* in this task.

About this task

The IBM WebSphere MQ Telemetry support for IBM WebSphere MQ Explorer includes a wizard, and a sample command procedure `sampleMQM`. They set up an initial configuration using the guest user ID; see

 Verifying the installation of IBM WebSphere MQ Telemetry by using IBM WebSphere MQ Explorer (*WebSphere MQ V7.1 Installing Guide*) and  MQ Telemetry Transport sample programs.

Follow the steps in this task to configure IBM WebSphere MQ Telemetry manually using different authorization schemes.

Procedure

1. Open a command window at the telemetry samples directory.
The telemetry samples directory is `WMQ program installation directory\mqxr\samples`.
2. Create the telemetry transmission queue.

```
echo DEFINE QLOCAL('SYSTEM.MQTT.TRANSMIT.QUEUE') USAGE(XMITQ)
MAXDEPTH(100000) | runmqsc qMgr
```

When the telemetry service is first started, it creates `SYSTEM.MQTT.TRANSMIT.QUEUE`.

It is created manually in this task, because `SYSTEM.MQTT.TRANSMIT.QUEUE` must exist before the telemetry service is started, to authorize access to it.

3. Set the default transmission queue for *qMgr*

```
echo ALTER QMGR DEFXMITQ('SYSTEM.MQTT.TRANSMIT.QUEUE') | runmqsc qMgr
```

Figure 19. Set default transmission queue

When the telemetry service is first started, it does not alter the queue manager to make SYSTEM.MQTT.TRANSMIT.QUEUE the default transmission queue.

To make SYSTEM.MQTT.TRANSMIT.QUEUE the default transmission queue alter the default transmission queue property. Alter the property using the IBM WebSphere MQ Explorer or with the command in Figure 19.

Altering the default transmission queue might interfere with your existing configuration. The reason for altering the default transmission queue to SYSTEM.MQTT.TRANSMIT.QUEUE is to make sending messages directly to MQTT clients easier. Without altering the default transmission queue you must add a remote queue definition for every client that receives IBM WebSphere MQ messages; see “Sending a message to a client directly” on page 128.

4. Follow a procedure in “Authorizing MQTT clients to access WebSphere MQ objects” on page 131 to create one or more user IDs. The user IDs have the authority to publish, subscribe, and send publications to MQTT clients.

5. Install the telemetry service

```
type installMQXRService_win.mqsc | runmqsc qMgr
```

6. Start the service

```
echo START SERVICE(SYSTEM.MQXR.SERVICE) | runmqsc qMgr
```

The telemetry service is started automatically when the queue manager is started.

It is started manually in this task, because the queue manager is already running.

7. Using IBM WebSphere MQ Explorer, configure telemetry channels to accept connections from MQTT clients.

The telemetry channels must be configured such that their identities are one of the user IDs defined in step 4.

8. Verify the configuration by running the sample client.

For the sample client to work with your telemetry channel, the channel must authorize the client to publish, subscribe, and receive publications. The sample client connects to the telemetry channel on port 1883 by default.

Creating SYSTEM.MQXR.SERVICE manually

Figure 20 shows the **runmqsc** command to create the SYSTEM.MQXR.SERVICE manually on Windows.

```
DEF SERVICE(SYSTEM.MQXR.SERVICE) +  
CONTROL(QMGR) +  
DESCR('Manages clients using MQXR protocols such as MQTT') +  
SERVTYPE(SERVER) +  
STARTCMD('+MQ_INSTALL_PATH+\mqxr\bin\runMQXRService.bat') +  
STARTARG('-m +QMNAME+ -d "+MQ_Q_MGR_DATA_PATH+\" -g "+MQ_DATA_PATH+\"') +  
STOPCMD('+MQ_INSTALL_PATH+\mqxr\bin\endMQXRService.bat') +  
STOPARG('-m +QMNAME+') +  
STDOUT('+MQ_Q_MGR_DATA_PATH+\mqxr.stdout') +  
STDERR('+MQ_Q_MGR_DATA_PATH+\mqxr.stderr')
```

Figure 20. installMQXRService_win.mqsc

Configure distributed queuing to send messages to MQTT clients

WebSphere MQ applications can send MQTT v3 clients messages by publishing to subscription created by a client, or by sending a message directly. Whichever method is used, the message is placed on SYSTEM.MQTT.TRANSMIT.QUEUE, and sent to the client by the telemetry service. There are a number of ways to place a message on SYSTEM.MQTT.TRANSMIT.QUEUE.

Publishing a message in response to an MQTT client subscription

The telemetry service creates a subscription on behalf of the MQTT client. The client is the destination for any publications that match the subscription sent by the client. The telemetry services forwards matching publications back to the client.

An MQTT client is connected to WebSphere MQ as a queue manager, with its queue manager name set to its *ClientIdentifier*. The destination for publications to be sent to the client is a transmission queue, SYSTEM.MQTT.TRANSMIT.QUEUE. The telemetry service forwards messages on SYSTEM.MQTT.TRANSMIT.QUEUE to MQTT clients, using the target queue manager name as the key to a specific client.

The telemetry service opens the transmission queue using *ClientIdentifier* as the queue manager name. The telemetry service passes the object handle of the queue to the MQSUB call, to forward publications that match the client subscription. In the object name resolution, the *ClientIdentifier* is created as the remote queue manager name, and the transmission queue must resolve to SYSTEM.MQTT.TRANSMIT.QUEUE. Using standard WebSphere MQ object name resolution, *ClientIdentifier* is resolved as follows; see Table 2.

1. *ClientIdentifier* matches nothing.

ClientIdentifier is a remote queue manager name. It does not match the local queue manager name, a queue manager alias, or a transmission queue name.

The queue name is not defined. Currently, the telemetry service sets SYSTEM.MQTT.PUBLICATION.QUEUE as the name of the queue. An MQTT v3 client does not support queues, so the resolved queue name is ignored by the client.

The local queue manager property, Default transmission queue, name must be set to SYSTEM.MQTT.TRANSMIT.QUEUE, so that the publication is put on SYSTEM.MQTT.TRANSMIT.QUEUE to be sent to the client.

2. *ClientIdentifier* matches a queue manager alias named *ClientIdentifier*.

ClientIdentifier is a remote queue manager name. It matches the name of a queue manager alias.

The queue manager alias must be defined with *ClientIdentifier* as the remote queue manager name.

By setting the transmission queue name in the queue manager alias definition it is not necessary for the default transmission to be set to SYSTEM.MQTT.TRANSMIT.QUEUE.

Table 2. Name resolution of an MQTT queue manager alias

	Input		Output		
	Queue manager name	Queue name	Queue manager name	Queue name	Transmission queue
<i>ClientIdentifier</i>	<i>ClientIdentifier</i>	<i>undefined</i>	<i>ClientIdentifier</i>	<i>undefined</i>	Default transmission queue. SYSTEM.MQTT.TRANSMIT.QUEUE

Table 2. Name resolution of an MQTT queue manager alias (continued)

	Input		Output		
	Queue manager name	Queue name	Queue manager name	Queue name	Transmission queue
<i>ClientIdentifier</i>	<i>ClientIdentifier</i>	<i>undefined</i>	<i>ClientIdentifier</i>	<i>undefined</i>	SYSTEM.MQTT. TRANSMIT.QUEUE

For further information about name resolution, see  Name resolution (*WebSphere MQ V7.1 Programming Guide*).

Any WebSphere MQ program can publish to the same topic. The publication is sent to its subscribers, including MQTT v3 clients that have a subscription to the topic.

If an administrative topic is created in a cluster, with the attribute `CLUSTER(clusterName)`, any application in the cluster can publish to the client; for example:

```
echo DEFINE TOPIC('MQTTEXAMPLES') TOPICSTR('MQTT Examples') CLUSTER(MQTT)
REPLACE | runmqsc qMgr
```

Figure 21. Defining a cluster topic on Windows

Note: Do not give `SYSTEM.MQTT.TRANSMIT.QUEUE` a cluster attribute.

MQTT client subscribers and publishers can connect to different queue managers. The subscribers and publishers can be part of the same cluster, or connected by a publish/subscribe hierarchy. The publication is delivered from the publisher to the subscriber using WebSphere MQ.

Sending a message to a client directly

An alternative to a client creating a subscription and receiving a publication that matches the subscription topic, send a message to an MQTT v3 client directly. MQTT V3 client applications cannot send messages directly, but other application, such as WebSphere MQ applications, can.

The WebSphere MQ application must know the `ClientIdentifier` of the MQTT v3 client. As MQTT v3 clients do not have queues, the target queue name is passed to the MQTT v3 application client `messageArrived` method as a topic name. For example, in an MQI program, create an object descriptor with the client as the `ObjectQmgrName`:

```
MQOD.ObjectQmgrName = ClientIdentifier;
MQOD.ObjectName = name;
```

Figure 22. MQI Object descriptor to send a message to an MQTT v3 client destination

If the application is written using JMS, create a point-to-point destination; for example:

```

javax.jms.Destination jmsDestination =
    (javax.jms.Destination)jmsFactory.createQueue
    ("queue://ClientIdentifier/name");

```

Figure 23. JMS destination to send a message to an MQTT v3 client

To send an unsolicited message to an MQTT client use a remote queue definition. The remote queue manager name must resolve to the `ClientIdentifier` of the client. The transmission queue must resolve to `SYSTEM.MQTT.TRANSMIT.QUEUE`; see Table 3. The remote queue name can be anything. The client receives it as a topic string.

Table 3. Name resolution of an MQTT client remote queue definition

Input		Output		
Queue name	Queue manager name	Queue name	Queue manager name	Transmission queue
Name of remote queue definition	Blank or local queue manager name	Remote queue name used as a topic string	<code>ClientIdentifier</code>	<code>SYSTEM.MQTT.TRANSMIT.QUEUE</code>

If the client is connected, the message is sent directly to the MQTT client, which calls the `messageArrived` method; see  `messageArrived` method.

If the client has disconnected with a persistent session, the message is stored in

`SYSTEM.MQTT.TRANSMIT.QUEUE`; see  MQTT stateless and stateful sessions. It is forwarded to the client when the client reconnects to the session again.

If you send a non-persistent message it is sent to the client with "at most once" quality of service, `QoS=0`. If you send a persistent message directly to a client, by default, it is sent with "exactly once" quality of service, `QoS=2`. As the client might not have a persistence mechanism, the client can lower the quality of service it accepts for messages sent directly. To lower the quality of service for messages sent directly to a client, make a subscription to the topic `DEFAULT.QoS`. Specify the maximum quality of service the client can support.

MQTT client identification, authorization, and authentication

The telemetry service publishes, or subscribes to, WebSphere MQ topics on behalf of MQTT clients, using MQTT channels. The WebSphere MQ administrator configures the MQTT channel identity that is used for WebSphere MQ authorization. The administrator can define a common identity for the channel, or use the `Username` or `ClientIdentifier` of a client connected to the channel.

The telemetry service can authenticate the client using the `Username` supplied by the client, or by using a client certificate. The `Username` is authenticated using a password provided by the client.

To summarize: Client identification is the selection of the client identity. Depending on the context, the client is identified by the `ClientIdentifier`, `Username`, a common client identity created by the administrator, or a client certificate. The client identifier used for authenticity checking does not have to be the same identifier that is used for authorization.

MQTT client programs set the `Username` and `Password` that are sent to the server using an MQTT channel. They can also set the SSL properties that are required to encrypt and authenticate the connection. The administrator decides whether to authenticate the MQTT channel, and how to authenticate the channel.

To authorize an MQTT client to access WebSphere MQ objects, authorize the `ClientIdentifier`, or `Username` of the client, or authorize a common client identity. To permit a client to connect to WebSphere MQ, authenticate the `Username`, or use a client certificate. Configure JAAS to authenticate the `Username`, and configure SSL to authenticate a client certificate.

If you set a `Password` at the client, either encrypt the connection using VPN, or configure the MQTT channel to use SSL, to keep the password private.

It is difficult to manage client certificates. For this reason, if the risks associated with password authentication are acceptable, password authentication is often used to authenticate clients.

If there is a secure way to manage and store the client certificate it is possible to rely on certificate authentication. However, it is rarely the case that certificates can be managed securely in the types of environments that telemetry is used in. Instead, the authentication of devices using client certificates is complemented by authenticating client passwords at the server. Because of the additional complexity, the use of client certificates is restricted to highly sensitive applications. The use of two forms of authentication is called two-factor authentication. You must know one of the factors, such as a password, and have the other, such as a certificate.

In a highly sensitive application, such as a chip-and-pin device, the device is locked down during manufacture to prevent tampering with the internal hardware and software. A trusted, time-limited, client certificate is copied to the device. The device is deployed to the location where it is to be used. Further authentication is performed each time the device is used, either using a password, or another certificate from a smart card.

MQTT client identity and authorization

Use the `ClientIdentifier`, `Username`, or a common client identity for authorization to access WebSphere MQ objects.

The WebSphere MQ administrator has three choices for selecting the identity of the MQTT channel. The administrator makes the choice when defining or modifying the MQTT channel used by the client. The identity is used to authorize access to WebSphere MQ topics. The choices are:

1. The client identifier.
2. An identity the administrator provides for the channel.
3. The `Username` passed from the MQTT client.

`Username` is an attribute of the `MqttConnectOptions` class. It must be set before the client connects to the service. Its default value is null.

Use the WebSphere MQ `setmqaut` command to select which objects, and which actions, are authorized to be used by the identity associated with the MQTT channel. For example, to authorize a channel identity, `MQTTClient`, provided by the administrator of queue manager, `QM1`:

```
setmqaut -m QM1 -t q -n SYSTEM.MQTT.TRANSMIT.QUEUE -p MQTTClient -all +put
setmqaut -m QM1 -t topic -n SYSTEM.BASE.TOPIC -p MQTTClient -all +pub +sub
```

Authorizing MQTT clients to access WebSphere MQ objects:

Follow these steps to authorize MQTT clients to publish and subscribe to WebSphere MQ Objects. The steps follow four alternative access control patterns.

Before you begin

MQTT clients are authorized to access objects in WebSphere MQ by being assigned an identity when they connect to a telemetry channel. The WebSphere MQ Administrator configures the telemetry channel using WebSphere MQ Explorer to give a client one of three types of identity:

1. ClientIdentifier
2. Username
3. A name the administrator assigns to the channel.

Whichever type is used, the identity must be defined to WebSphere MQ as a principal by the installed authorization service. The default authorization service on Windows or Linux is called the Object Authority Manager (OAM). If you are using the OAM, the identity must be defined as a user ID.

Use the identity to give a client, or collection of clients, permission to publish or subscribe to topics defined in WebSphere MQ. If an MQTT client has subscribed to a topic, use the identity to give it permission to receive the resulting publications.

It is hard to manage a system with tens of thousands of MQTT clients, each requiring individual access permissions. One solution is to define common identities, and associate individual MQTT clients with one of the common identities. Define as many common identities as you require to define different combinations of permissions. Another solution is to write your own authorization service that can deal more easily with thousands of users than the operating system.

You can combine MQTT clients into common identities in two ways, using the OAM:

1. Define multiple telemetry channels, each with a different user ID that the administrator allocates using WebSphere MQ Explorer. Clients connecting using different TCP/IP port numbers are associated with different telemetry channels, and are assigned different identities.
2. Define a single telemetry channel, but have each client select a Username from a small set of user IDs. The administrator configures the telemetry channel to select the client Username as its identity.

In this task, the identity of the telemetry channel is called *mqttUser*, regardless of how it is set. If collections of clients use different identities, use multiple *mqttUsers*, one for each collection of clients. As the task uses the OAM, each *mqttUser* must be a user ID.

About this task

In this task, you have a choice of four access control patterns that you can tailor to specific requirements. The patterns differ in their granularity of access control.

- “No access control” on page 132
- “Coarse-grained access control” on page 132
- “Medium-grained access control” on page 132
- “Fine-grained access control” on page 132

The result of the models is to assign *mqttUsers* sets of permissions to publish and subscribe to WebSphere MQ, and receive publications from WebSphere MQ.

No access control:

MQTT clients are given WebSphere MQ administrative authority, and can perform any action on any object.

Procedure

1. Create a user ID *mqttUser* to act as the identity of all MQTT clients.
2. Add *mqttUser* to the *mqm* group; see Adding a user to a group on Windows, or Adding a user to a group on Linux

Coarse-grained access control:

MQTT clients have authority to publish and subscribe, and to send messages to MQTT clients. They do not have authority to perform other actions, or to access other objects.

Procedure

1. Create a user ID *mqttUser* to act as the identity of all MQTT clients.
2. Authorize *mqttUser* to publish and subscribe to all topics and to send publications to MQTT clients.

```
setmqaut -m qMgr -t topic -n SYSTEM.BASE.TOPIC -p mqttUser -all +pub +sub  
setmqaut -m qMgr -t q -n SYSTEM.MQTT.TRANSMIT.QUEUE -p mqttUser -all +put
```

Medium-grained access control:

MQTT clients are divided into different groups to publish and subscribe to different sets of topics, and to send messages to MQTT clients.

Procedure

1. Create multiple user IDs, *mqttUsers*, and multiple administrative topics in the publish/subscribe topic tree.
2. Authorize different *mqttUsers* to different topics.

```
setmqaut -m qMgr -t topic -n topic1 -p mqttUserA -all +pub +sub  
setmqaut -m qMgr -t topic -n topic2 -p mqttUserB -all +pub +sub
```
3. Create a group *mqtt*, and add all *mqttUsers* to the group.
4. Authorize *mqtt* to send topics to MQTT clients.

```
setmqaut -m qMgr -t q -n SYSTEM.MQTT.TRANSMIT.QUEUE -p mqtt -all +put
```

Fine-grained access control:

MQTT clients are incorporated into an existing system of access control, that authorizes groups to perform actions on objects.

About this task

A user ID is assigned to one or more operating system groups depending on the authorizations it requires. If WebSphere MQ applications are publishing and subscribing to the same topic space as MQTT clients, use this model. The groups are referred to as Publish X , Subscribe Y , and *mqtt*

Publish X

Members of Publish X groups can publish to *topic X* .

Subscribe Y

Members of Subscribe Y groups can subscribe to *topic Y* .

mqtt Members of the *mqtt* group can send publications to MQTT clients.

Procedure

1. Create multiple groups, *PublishX* and *SubscribeY* that are allocated to multiple administrative topics in the publish/subscribe topic tree.
2. Create a group *mqtt*.
3. Create multiple user IDs, *mqttUsers*, and add the users to any of the groups, depending on what they are authorized to do.
4. Authorize different *PublishX* and *SubscribeX* groups to different topics, and authorize the *mqtt* group to send messages to MQTT clients.

```
setmqaut -m qMgr -t topic -n topic1 -p PublishX -all +pub
setmqaut -m qMgr -t topic -n topic1 -p SubscribeX -all +pub +sub
setmqaut -m qMgr -t q -n SYSTEM.MQTT.TRANSMIT.QUEUE -p mqtt -all +put
```

MQTT client authentication using a password

Authenticate the Username using the client password. You can authenticate the client using a different identity to the identity used to authorize the client to publish and subscribe to topics.

The telemetry service uses JAAS to authenticate the client Username. JAAS uses the Password supplied by the MQTT client.

The WebSphere MQ administrator decides whether to authenticate the Username, or not to authenticate at all, by configuring the MQTT channel a client connects to. Clients can be assigned to different channels, and each channel can be configured to authenticate its clients in different ways. Using JAAS, you can configure which methods must authenticate the client, and which can optionally authenticate the client.

The choice of identity for authentication does not affect the choice of identity for authorization. You might want to set up a common identity for authorization for administrative convenience, but authenticate each user to use that identity. The following procedure outlines the steps to authenticate individual users to use a common identity:

1. The WebSphere MQ administrator sets the MQTT channel identity to any name, such as *MQTTClientUser*, using WebSphere MQ Explorer.
2. The WebSphere MQ administrator authorizes *MQTTClient* to publish and subscribe to any topic:

```
setmqaut -m QM1 -t q -n SYSTEM.MQTT.TRANSMIT.QUEUE -p MQTTClient -all +put
setmqaut -m QM1 -t topic -n SYSTEM.BASE.TOPIC -p MQTTClient -all +pub +sub
```
3. The MQTT client application developer creates an *MqttConnectOptions* object and sets Username and Password before connecting to the server.
4. The security developer creates a JAAS LoginModule to authenticate the Username with the Password and includes it in the JAAS configuration file.
5. The WebSphere MQ administrator configures the MQTT channel to authenticate the Username of the client using JAAS.

MQTT client authentication using SSL

Connections between the MQTT client and the queue manager are always initiated by the MQTT client. The MQTT client is always the SSL client. Client authentication of the server and server authentication of the MQTT client are both optional.

By providing the client with a private signed digital certificate, you can authenticate the MQTT client to WebSphere MQ. The WebSphere MQ Administrator can force MQTT clients to authenticate themselves to the queue manager using SSL. You can only request client authentication as part of mutual authentication.

As an alternative to using SSL, some kinds of Virtual Private Network (VPN), such as IPsec, authenticate the endpoints of a TCP/IP connection. VPN encrypts each IP packet that flows over the network. Once such a VPN connection is established, you have established a trusted network. You can connect MQTT clients to telemetry channels using TCP/IP over the VPN network.

Client authentication using SSL relies upon the client having a secret. The secret is the private key of the client in the case of a self-signed certificate, or a key provided by a certificate authority. The key is used to sign the digital certificate of the client. Anyone in possession of the corresponding public key can verify the digital certificate. Certificates can be trusted, or if they are chained, traced back through a certificate chain to a trusted root certificate. Client verification sends all the certificates in the certificate chain provided by the client to the server. The server checks the certificate chain until it finds a certificate it trusts. The trusted certificate is either the public certificate generated from a self-signed certificate, or a root certificate typically issued by a certificate authority. As a final, optional, step the trusted certificate can be compared with a "live" certificate revocation list.

The trusted certificate might have been issued by a certificate authority and already included in the JRE certificate store. It might be a self-signed certificate, or any certificate that has been added to the telemetry channel keystore as a trusted certificate.

Note: The telemetry channel has a combined keystore/truststore that holds both the private keys to one or more telemetry channels, and any public certificates needed to authenticate clients. Because an SSL channel must have a keystore, and it is the same file as the channel truststore, the JRE certificate store is never referenced. The implication is that if authentication of a client requires a CA root certificate, you must place the root certificate in the keystore for the channel, even if the CA root certificate is already in the JRE certificate store. The JRE certificate store is never referenced.

Think about the threats that client authentication is intended to counter, and the roles the client and server play in countering the threats. Authenticating the client certificate alone is insufficient to prevent unauthorized access to a system. If someone else has got hold of the client device, the client device is not necessarily acting with the authority of the certificate holder. Never rely on a single defense against unwanted attacks. At least use a two-factor authentication approach and supplement possession of a certificate with knowledge of private information. For example, use JAAS, and authenticate the client using a password issued by the server.

The primary threat to the client certificate is that it gets into the wrong hands. The certificate is held in a password protected keystore at the client. How does it get placed in the keystore? How does the MQTT client get the password to the keystore? How secure is the password protection? Telemetry devices are often easy to remove, and then can be hacked in private. Must the device hardware be tamper-proof? Distributing and protecting client-side certificates is recognized to be hard; it is called the key-management problem.

A secondary threat is that the device is misused to access servers in unintended ways. For example, if the MQTT application is tampered with, it might be possible to use a weakness in the server configuration using the authenticated client identity.

To authenticate an MQTT client using SSL, configure the telemetry channel, and the client.

Telemetry channel configuration

The WebSphere MQ administrator configures telemetry channels at the server. Each channel is configured to accept a TCP/IP connection on a different port number. The channels are configured either as `com.ibm.mq.MQTT.channel/PlainText` or `com.ibm.mq.MQTT.channel/SSL`. SSL channels are configured with passphrase protected access to key files. If an SSL channel is defined with no passphrase or key file, the channel does not accept SSL connections.

Set the property, `com.ibm.mq.MQTT.ClientAuth` of an SSL telemetry channel to `REQUIRED` to force all clients connecting on that channel to provide proof that they have verified digital certificates. The client certificates are authenticated using certificates from certificate authorities, leading to a trusted root certificate. If the client certificate is self-signed, or is signed by a certificate that is from a certificate authority, the publicly signed certificates of the client, or certificate authority, must be stored securely at the server.

Place the publicly signed client certificate or the certificate from the certificate authority in the telemetry channel keystore. At the server, publicly signed certificates are stored in the same key file as privately signed certificates, rather than in a separate truststore.

The server verifies the signature of any client certificates it is sent using all the public certificates and cipher suites it has. The server verifies the key chain. The queue manager can be configured to test the certificate against the certificate revocation list. The queue manager revocation namelist property is SSLCRLNL.

If any of the certificates a client sends is verified by a certificate in the server keystore, then the client is authenticated.

The WebSphere MQ administrator can configure the same telemetry channel to use JAAS to check the Username or ClientIdentifier of the client with the client Password.

You can use the same keystore for multiple telemetry channels.

Verification of at least one digital certificate in the password protected client keystore on the device authenticates the client to the server. The digital certificate is only used for authentication by WebSphere MQ. It is not used to verify the TCP/IP address of the client, set the identity of the client for authorization or accounting. The identity of the client adopted by the server is either the Username or ClientIdentifier of the client, or an identity created by the WebSphere MQ administrator.

MQTT Client configuration

To authenticate the MQTT client using SSL, the client connects to a telemetry channel using SSL. It must specify a TCP port that corresponds to a telemetry channel that is configured to authenticate SSL clients. For example, at the client:

```
MQTTClient mqttClient = new MqttClient( "ssl://www.example.org:8884", "clientId");
mqttClient.connect();
```

The client JVM must use the standard socket factory from JSSE. If you are using Java ME, you must ensure that the JSSE package is loaded. If you are using Java SE, JSSE has been included with the JRE since Java version 1.4.1.

The SSL connection requires a number of SSL properties to be set before connecting. You can set the properties either by passing them to the JVM using the `-D` switch, or you can set the properties using the `MqttConnectionOptions.setSSLProperties` method.

If you load a non-standard socket factory, by calling the method `MqttConnectOptions.setSocketFactory(javax.net.SocketFactory)`, then the way SSL settings are passed to the network socket is application defined.

Add the digital certificate of the client, signed either using the private key of the client, or by a CA, to the password protected keystore on the client. If the certificate has a key chain, you can add the certificates from the key chain to the store. When the server verifies the client certificate, it uses the certificates sent by the client to match against certificates in its keystore. It is looking for the first match in the key chain with a certificate that it has. The remainder of the key chain is ignored.

The MQTT client sends all the certificates in its keystore to the server. If the server authenticates any of the key chains the client sends, then the client is authenticated.

You can also use SSL cipher suites for client authentication. Here is an alphabetic list of the SSL cipher suites that are currently supported:

- SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA

- SSL_DH_anon_EXPORT_WITH_RC4_40_MD5
- SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
- SSL_DH_anon_WITH_AES_128_CBC_SHA
- SSL_DH_anon_WITH_DES_CBC_SHA
- SSL_DH_anon_WITH_RC4_128_MD5
- SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_DSS_WITH_AES_128_CBC_SHA
- SSL_DHE_DSS_WITH_DES_CBC_SHA
- SSL_DHE_DSS_WITH_RC4_128_SHA
- SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_RSA_WITH_AES_128_CBC_SHA
- SSL_DHE_RSA_WITH_DES_CBC_SHA
- SSL_KRB5_EXPORT_WITH_DES_CBC_40_MD5
- SSL_KRB5_EXPORT_WITH_DES_CBC_40_SHA
- SSL_KRB5_EXPORT_WITH_RC4_40_MD5
- SSL_KRB5_EXPORT_WITH_RC4_40_SHA
- SSL_KRB5_WITH_3DES_EDE_CBC_MD5
- SSL_KRB5_WITH_3DES_EDE_CBC_SHA
- SSL_KRB5_WITH_DES_CBC_MD5
- SSL_KRB5_WITH_DES_CBC_SHA
- SSL_KRB5_WITH_RC4_128_MD5
- SSL_KRB5_WITH_RC4_128_SHA
- SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_RSA_EXPORT_WITH_RC4_40_MD5
- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_AES_128_CBC_SHA256
- SSL_RSA_FIPS_WITH_AES_256_CBC_SHA256
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_WITH_AES_128_CBC_SHA
- SSL_RSA_WITH_AES_128_CBC_SHA256
- SSL_RSA_WITH_AES_256_CBC_SHA256
- SSL_RSA_WITH_DES_CBC_SHA
- SSL_RSA_WITH_NULL_MD5
- SSL_RSA_WITH_NULL_SHA
- SSL_RSA_WITH_NULL_SHA256
- SSL_RSA_WITH_RC4_128_MD5
- SSL_RSA_WITH_RC4_128_SHA

See also “System requirements for using SHA-2 cipher suites with MQTT channels and clients” on page 1211.

Telemetry channel authentication using SSL

Connections between the MQTT client and the queue manager are always initiated by the MQTT client. The MQTT client is always the SSL client. Client authentication of the server and server authentication of the MQTT client are both optional.

The client always attempts to authenticate the server, unless the client is configured to use a CipherSpec that supports anonymous connection. If the authentication fails, then the connection is not established.

As an alternative to using SSL, some kinds of Virtual Private Network (VPN), such as IPsec, authenticate the endpoints of a TCP/IP connection. VPN encrypts each IP packet that flows over the network. Once such a VPN connection is established, you have established a trusted network. You can connect MQTT clients to telemetry channels using TCP/IP over the VPN network.

Server authentication using SSL authenticates the server to which you are about to send confidential information to. The client performs the checks matching the certificates sent from the server, against certificates placed in its truststore, or in its JRE cacerts store.

The JRE certificate store is a JKS file, cacerts. It is located in JRE InstallPath\lib\security\. It is installed with the default password changeit. You can either store certificates you trust in the JRE certificate store, or in the client truststore. You cannot use both stores. Use the client truststore if you want to keep the public certificates the client trusts separate from certificates other Java applications use. Use the JRE certificate store if you want to use a common certificate store for all Java applications running on the client. If you decide to use the JRE certificate store review the certificates it contains, to make sure you trust them.

You can modify the JSSE configuration by supplying a different trust provider. You can customize a trust provider to perform different checks on a certificate. In some OGSi environments that have used the MQTT client, the environment provides a different trust provider.

To authenticate the WebSphere MQ telemetry channel using SSL, configure the server, and the client.

Telemetry channel configuration

The WebSphere MQ administrator configures telemetry channels at the server. Each channel is configured to accept a TCP/IP connection on a different port number. The channels are configured either as com.ibm.mq.MQTT.channel/PlainText or com.ibm.mq.MQTT.channel/SSL. SSL channels are configured with passphrase protected access to key files. If an SSL channel is defined with no passphrase or key file, the channel does not accept SSL connections.

Store the digital certificate of the server, signed with its private key, in the keystore that the telemetry channel is going to use at the server. Store any certificates in its key chain in the keystore, if you want to transmit the key chain to the client. Configure the telemetry channel using WebSphere MQ explorer to use SSL. Provide it with the path to the keystore, and the passphrase to access the keystore. If you do not set the TCP/IP port number of the channel, the SSL telemetry channel port number defaults to 8883.

MQTT Client configuration

To authenticate the telemetry channel using SSL, the client must connect to the telemetry channel using SSL. It must specify a port that corresponds to a telemetry channel that is configured for SSL. The configuration must include a passphrase protected keystore that contains the privately signed digital certificate of the server. For example, at the client:

```
MQTTClient mqttClient = new MqttClient( "ssl://www.example.org:8884", "clientId");  
mqttClient.connect();
```

The client JVM must use the standard socket factory from JSSE. If you are using Java ME, you must ensure that the JSSE package is loaded. If you are using Java SE, JSSE has been included with the JRE since Java version 1.4.1.

The SSL connection requires a number of SSL properties to be set before connecting. You can set the properties either by passing them to the JVM using the `-D` switch, or you can set the properties using the `MqttConnectionOptions.setSSLProperties` method.

If you load a non-standard socket factory, by calling the method `MqttConnectOptions.setSocketFactory(javax.net.SocketFactory)`, then the way SSL settings are passed to the network socket is application defined.

Code the client to connect to the telemetry channel using SSL, and configure the client to trust a server certificate in one of three ways:

Using a server certificate signed by well-known certificate authority in the cacerts store.

No additional configuration, if the server sends all the intermediate keys in the key-chain. You are advised to review the certificates in the cacerts store of the client JRE and change the password to the cacerts store

Other certificates

Store the certificates you trust in the truststore at the client. You must store at least one of the certificates in the key-chain in the truststore, Set the truststore parameters in `MqttConnectionOptions.SSLProperty`.

- `com.ibm.ssl.trustStore`
- `com.ibm.ssl.trustStorePassword`

Using a custom trust manager

Implement a trust provider and pass it the name of the algorithm that is used. Set the name of the provider class and the algorithm to use in `MqttConnectionOptions.SSLProperty`.

- `com.ibm.ssl.trustStoreProvider`
- `com.ibm.ssl.trustStoreManager`

You can also use SSL cipher suites for client authentication. Here is an alphabetic list of the SSL cipher suites that are currently supported:

- `SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA`
- `SSL_DH_anon_EXPORT_WITH_RC4_40_MD5`
- `SSL_DH_anon_WITH_3DES_EDE_CBC_SHA`
- `SSL_DH_anon_WITH_AES_128_CBC_SHA`
- `SSL_DH_anon_WITH_DES_CBC_SHA`
- `SSL_DH_anon_WITH_RC4_128_MD5`
- `SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA`
- `SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA`
- `SSL_DHE_DSS_WITH_AES_128_CBC_SHA`
- `SSL_DHE_DSS_WITH_DES_CBC_SHA`
- `SSL_DHE_DSS_WITH_RC4_128_SHA`
- `SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA`
- `SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA`
- `SSL_DHE_RSA_WITH_AES_128_CBC_SHA`
- `SSL_DHE_RSA_WITH_DES_CBC_SHA`
- `SSL_KRB5_EXPORT_WITH_DES_CBC_40_MD5`

- SSL_KRB5_EXPORT_WITH_DES_CBC_40_SHA
- SSL_KRB5_EXPORT_WITH_RC4_40_MD5
- SSL_KRB5_EXPORT_WITH_RC4_40_SHA
- SSL_KRB5_WITH_3DES_EDE_CBC_MD5
- SSL_KRB5_WITH_3DES_EDE_CBC_SHA
- SSL_KRB5_WITH_DES_CBC_MD5
- SSL_KRB5_WITH_DES_CBC_SHA
- SSL_KRB5_WITH_RC4_128_MD5
- SSL_KRB5_WITH_RC4_128_SHA
- SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_RSA_EXPORT_WITH_RC4_40_MD5
- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_AES_128_CBC_SHA256
- SSL_RSA_FIPS_WITH_AES_256_CBC_SHA256
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_WITH_AES_128_CBC_SHA
- SSL_RSA_WITH_AES_128_CBC_SHA256
- SSL_RSA_WITH_AES_256_CBC_SHA256
- SSL_RSA_WITH_DES_CBC_SHA
- SSL_RSA_WITH_NULL_MD5
- SSL_RSA_WITH_NULL_SHA
- SSL_RSA_WITH_NULL_SHA256
- SSL_RSA_WITH_RC4_128_MD5
- SSL_RSA_WITH_RC4_128_SHA

See also “System requirements for using SHA-2 cipher suites with MQTT channels and clients” on page 1211.

Publication privacy on telemetry channels

The privacy of MQTT publications sent in either direction across telemetry channels is secured by using SSL to encrypt transmissions over the connection.

MQTT clients that connect to telemetry channels use SSL to secure the privacy of publications transmitted on the channel using symmetric key cryptography. Because the endpoints are not authenticated, you cannot trust channel encryption alone. Combine securing privacy with server or mutual authentication.

As an alternative to using SSL, some kinds of Virtual Private Network (VPN), such as IPsec, authenticate the endpoints of a TCP/IP connection. VPN encrypts each IP packet that flows over the network. Once such a VPN connection is established, you have established a trusted network. You can connect MQTT clients to telemetry channels using TCP/IP over the VPN network.

For a typical configuration, which encrypts the channel and authenticates the server, consult “Telemetry channel authentication using SSL” on page 137.

Encrypting SSL connections without authenticating the server exposes the connection to man-in-the-middle attacks. Although the information you exchange is protected against eavesdropping, you do not know who you are exchanging it with. Unless you control the network, you are exposed to someone intercepting your IP transmissions, and masquerading as the endpoint.

You can create an encrypted SSL connection, without authenticating the server, by using a Diffie-Hellman key exchange CipherSpec that supports anonymous SSL. The master secret, shared between the client and server, and used to encrypt SSL transmissions, is established without exchanging a privately signed server certificate.

Because anonymous connections are insecure, most SSL implementations do not default to using anonymous CipherSpecs. If a client request for SSL connection is accepted by a telemetry channel, the channel must have a keystore protected by a passphrase. By default, since SSL implementations do not use anonymous CipherSpecs, the keystore must contain a privately signed certificate that the client can authenticate.

If you use anonymous CipherSpecs, the server keystore must exist, but it need not contain any privately signed certificates.

Another way to establish an encrypted connection is to replace the trust provider at the client with your own implementation. Your trust provider would not authenticate the server certificate, but the connection would be encrypted.

SSL configuration of MQTT clients and telemetry channels

MQTT clients and the WebSphere MQ Telemetry (MQXR) service use Java Secure Socket Extension (JSSE) to connect telemetry channels using SSL. The WebSphere MQ Telemetry daemon for devices does not support SSL.

Configure SSL to authenticate the telemetry channel, the MQTT client, and encrypt the transfer of messages between clients and the telemetry channel.

As an alternative to using SSL, some kinds of Virtual Private Network (VPN), such as IPsec, authenticate the endpoints of a TCP/IP connection. VPN encrypts each IP packet that flows over the network. Once such a VPN connection is established, you have established a trusted network. You can connect MQTT clients to telemetry channels using TCP/IP over the VPN network.

You can configure the connection between a Java MQTT client and a telemetry channel to use the SSL protocol over TCP/IP. What is secured depends on how you configure SSL to use JSSE. Starting with the most secured configuration, you can configure three different levels of security:

1. Permit only trusted MQTT clients to connect. Connect an MQTT client only to a trusted telemetry channel. Encrypt messages between the client and the queue manager; see “MQTT client authentication using SSL” on page 133
2. Connect an MQTT client only to a trusted telemetry channel. Encrypt messages between the client and the queue manager; see “Telemetry channel authentication using SSL” on page 137.
3. Encrypt messages between the client and the queue manager; see “Publication privacy on telemetry channels” on page 139.

JSSE configuration parameters

Modify JSSE parameters to alter the way an SSL connection is configured. The JSSE configuration parameters are arranged into three sets:

1. WebSphere MQ Telemetry channel
2. MQTT Java client
3. JRE

Configure the telemetry channel parameters using WebSphere MQ Explorer. Set the MQTT Java Client parameters in the `MqttConnectionOptions.SSLProperties` attribute. Modify JRE security parameters by editing files in the JRE security directory on both the client and server.

WebSphere MQ Telemetry channel

Set all the telemetry channel SSL parameters using WebSphere MQ Explorer.

ChannelName

ChannelName is a required parameter on all channels.

The channel name identifies the channel associated with a particular port number. Name channels to help you administer sets of MQTT clients.

PortNumber

PortNumber is an optional parameter on all channels. It defaults to 1883 for TCP channels, and 8883 for SSL channels.

The TCP/IP port number associated with this channel. MQTT clients are connected to a channel by specifying the port defined for the channel. If the channel has SSL properties, the client must connect using the SSL protocol; for example:

```
MQTTClient mqttClient = new MqttClient( "ssl://www.example.org:8884", "clientId");
mqttClient.connect();
```

KeyFileName

KeyFileName is a required parameter for SSL channels. It must be omitted for TCP channels.

KeyFileName is the path to the Java keystore containing digital certificates that you provide. Use JKS, JCEKS or PKCS12 as the type of keystore on the server.

Identify the keystore type by using one of the following file extensions:

- .jks
- .jceks
- .p12
- .pkcs12

A keystore with any other file extension is assumed to be a JKS keystore.

You can combine one type of keystore at the server with other types of keystore at the client.

Place the private certificate of the server in the keystore. The certificate is known as the server certificate. The certificate can be self-signed, or part of a key-chain.

If you are using a key-chain, place the key-chain certificates in the server keystore.

The server certificate, and any certificates in its certificate chain, are sent to clients to authenticate the identity of the server.

If you have set ClientAuth to Required, the keystore must contain any certificates necessary to authenticate the client. Place trusted public self-signed certificates of clients, or trusted certificates in the key-chain of clients in the keystore. The first verification of a self-signed certificate, or key-chain, sent by a client against a certificate in the keystore, authenticates the client. Using a key-chain, one certificate can verify many clients, even if they are issued with different client certificates.

PassPhrase

PassPhrase is a required parameter for SSL channels. It must be omitted for TCP channels.

The passphrase is used to protect the keystore.

ClientAuth

`ClientAuth` is an optional SSL parameter. It defaults to no client authentication. It must be omitted for TCP channels.

Set `ClientAuth` if you want the telemetry service to authenticate the client, before permitting the client to connect to the telemetry channel.

If you set `ClientAuth`, the client must connect to the server using SSL, and authenticate the server. In response to setting `ClientAuth`, the client sends its digital certificate to the server, and any other certificates in its keystore. Its digital certificate is known as the client certificate. These certificates are authenticated against certificates held in the channel keystore, and in the JRE cacerts store.

CipherSuite

`CipherSuite` is an optional SSL parameter. It defaults to try all the enabled `CipherSpecs`. It must be omitted for TCP channels.

If you want to use a particular `CipherSpec`, set `CipherSuite` to the name of the `CipherSpec` that must be used to establish the SSL connection.

The telemetry service and MQTT client negotiate a common `CipherSpec` from all the `CipherSpecs` that are enabled at each end. If a specific `CipherSpec` is specified at either or both ends of the connection, it must match the `CipherSpec` at the other end.

Install additional ciphers by adding additional providers to JSSE.

Federal Information Processing Standards (FIPS)

FIPS is an optional setting. By default it is not set.

Either in the properties panel of the queue manager, or using `runmqsc`, set `SSLFIPS`. `SSLFIPS` specifies whether only FIPS-certified algorithms are to be used.

Revocation namelist

Revocation namelist is an optional setting. By default it is not set.

Either in the properties panel of the queue manager, or using `runmqsc`, set `SSLCRLNL`. `SSLCRLNL` specifies a namelist of authentication information objects which are used to provide certificate revocation locations.

No other queue manager parameters that set SSL properties are used.

MQTT Java client

Set SSL properties for the Java client in `MqttConnectionOptions.SSLProperties`; for example:

```
java.util.Properties sslClientProperties = new Properties();
sslClientProperties.setProperty("com.ibm.ssl.keyStoreType", "JKS");
com.ibm.micro.client.mqttv3.MqttConnectOptions conOptions = new MqttConnectOptions();
conOptions.setSSLProperties(sslClientProperties);
```

The names and values of specific properties are described in the `MqttConnectOptions` class. For links to client API documentation for the MQTT version 3.1 protocol, see ["MQTT client programming reference"](#) in the WebSphere MQ Version 7.5 product documentation.

Protocol

Protocol is optional.

The protocol is selected in negotiation with the telemetry server. If you require a specific protocol you can select one. If the telemetry server does not support the protocol the connection fails.

ContextProvider

`ContextProvider` is optional.

KeyStore

KeyStore is optional. Configure it if ClientAuth is set at the server to force authentication of the client.

Place the digital certificate of the client, signed using its private key, into the keystore. Specify the keystore path and password. The type and provider are optional. JKS is the default type, and IBMJCE is the default provider.

Specify a different keystore provider to reference a class that adds a new keystore provider. Pass the name of the algorithm used by the keystore provider to instantiate the KeyManagerFactory by setting the key manager name.

TrustStore

TrustStore is optional. You can place all the certificates you trust in the JRE cacerts store.

Configure the truststore if you want to have a different truststore for the client. You might not configure the truststore if the server is using a certificate issued by a well known CA that already has its root certificate stored in cacerts.

Add the publicly signed certificate of the server or the root certificate to the truststore, and specify the truststore path and password. JKS is the default type, and IBMJCE is the default provider.

Specify a different truststore provider to reference a class that adds a new truststore provider. Pass the name of the algorithm used by the truststore provider to instantiate the TrustManagerFactory by setting the trust manager name.

JRE

Other aspects of Java security that affect the behavior of SSL on both the client and server are configured in the JRE. The configuration files on Windows are in *Java Installation Directory*\jre\lib\security. If you are using the JRE shipped with IBM WebSphere MQ the path is as shown in the following table:

Platform	Filepath
Windows	<i>WMQ Installation Directory</i> \java\jre\lib\security
Linux for System x 32 bit	<i>WMQ Installation Directory</i> /java/jre/lib/security
Other UNIX and Linux platforms	<i>WMQ Installation Directory</i> /java/jre64/jre/lib/security

Well-known certificate authorities

The cacerts file contains the root certificates of well-known certificate authorities. The cacerts is used by default, unless you specify a truststore. If you use the cacerts store, or do not provide a truststore, you must review and edit the list of signers in cacerts to meet your security requirements.

You can open cacerts using the WebSphere MQ command `strmqikm`, which runs the IBM Key Management utility. Open cacerts as a JKS file, using the password `changeit`. Modify the password to secure the file.

Configuring security classes

Use the `java.security` file to register additional security providers and other default security properties.

Permissions

Use the `java.policy` file to modify the permissions granted to resources. `javaws.policy` grants permissions to `javaws.jar`

Encryption strength

Some JREs ship with reduced strength encryption. If you cannot import keys into keystores, reduced strength encryption might be the cause. Either, try starting **ikeyman** using the **strmqikm** command, or download strong, but limited jurisdiction files from

[IBM developer kits, Security information.](#)

Important: Your country of origin might have restrictions on the import, possession, use, or re-export to another country, of encryption software. Before downloading or using the unrestricted policy files, you must check the laws of your country. Check its regulations, and its policies concerning the import, possession, use, and re-export of encryption software, to determine if it is permitted.

Modify the trust provider to permit the client to connect to any server

The example illustrates how to add a trust provider and reference it from the MQTT client code. The example performs no authentication of the client or server. The resulting SSL connection is encrypted without being authenticated.

The code snippet in Figure 24 sets the AcceptAllProviders trust provider and trust manager for the MQTT client.

```
java.security.Security.addProvider(new AcceptAllProvider());
java.util.Properties sslClientProperties = new Properties();
sslClientProperties.setProperty("com.ibm.ssl.trustManager", "TrustAllCertificates");
sslClientProperties.setProperty("com.ibm.ssl.trustStoreProvider", "AcceptAllProvider");
conOptions.setSSLProperties(sslClientProperties);
```

Figure 24. MQTT Client code snippet

```
package com.ibm.mq.id;
public class AcceptAllProvider extends java.security.Provider {
    private static final long serialVersionUID = 1L;
    public AcceptAllProvider() {
        super("AcceptAllProvider", 1.0, "Trust all X509 certificates");
        put("TrustManagerFactory.TrustAllCertificates",
            AcceptAllTrustManagerFactory.class.getName());
    }
}
```

Figure 25. AcceptAllProvider.java

```
protected static class AcceptAllTrustManagerFactory extends
    javax.net.ssl.TrustManagerFactorySpi {
    public AcceptAllTrustManagerFactory() {}
    protected void engineInit(java.security.KeyStore keystore) {}
    protected void engineInit(
        javax.net.ssl.ManagerFactoryParameters parameters) {}
    protected javax.net.ssl.TrustManager[] engineGetTrustManagers() {
        return new javax.net.ssl.TrustManager[] { new AcceptAllX509TrustManager() };
    }
}
```

Figure 26. AcceptAllTrustManagerFactory.java

```

protected static class AcceptAllX509TrustManager implements
    javax.net.ssl.X509TrustManager {
public void checkClientTrusted(
    java.security.cert.X509Certificate[] certificateChain,
    String authType) throws java.security.cert.CertificateException {
    report("Client authtype=" + authType);
    for (java.security.cert.X509Certificate certificate : certificateChain) {
        report("Accepting:" + certificate);
    }
}
public void checkServerTrusted(
    java.security.cert.X509Certificate[] certificateChain,
    String authType) throws java.security.cert.CertificateException {
    report("Server authtype=" + authType);
    for (java.security.cert.X509Certificate certificate : certificateChain) {
        report("Accepting:" + certificate);
    }
}
public java.security.cert.X509Certificate[] getAcceptedIssuers() {
    return new java.security.cert.X509Certificate[0];
}
    private static void report(String string) {
        System.out.println(string);
    }
}
}

```

Figure 27. *AcceptAllX509TrustManager.java*

Telemetry channel JAAS configuration

Configure JAAS to authenticate the Username sent by the client.

The WebSphere MQ administrator configures which MQTT channels require client authentication using JAAS. Specify the name of a JAAS configuration for each channel that is to perform JAAS authentication. Channels can all use the same JAAS configuration, or they can use different JAAS configurations. The configurations are defined in *WMQData directory\qmgrs\qMgrName\mqxr\jaas.config*.

The *jaas.config* file is organized by JAAS configuration name. Under each configuration name is a list of Login configurations; see Figure 28 on page 146.

JAAS provides four standard Login modules. The standard NT and UNIX Login modules are of limited value.

JndiLoginModule

Authenticates against a directory service configured under JNDI (Java Naming and Directory Interface).

Krb5LoginModule

Authenticates using Kerberos protocols.

NTLoginModule

Authenticates using the NT security information for the current user.

UnixLoginModule

Authenticates using the UNIX security information for the current user.

The problem with using *NTLoginModule* or *UnixLoginModule* is that the telemetry service runs with the *mqm* identity, and not the identity of the MQTT channel. *mqm* is the identity passed to *NTLoginModule* or *UnixLoginModule* for authentication, and not the identity of the client.

To overcome this problem, write your own Login module, or use the other standard Login modules. A sample *JAASLoginModule.java* is supplied with WebSphere MQ Telemetry. It is an implementation of the *javax.security.auth.spi.LoginModule* interface. Use it to develop your own authentication method.

Any new LoginModule classes you provide must be on the class path of the telemetry service. Do not place your classes in WebSphere MQ directories that are in the class path. Create your own directories, and define the whole class path for the telemetry service.

You can augment the class path used by the telemetry service by setting class path in the `service.env` file. `CLASSPATH` must be capitalized, and the class path statement can only contain literals. You cannot use variables in the `CLASSPATH`; for example `CLASSPATH=%CLASSPATH%` is incorrect. The telemetry service sets its own classpath. The `CLASSPATH` defined in `service.env` is added to it.

The telemetry service provides two callbacks that return the Username and the Password for a client connected to the MQTT channel. The Username and Password are set in the `MqttConnectOptions` object. See Figure 29 for an example of how to access Username and Password.

Examples

An example of a JAAS configuration file with one named configuration, `MQXRConfig`.

```
MQXRConfig {
  samples.JAASLoginModule required debug=true;
  //com.ibm.security.auth.module.NTLoginModule required;
  //com.ibm.security.auth.module.Krb5LoginModule required
  //      principal=principal@your_realm
  //      useDefaultCcache=TRUE
  //      renewTGT=true;
  //com.sun.security.auth.module.NTLoginModule required;
  //com.sun.security.auth.module.UnixLoginModule required;
  //com.sun.security.auth.module.Krb5LoginModule required
  //      useTicketCache="true"
  //      ticketCache="${user.home}${/}tickets";
};
```

Figure 28. Sample `jaas.config` file

An example of a JAAS Login module coded to receive the Username and Password provided by an MQTT client.

```
public boolean login()
  throws javax.security.auth.login.LoginException {
  javax.security.auth.callback.Callback[] callbacks =
    new javax.security.auth.callback.Callback[2];
  callbacks[0] = new javax.security.auth.callback.NameCallback("NameCallback");
  callbacks[1] = new javax.security.auth.callback.PasswordCallback(
    "PasswordCallback", false);
  try {
    callbackHandler.handle(callbacks);
    String username = ((javax.security.auth.callback.NameCallback) callbacks[0])
      .getName();
    char[] password = ((javax.security.auth.callback.PasswordCallback) callbacks[1])
      .getPassword(); // Accept everything.
    if (true) {
      loggedIn = true;
    } else {
      throw new javax.security.auth.login.LoginException("Login failed");
    }
    principal = new JAASPrincipal(username);
  } catch (java.io.IOException exception) {
    throw new javax.security.auth.login.LoginException(exception.toString());
  } catch (javax.security.auth.callback.UnsupportedCallbackException exception) {
    throw new javax.security.auth.login.LoginException(exception.toString());
  }
  return loggedIn;
}
```

Figure 29. Sample `JAASLoginModule.Login()` method

WebSphere MQ Telemetry daemon for devices concepts

The WebSphere MQ Telemetry daemon for devices is an advanced MQTT V3 client application. Use it to store and forward messages from other MQTT clients. It connects to WebSphere MQ like an MQTT client, but you can also connect other MQTT clients to it.

The daemon is a publish/subscribe broker. MQTT V3 clients connect to it to publish and subscribe to topics, using topic strings to publish, and topic filters to subscribe. The topic string is hierarchical, with topic levels divided by /. Topic filters are topic strings that can include single level + wildcards and a multilevel # wildcard as the last part of the topic string.

Note: Wildcards in the daemon follow the more restrictive rules of WebSphere Message Broker, v6. WebSphere MQ is different. It supports multiple multilevel wildcards; wildcards can stand in for any number of levels of the hierarchy, anywhere in the topic string.

Multiple MQTT v3 clients connect to the daemon using a listener port. The default listener port is modifiable. You can define multiple listener ports and allocate different namespaces to them, see “WebSphere MQ Telemetry daemon for devices listener ports” on page 155. The daemon is itself an MQTT v3 client. Configure a daemon bridge connection to connect the daemon to the listener port of another daemon, or to a WebSphere MQ Telemetry service.

You can configure multiple bridges for the WebSphere MQ Telemetry daemon for devices. Use the bridges to connect together a network of daemons that can exchange publications.

Each bridge can publish and subscribe to topics at its local daemon. It can also publish and subscribe to topics at another daemon, a WebSphere MQ publish/subscribe broker, or any other MQTT v3 broker it is connected to. Using a topic filter, you can select the publications to propagate from one broker to another. You can propagate publications in either direction. You can propagate publications from the local daemon to each of its attached remote brokers, or from any of the attached brokers to the local daemon; see “WebSphere MQ Telemetry daemon for devices bridges.”

WebSphere MQ Telemetry daemon for devices bridges

A WebSphere MQ Telemetry daemon for devices bridge connects two publish/subscribe brokers using the MQTT v3 protocol. The bridge propagates publications from one broker to the other, in either direction. At one end is a WebSphere MQ Telemetry daemon for devices bridge connection, and at the other might be a queue manager, or another daemon. A queue manager is connected to the bridge connection using a telemetry channel. A daemon is connected to the bridge connection using a daemon listener.

WebSphere MQ Telemetry daemon for devices supports one or more simultaneous connections to other brokers. The connections from the daemon are called bridges and are defined by connection entries in the daemon configuration file. The connections to WebSphere MQ are made using WebSphere MQ telemetry channels; see Figure 30 on page 148.

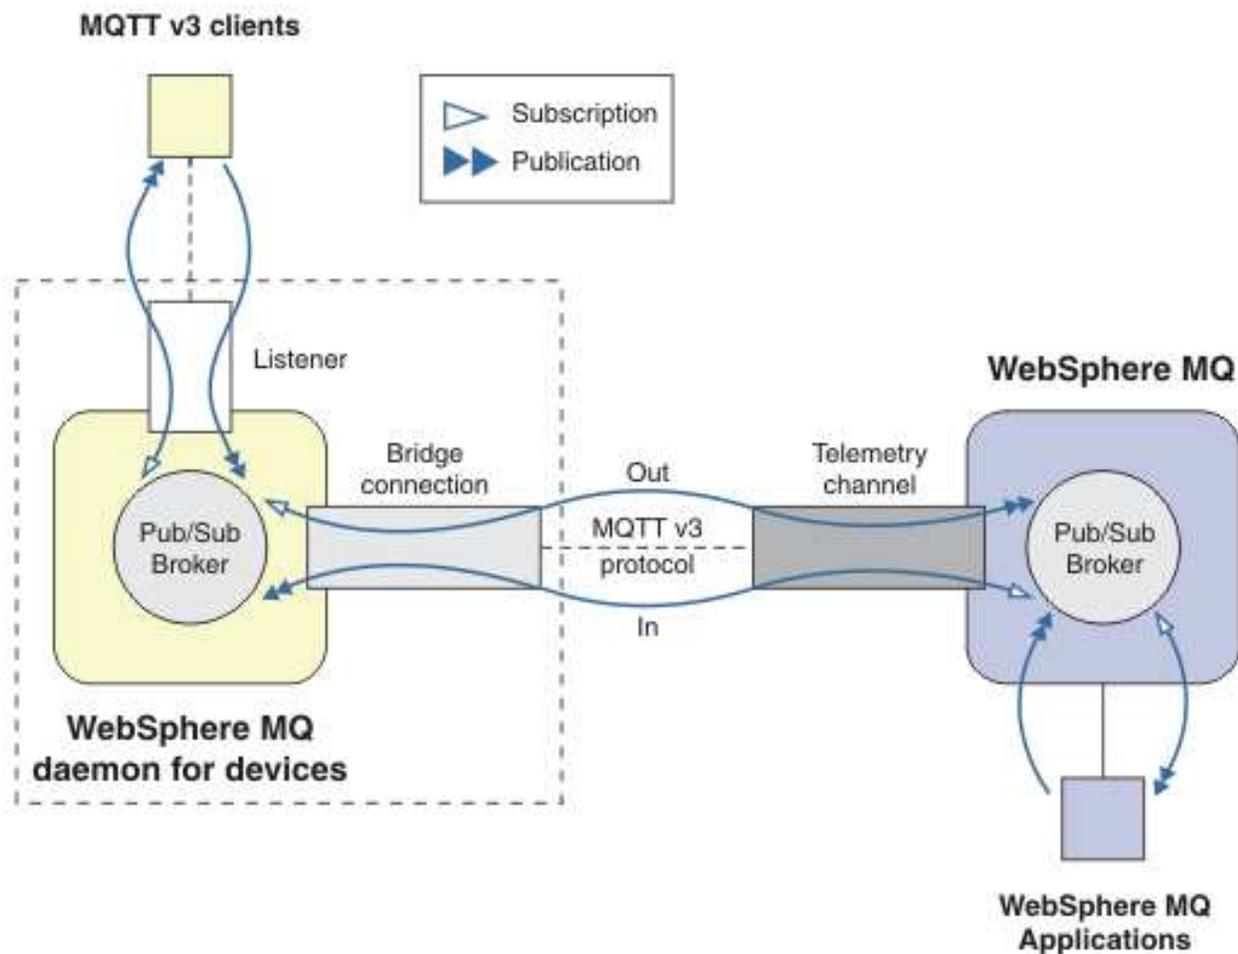


Figure 30. Connecting WebSphere MQ Telemetry daemon for devices to WebSphere MQ

A bridge connects the daemon to another broker as an MQTT v3 client. The bridge parameters mirror the attributes of an MQTT v3 client.

A bridge is more than a connection. It acts as a publish and subscribe agent situated between two publish/subscribe brokers. The local broker is the WebSphere MQ Telemetry daemon for devices, and the remote broker is any publish/subscribe broker that supports the MQTT v3 protocol. Typically the remote broker is another daemon or WebSphere MQ.

The job of the bridge is to propagate publications between the two brokers. The bridge is bidirectional. It propagates publications in either direction. Figure 30 illustrates the way the bridge connects WebSphere MQ Telemetry daemon for devices to WebSphere MQ. “Example topic settings for the bridge” on page 149 uses examples to illustrate how to use the topic parameter to configure the bridge.

The In and Out arrows in Figure 30 indicate the bidirectionality of the bridge. At one end of the arrow, a subscription is created. The publications that match the subscription are published to the broker at the opposite end of the arrow. The arrow is labeled according to the flow of publications. Publications flow In to the daemon and Out from the daemon. The importance of the labels is they are used in the command syntax. Remember that In and Out refer to where the publications flow, and not to where the subscription is sent.

Other clients, applications, or brokers might be connected either to WebSphere MQ or to WebSphere MQ Telemetry daemon for devices. They publish and subscribe to topics at the broker they are connected to. If the broker is WebSphere MQ, the topics might be clustered or distributed, and not explicitly defined at the local queue manager.

Uses of bridges

Connect daemons together using bridge connections and listeners. Connect daemons and queue managers together using bridge connections and telemetry channels. When you connect multiple brokers together it is possible to create loops. Be careful: Publications might circulate endlessly around a loop of brokers, undetected.

Some of the reasons for using daemons bridged to WebSphere MQ are as follows:

Reduce the number of MQTT client connections to WebSphere MQ

Using a hierarchy of daemons you can connect many clients to WebSphere MQ; more clients than the number a single queue manager can connect at one time.

Store and forward messages between MQTT clients and WebSphere MQ

You might use store and forward to avoid maintaining continuous connections between clients and WebSphere MQ, if the clients do not have their own storage. You might use multiple types of connections between the MQTT client and WebSphere MQ; see  Telemetry concepts and scenarios for monitoring and control (*WebSphere MQ V7.1 Product Overview Guide*).

Filter the publications exchanged between MQTT clients and WebSphere MQ

Commonly, publications divide into messages that are processed locally and messages that involve other applications. Local publications might include control flows between sensors and actuators, and remote publications include requests for readings, status, and configuration commands.

Change the topic spaces of publications

Avoid topic strings from clients attached to different listener ports from colliding with one another. The example uses the daemon to label meter readings coming from different buildings; see *Separating the topic spaces of different groups of clients*.

Example topic settings for the bridge

Publish everything to the remote broker - using defaults

The default direction is called out, and the bridge publishes topics to the remote broker. The topic parameter controls what topics are propagated using topic filters.

The bridge uses the topic parameter in Figure 31 to subscribe to everything published to the local daemon by MQTT clients, or by other brokers. The bridge publishes the topics to the remote broker connected by the bridge.

Figure 31. Publish everything to the remote broker

```
connection Daemon1
topic #
```

Publish everything to the remote broker - explicit

The topic setting in Figure 32 on page 150 gives the same result. The only difference is direction parameter is explicit. Use the out direction to subscribe to the local broker, the daemon, and publish to the remote broker. Publications created on the local daemon that the bridge has subscribed to, are published at the remote broker.

Figure 32. Publish everything to the remote broker - explicit

```
connection Daemon1
topic # out
```

Publish everything to the local broker

Instead of using the direction, out, you can set the opposite direction, in. Figure 33 configures the bridge to subscribe to everything published at the remote broker connected by the bridge. The bridge publishes the topics to the local broker, the daemon.

Figure 33. Publish everything to the local broker

```
connection Daemon1
topic # in
```

Publish everything from the export topic at the local broker to the import topic at the remote broker

Use two additional topic parameters, *local_prefix* and *remote_prefix*, to modify the topic filter, # in the previous examples. One parameter is used to modify the topic filter used in the subscription, and the other parameter is used to modify the topic the publication is published to. The effect is to replace the beginning of the topic string used in one broker with another topic string on the other broker.

Depending on the direction of the topic command the meaning of *local_prefix* and *remote_prefix* reverses. If the direction is out, the default, *local_prefix* is used as part of the topic subscription, and *remote_prefix* replaces the *local_prefix* part of the topic string in the remote publication. If the direction is in, *remote_prefix* becomes part of the remote subscription, and *local_prefix* replaces the *remote_prefix* part of the topic string.

The first part of a topic string is often thought of as defining a topic space. Use the additional parameters to change the topic space a topic is published to. You might do this to avoid the topic being propagated colliding with another the topic on the target broker, or to remove a mount point topic string.

As an example, in Figure 34, all the publications to the topic string export/# at the daemon are republished to import/# at the remote broker.

Figure 34. Publish everything from the export topic at the local broker to the import topic at the remote broker

```
topic # out export/ import/
```

Publish everything to the import topic at the local broker from the export topic at the remote broker

Figure 35 shows the configuration reversed; the bridge subscribes to everything published with the export/# topic string at the remote broker and publishes it to import/# at the local broker.

Figure 35. Publish everything to the import topic at the local broker from the export topic at the remote broker

```
connection Daemon1
topic # in import/ export/
```

Publish everything from the 1884/ mount point to the remote broker with the original topic strings

In the example in Figure 36 on page 151 the bridge subscribes to everything published by clients connected to the mount point 1884/ at the local daemon. The bridge publishes everything published to the mount point to the remote broker. The mount point string 1884/ is removed from the topics published to the remote broker. The *local_prefix* is the same as the mount point string 1884/, and the *remote_prefix* is a blank string.

Figure 36. Publish everything from the 1884/ mount point to the remote broker with the original topic strings.

```
listener 1884
mount_point 1884/
connection Daemon1
topic # out 1884/ ""
```

Separating the topic spaces of different clients connected to different daemons

An application is written for electrical power meters to publish meter readings for a building. The readings are published using MQTT clients to a daemon hosted in the same building. The topic selected for the publications is power. The same application is deployed to a number of buildings in a complex. For site monitoring and data storage, readings from all buildings are aggregated using bridge connections. The connections link the building daemons to WebSphere MQ at a central location.

The client applications in each building are identical, but the data must be differentiated by building. Each reading has a power topic and must be prefixed with the building number to distinguish it. The bridge from the first building in the complex uses the prefix meters/building01/, from building two the prefix is meters/building02/. The readings from the other buildings follow the same pattern. WebSphere MQ receives the readings with topics like meters/building01/power.

The example is contrived; in practice the topic space the application publishes to is likely to be configurable.

The configuration file for each daemon has a topic statement that follows the pattern in Figure 37.

Figure 37. Separate the topic spaces of clients connected to different daemons

```
connection Daemon1
topic power out "" meters/building01/
```

Specify an empty string as a placeholder for the unused local_prefix parameter.

Separate the topic spaces of clients connected to the same daemon

Suppose that a single daemon is used to connect all the power meters. Assuming that in the application can be configured to connect to different ports, you might distinguish the buildings by attaching the meters from different buildings to different listener ports; see Figure 38. Again, the example is contrived; it illustrates how mount points might be used.

Figure 38. Separate the topic spaces of clients connected to the same daemon

```
listener 1884
mount_point meters/building01/
listener 1885
mount_point meters/building02/
connection Daemon1
topic meters/+ /power out
```

Remap different topics for publications flowing in both directions

In the configuration in Figure 39 on page 152 the bridge subscribes to the single topic b at the remote broker and forwards publications about b to the local daemon, changing the topic to a. The bridge also subscribes to the single topic x at the local broker and forwards publications about x to the remote broker, changing the topic to y.

Figure 39. Remap different topics for publications flowing in both directions

```
connection Daemon1
topic "" in a b
topic "" out x y
```

An important point about this example is that different topics are subscribed to and published to at both brokers. The topics spaces at both brokers are disjoint.

Remap the same topics for publications flowing in both directions (looping)

Unlike the previous example, the configuration in Figure 40, in general, results in a loop. In the topic statement `topic "" in a b`, the bridge subscribes to `b` remotely, and publishes to `a` locally. In the other topic statement, the bridge subscribes to `a` locally, and publishes to `b` remotely. The same configuration can be written as Figure 41.

The general result is that if a client publishes to `b` remotely, the publication is transferred to the local daemon as a publication on topic `a`. However, on being published by the bridge to the local daemon on the topic `a`, the publication matches the subscription made by the bridge to local topic `a`. The subscription is `topic "" out a b`. As a result, the publication is transferred back to the remote broker as a publication on topic `b`. The bridge is now subscribed to the remote topic `b`, and the cycle begins again.

Some brokers implement loop detection to prevent the loop happening. But the loop detection mechanism must work when different types of brokers are bridged together. Loop detection does not work if WebSphere MQ is bridged to the WebSphere MQ Telemetry daemon for devices. It does work if two WebSphere MQ Telemetry daemon for devices are bridged together. By default loop detection is turned on; see  `try_private`.

Figure 40. !Remap the same topics for publications flowing in both directions

```
connection Daemon1
topic "" in a b
topic "" out a b
```

Figure 41. !Remap the same topics for publications flowing in both directions, using both.

```
connection Daemon1
topic "" both a b
```

The configuration in Figure 39 is the same as Figure 40.

Availability of IBM WebSphere MQ Telemetry daemon for devices bridge connections

Configure multiple IBM WebSphere MQ Telemetry daemon for devices bridge connection addresses to connect to the first available remote broker. If the broker is a multi-instance queue manager, provide both of its TCP/IP addresses. Configure a primary connection to connect, or reconnect, to the primary server, when it is available.

The connection bridge parameter,  `addresses`, is a list of TCP/IP socket addresses. The bridge attempts to connect to each address in turn, until it makes a successful connection. The  `round_robin` and  `start_type` connection parameters control how the addresses are used once a successful connection has been made.

If `start_type` is `auto`, `manual`, or `lazy`, then if the connection fails, the bridge attempts to reconnect. It uses each address in turn, with about a 20 second delay between each connection attempt. If `start_type` is `once`, then if the connection fails, the bridge does not attempt to reconnect automatically.

If `round_robin` is `true`, the bridge connection attempts start at the first address in the list and tries each address in the list in turn. It starts at the first address again, when the list is exhausted. If there is only one address in the list, it tries it again every 20 seconds.

If `round_robin` is `false`, the first address in the list, which is called the primary server, is given preference. If the first attempt to connect to the primary server fails, the bridge continues to try to reconnect to the primary server in the background. At the same time, the bridge tries to connect using the other addresses in the list. When the background attempts to connect to the primary server succeed, the bridge disconnects from the current connection, and switches to the primary server connection.

If a connection is disconnected voluntarily, for example by issuing a `connection_stop` command, then if the connection is restarted, it tries to use the same address again. If the connection is disconnected due to a failure to connect, or to the remote broker dropping the connection, the bridge waits 20 seconds. It then tries to connect to the next address in the list, or the same address, if there is only one address in the list.

Connecting to a multi-instance queue manager

In a multi-instance queue manager configuration, the queue manager runs on two different servers with different IP addresses. Typically telemetry channels are configured without a specific IP address. They are configured only with a port number. When the telemetry channel is started, by default it selects the first available network address on the local server.

Configure the `addresses` parameter of the bridge connection with the two IP addresses used by the queue manager. Set `round_robin` to `true`.

If the active queue manager instance fails, the queue manager switches over to the standby instance. The daemon detects that the connection to the active instance has broken and tries to reconnect to the standby instance. It uses the other IP address in the list of addresses configured for the bridge connection.

The queue manager to which the bridge connects is still the same queue manager. The queue manager recovers its own state. If `cleansession` is set to `false`, the bridge connection session is restored to the same state as before the failover. The connection resumes after a delay. Messages with “at least once” or “at most once” quality of service are not lost, and subscriptions continue to work.

The reconnection time depends on the number of channels and clients that restart when the standby instance starts, and how many messages were inflight. The bridge connection might try to reconnect to both IP addresses a number of times before the connection is reestablished.

Do not configure a multi-instance queue manager telemetry channel with a specific IP address. The IP address is only valid on one server.

If you are using an alternative high-availability solution, that manages the IP address, then it might be correct to configure a telemetry channel with a specific IP address.

cleansession

A bridge connection is an MQTT v3 client session. You can control whether a connection starts a new session, or whether it restores an existing session. If it restores an existing session, the bridge connection preserves the subscriptions and retained publications from the previous session.

Do not set `cleansession` to `false` if `addresses` lists multiple IP addresses, and the IP addresses connect to telemetry channels hosted by different queue managers, or to different telemetry daemons. Session state

is not transferred between queue managers or daemons. Trying to restart an existing session on a different queue manager or daemon results in a new session being started. In-doubt messages are lost, and subscriptions might not behave as expected.

notifications

An application can keep track of whether the bridge connection is running by using notifications. A notification is a publication that has the value 1, connected, or 0, disconnected. It is published to

topicString defined by the  *notification_topic* parameter. The default value of *topicString* is **\$\$SYS/broker/connection/clientIdentifier/state**. The default *topicString* contains the prefix **\$\$SYS**. Subscribe to topics beginning with **\$\$SYS** by defining a topic filter beginning with **\$\$SYS**. The topic filter **#**, subscribe to everything, does not subscribe to topics beginning with **\$\$SYS** on the daemon. Think of **\$\$SYS** as defining a special system topic space distinct from the application topic space.

Notifications enable IBM WebSphere MQ Telemetry daemon for devices to notify MQTT clients when a bridge is connected or disconnected.

keepalive_interval

The *keepalive_interval* bridge connection parameter sets the interval between the bridge sending a TCP/IP ping to the remote server. The default interval is 60 seconds. The ping prevents the TCP/IP session being closed by the remote server, or by a firewall, that detects a period of inactivity on the connection.

clientid

A bridge connection is an MQTT v3 client session and has a *clientIdentifier* that is set by the bridge connection parameter *clientid*. If you intend reconnections to resume a previous session by setting the *cleansession* parameter to false, the *clientIdentifier* used in each session must be the same. The default value of *clientid* is *hostname.connectionName*, which remains the same.

Installation, verification, configuration, and control of the WebSphere MQ Telemetry daemon for devices

Installation, configuration, and control of the daemon is file-based.

Install the daemon by copying the Software Development Kit to the device where you are going to run the daemon.

As an example, run the MQTT client utility and connect to the WebSphere MQ Telemetry daemon for devices as the publish/subscribe broker; see  Using the WebSphere MQ Telemetry daemon for devices as the publish/subscribe broker.

Configure the daemon by creating a configuration file; see  WebSphere MQ Telemetry daemon for devices configuration file (*WebSphere MQ V7.1 Reference*).

Control a running daemon by creating commands in the file, *amqtd.d.upd*. Every 5 seconds the daemon reads the file, runs the commands, and deletes the file; see  WebSphere MQ Telemetry daemon for devices command file (*WebSphere MQ V7.1 Reference*)

WebSphere MQ Telemetry daemon for devices listener ports

Connect MQTT V3 clients to the WebSphere MQ Telemetry daemon for devices using listener ports. You can qualify a listener port with a mount point and a maximum number of connections.

A listener port must correspond to the port number specified on the MQTT client connect(serverURI) method of a client connecting to this port. It defaults on both the client and the daemon to 1883.

You can change the default port for the daemon by setting the global definition port in the daemon configuration file. You can set specific ports by adding a listener definition to the daemon configuration file.

For each listener port, other than the default port, you can specify a mount point to isolate clients. Clients connected to a port with a mount point are isolated from other clients; see “WebSphere MQ Telemetry daemon for devices mount points.”

You can limit the number of clients that can connect to any port. Set the global definition max_connections to limit connections to the default port, or qualify each listener port with max_connections.

Example

An example of a configuration file that changes the default port from 1883 to 1880, and limits connections to port 1880 to 10000. Connections to port 1884 are limited to 1000. Clients attached to port 1884 are isolated from clients attached to other ports.

```
port 1880
max_connections 10000
listener 1884
mount_point 1884/
max_connections 1000
```

WebSphere MQ Telemetry daemon for devices mount points

You can associate a mount point with a listener port used by MQTT clients to connect to a WebSphere MQ Telemetry daemon for devices. A mount point isolates the publications and subscriptions exchanged by MQTT clients using one listener port from MQTT clients connected to a different listener port.

Clients attached to a listener port with a mount point can never directly exchange topics with clients attached to any other listener ports. Clients attached to a listener port without a mount point can publish or subscribe to topics of any client. Clients are not aware of whether they are attached through a mount point or not; it makes no difference to the topics strings created by clients.

A mount point is a string of text that is prefixed to the topic string of publications and subscriptions. It is prefixed to all the topic strings created by clients attached to listener port with a mount point. The string of text is removed from all topic strings sent to clients attached to the listener port.

If a listener port has no mount point, the topic strings of publications and subscriptions created and received by clients attached to the port are not altered.

Create mount point strings with a trailing /. That way the mount point is the parent topic of the topic tree for the mount point.

Example

A configuration file contains the following listener ports:

```
listener 1883
mount_point 1883/
listener 1884 127.0.0.1
mount_point 1884/
listener 1885
```

A client, attached to port 1883, creates a subscription to MyTopic. The daemon registers the subscription as 1883/MyTopic. Another client attached to port 1883 publishes a message on the topic, MyTopic. The daemon changes the topic string to 1883/MyTopic and searches for matching subscriptions. The subscriber on port 1883 receives the publication with the original topic string MyTopic. The daemon has removed the mount point prefix from the topic string.

Another client, attached to port 1884, also publishes on the topic MyTopic. This time the daemon registers the topic as 1884/MyTopic. The subscriber on port 1883 does not receive the publication, because the different mount point results in a subscription with a different topic string.

A client, attached to port 1885, publishes on the topic, 1883/MyTopic. The daemon does not change the topic string. The subscriber on port 1883 receives the publication to MyTopic.

WebSphere MQ Telemetry daemon for devices quality of service, durable subscriptions and retained publications

Quality of service settings apply only to a running daemon. If a daemon stops, whether in a controlled manner, or because of a failure, the state of inflight messages is lost. The delivery of a message at least once, or at most once, cannot be guaranteed if the daemon stops. WebSphere MQ Telemetry daemon for devices supports limited persistence. Set the **retained_persistence** configuration parameter to save retained publications and subscriptions when the daemon is shut down.

Unlike WebSphere MQ, the WebSphere MQ Telemetry daemon for devices does not journal persistent data. Session state, message state, and retained publications are not saved transactionally. By default, the daemon discards all data when it stops. You can set an option to periodically checkpoint subscriptions and retained publications. Message status is always lost when the daemon stops. All non-retained publications are lost.

Set the daemon configuration option, Retained_persistence to true, to save retained publications periodically to a file. When the daemon restarts, the retained publications that were last autosaved are reinstated. By default, retained messages created by clients are not reinstated when the daemon restarts.

Set the daemon configuration option, Retained_persistence to true, to save subscriptions created in a persistent session periodically to a file. If Retained_persistence is set to true, subscriptions that clients create in a session with CleanSession set to false, a “persistent session”, are restored. The daemon restores the subscriptions when it restarts, which start receiving publications. The client receives the publications when it restarts with CleanSession to false. By default, client session state is not saved when a daemon stops, and so subscriptions are not restored, even if the client sets CleanSession to false.

Retained_persistence is an autosave mechanism. It might not save the most recent retained publications or subscriptions. You can change how often retained publications and subscriptions are saved. Set the interval between saves, or the number of changes between saves, using the configuration options autosave_on_changes and autosave_interval.

Example configuration for setting persistence

```
# Sample configuration
# Daemon listens on port 1882 with persistence in /tmp
# Autosave every minute
port 1882
```

```
persistence_location /tmp/  
retained_persistence true  
autosave_on_changes false  
autosave_interval 60
```

WebSphere MQ Telemetry daemon for devices security

The WebSphere MQ Telemetry daemon for devices can authenticate clients that connect to it, use credentials to connect to other brokers, and control access to topics. The security the daemon provides is limited by being built using the WebSphere MQ Telemetry C client, which does not provide SSL support. Consequently, connections to and from the daemon are not encrypted, and cannot be authenticated using certificates.

By default, no security is switched on.

Authentication of clients

MQTT clients can set a username and password using the methods `MqttConnectOptions.setUsername` and `MqttConnectOptions.setPassword`.

Authenticate a client that connects to the daemon by checking the username and password provided by a client against entries in the password file. To enable authentication, create a password file and set the `password_file` parameter in the daemon configuration file; see  `password_file`.

Set the `allow_anonymous` parameter in the daemon configuration file to allow clients connecting without usernames or passwords to connect to a daemon that is checking authentication; see

 `allow_anonymous`. If a client does provide a username or password it is always checked against the password file, if the `password_file` parameter is set.

Set the `clientid_prefixes` parameter in the daemon configuration file to limit connections to specific clients. The clients must have `clientIdentifiers` that start with one of the prefixes listed in the

`clientid_prefixes` parameter; see  `clientid_prefixes`.

Bridge connection security

Each WebSphere MQ Telemetry daemon for devices bridge connection is an MQTT V3 client. You can set the username and password for each bridge connection as a bridge connection parameter in the daemon configuration file; see  `username` and  `password`. A bridge can then authenticate itself to a broker.

Access control of topics

If clients are being authenticated, the daemon can also provide control access to topics for each user. The daemon grants access control based on matching the topic to which a client is either publishing or subscribing with an access topic string in the access control file; see  `acl_file`.

The access control list has two parts. The first part controls access for all clients, including anonymous clients. The second part has a section for any user in the password file. It lists specific access control for each user.

Example

The security parameters are shown in the following example.

Daemon configuration file

```
acl_file c:\WMQTDaemon\config\acl.txt
password_file c:\WMQTDaemon\config\passwords.txt
allow_anonymous true
connection Daemon1
username daemon1
password deamonpassword
```

Password file, passwords.txt

```
Fred:Fredpassword
Barney:Barneypassword
```

Access control file, acl.txt

```
topic home/public/#
topic read meters/#
user Fred
topic write meters/fred
topic home/fred/#
user Barney
topic write meters/barney
topic home/barney/#
```

Administering multicast

Use this information to learn about the WebSphere MQ Multicast administration tasks such as reducing the size of multicast messages and enabling data conversion.

Getting started with multicast

Use this information to get started with WebSphere MQ Multicast topics and communication information objects.

About this task

WebSphere MQ Multicast messaging uses the network to deliver messages by mapping topics to group addresses. The following tasks are a quick way to test if the required IP address and port are correctly configured for multicast messaging.

Creating a COMMINFO object for multicast

The communication information (COMMINFO) object contains the attributes associated with multicast transmission. For more information about the COMMINFO object parameters, see

 [DEFINE COMMINFO \(WebSphere MQ V7.1 Reference\)](#).

Use the following command-line example to define a COMMINFO object for multicast:

```
DEFINE COMMINFO(MC1) GRPADDR(group address) PORT(port number)
```

where *MC1* is the name of your COMMINFO object, *group address* is your group multicast IP address or DNS name, and the *port number* is the port to transmit on (The default value is 1414).

A new COMMINFO object called *MC1* is created; This name is the name that you must specify when defining a TOPIC object in the next example.

Creating a TOPIC object for multicast

A topic is the subject of the information that is published in a publish/subscribe message, and a topic is defined by creating a TOPIC object. TOPIC objects have two parameters which define whether they can be used with multicast or not. These parameters are: **COMMINFO** and **MCAST**.

- **COMMINFO** This parameter specifies the name of the multicast communication information object.

For more information about the COMMINFO object parameters, see  [DEFINE COMMINFO \(WebSphere MQ V7.1 Reference\)](#).

- **MCAST** This parameter specifies whether multicast is allowable at this position in the topic tree.

Use the following command-line example to define a TOPIC object for multicast:

```
DEFINE TOPIC(ALLSPORTS) TOPICSTR('Sports') COMMINFO(MC1) MCAST(ENABLED)
```

A new TOPIC object called *ALLSPORTS* is created. It has a topic string *Sports*, its related communication information object is called *MC1* (which is the name you specified when defining a COMMINFO object in the previous example), and multicast is enabled.

Testing the multicast publish/subscribe

After the TOPIC and COMMINFO objects have been created, they can be tested using the `amqspubc` sample and the `amqssubc` sample. For more information about these samples see  The Publish/Subscribe sample programs (*WebSphere MQ V7.1 Programming Guide*).

1. Open two command-line windows; The first command line is for the `amqspubc` publish sample, and the second command line is for the `amqssubc` subscribe sample.
2. Enter the following command at command line 1:

```
amqspubc Sports QM1
```

where *Sports* is the topic string of the TOPIC object defined in an earlier example, and *QM1* is the name of the queue manager.

3. Enter the following command at command line 2:

```
amqssubc Sports QM1
```

where *Sports* and *QM1* are the same as used in step 2.

4. Enter `Hello world` at command line 1. If the port and IP address that are specified in the COMMINFO object are configured correctly; the `amqssubc` sample, which is listening on the port for publications from the specified address, outputs `Hello world` at command line 2.

WebSphere MQ Multicast topic topology

Use this example to understand the WebSphere MQ Multicast topic topology.

WebSphere MQ Multicast support requires that each subtree has its own multicast group and data stream within the total hierarchy.

The *classful network* IP addressing scheme has designated address space for multicast address. The full multicast range of IP address is 224.0.0.0 to 239.255.255.255, but some of these addresses are reserved.

For a list of reserved address either contact your system administrator or see  <http://www.iana.org/assignments/multicast-addresses> for more information. It is recommended that you use the locally scoped multicast address in the range of 239.0.0.0 to 239.255.255.255.

In the following diagram, there are two possible multicast data streams:

```
DEF COMMINFO(MC1) GRPADDR(239.XXX.XXX.XXX)
```

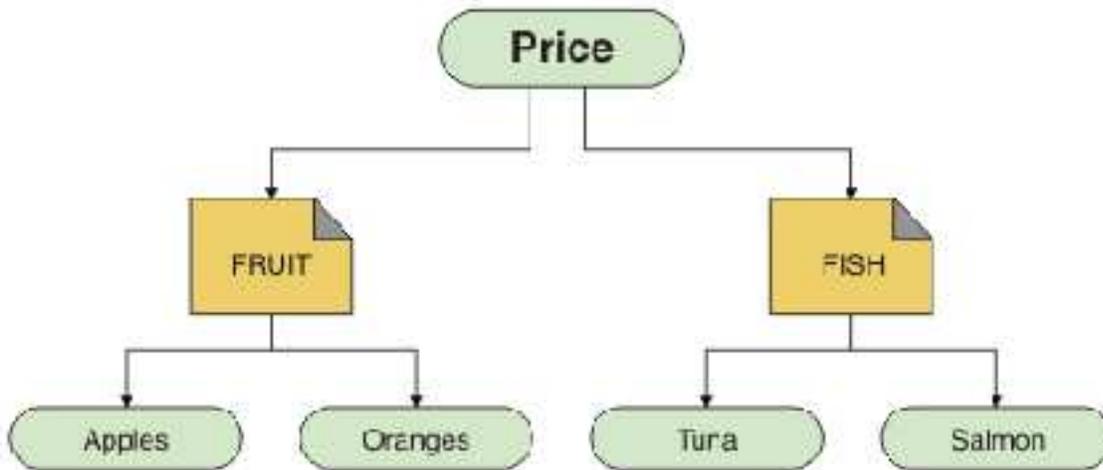
```
DEF COMMINFO(MC2) GRPADDR(239.YYY.YYY.YYY)
```

where *239.XXX.XXX.XXX* and *239.YYY.YYY.YYY* are valid multicast addresses.

These topic definitions are used to create a topic tree as shown in the following diagram:

```
DEFINE TOPIC(FRUIT) TOPICSTRING('Price/FRUIT') MCAST(ENABLED) COMMINFO(MC1)
```

```
DEFINE TOPIC(FISH) TOPICSTRING('Price/FISH') MCAST(ENABLED) COMMINFO(MC2)
```



Each multicast communication information (COMMINFO) object represents a different stream of data because their group addresses are different. In this example, the FRUIT topic is defined to use COMMINFO object MC1, the FISH topic is defined to use COMMINFO object MC2, and the Price node has no multicast definitions.

WebSphere MQ Multicast has a 255 character limit for topic strings. This limitation means that care must be taken with the names of nodes and leaf-nodes within the tree; if the names of nodes and leaf-nodes are too long, the topic string might exceed 255 characters and return the “2425 (0979) (RC2425): MQRC_TOPIC_STRING_ERROR” on page 1534 reason code. It is recommended to make topic strings as short as possible because longer topic strings might have a detrimental effect on performance.

Controlling the size of multicast messages

Use this information to learn about the WebSphere MQ message format, and reduce the size of WebSphere MQ messages.

WebSphere MQ messages have a number of attributes associated with them which are contained in the message descriptor. For small messages, these attributes might represent most of the data traffic and can have a significant detrimental effect on the transmission rate. WebSphere MQ Multicast enables the user to configure which, if any, of these attributes are transmitted along with the message.

The presence of message attributes, other than topic string, depends on whether the COMMINFO object states that they must be sent or not. If an attribute is not transmitted, the receiving application applies a default value. The default MQMD values are not necessarily the same as the MQMD_DEFAULT value, and are described later in this topic “Multicast message attributes” on page 161.

The COMMINFO object contains the MCPROP attribute which controls how many of the MQMD fields and user properties flow with the message. By setting the value of this attribute to an appropriate level, you can control the size of the WebSphere MQ Multicast messages:

MCPROP

The multicast properties control how many of the MQMD properties and user properties flow with the message.

ALL

All user properties and all the fields of the MQMD are transmitted.

REPLY

Only user properties, and MQMD fields that deal with replying to the messages, are transmitted. These properties are:

- MsgType
- MessageId
- CorrelId
- ReplyToQ
- ReplyToQmgr

USER

Only the user properties are transmitted.

NONE

No user properties or MQMD fields are transmitted.

COMPAT

This value causes the transmission of the message to be done in a compatible mode to RMM, which allows some inter-operation with the current XMS applications and WebSphere Message Broker RMM applications.

Multicast message attributes

Use this reference information to understand WebSphere MQ Multicast message attributes.

Message attributes can come from various places, such as the MQMD, the fields in the MQRFH2, and message properties.

The following table shows what happens when messages are sent subject to the value of MCPROP (described previously in this section), and the default value used when an attribute is not sent.

Attribute	Action when using multicast	Default if not transmitted
TopicString	Always Included	Not applicable
MQMQ StrucId	Not transmitted	Not applicable
MQMD Version	Not transmitted	Not applicable
Report	Included if not default	0
MsgType	Included if not default	MQMT_DATAGRAM
Expiry	Included if not default	0
Feedback	Included if not default	0
Encoding	Included if not default	MQENC_NORMAL(equiv)
CodedCharSetId	Included if not default	1208
Format	Included if not default	MQRFH2
Priority	Included if not default	4
Persistence	Included if not default	MQPER_NOT_PERSISTENT
MsgId	Included if not default	Null
CorrelId	Included if not default	Null
BackoutCount	Included if not default	0
ReplyToQ	Included if not default	Blank
ReplyToQMgr	Included if not default	Blank
UserIdentifier	Included if not default	Blank
AccountingToken	Included if not default	Null
PutAppIType	Included if not default	MQAT_JAVA
PutAppIName	Included if not default	Blank
PutDate	Included if not default	Blank

Attribute	Action when using multicast	Default if not transmitted
PutTime	Included if not default	Blank
ApplOriginData	Included if not default	Blank
GroupID	Excluded	Not applicable
MsgSeqNumber	Excluded	Not applicable
Offset	Excluded	Not applicable
MsgFlags	Excluded	Not applicable
OriginalLength	Excluded	Not applicable
UserProperties	Included	Not applicable

Related Links

 ALTER COMMINFO (*WebSphere MQ V7.1 Reference*)

 DEFINE COMMINFO (*WebSphere MQ V7.1 Reference*)

Enabling data conversion for Multicast messaging

Use this information to understand how data conversion works for WebSphere MQ Multicast messaging.

WebSphere MQ Multicast is a shared, connectionless protocol, and so it is not possible for each client to make specific requests for data conversion. Every client subscribed to the same multicast stream receives the same binary data; therefore, if WebSphere MQ data conversion is required, the conversion is performed locally at each client.

In a mixed platform installation, it might be that most of the clients require the data in a format that is not the native format of the transmitting application. In this situation the **CCSID** and **ENCODING** values of the multicast COMMINFO object can be used to define the encoding of the message transmission for efficiency.

WebSphere MQ Multicast supports data conversion of the message payload for the following built in formats:

- MQADMIN
- MQEVENT
- MQPCF
- MQRFH
- MQRFH2
- MQSTR

In addition to these formats, you can also define your own formats and use an  MQDXP – Data-conversion exit parameter (*WebSphere MQ V7.1 Reference*) data conversion exit.

For information about programming data conversions, see  Data conversion in the MQI for multicast messaging (*WebSphere MQ V7.1 Programming Guide*).

For more information about data conversion, see  Data conversion (*WebSphere MQ V7.1 Reference*).

For more information about data conversion exits and ClientExitPath, see  ClientExitPath stanza of the client configuration file (*WebSphere MQ V7.1 Installing Guide*).

Multicast application monitoring

Use this information to learn about administering and monitoring WebSphere MQ Multicast.

The status of the current publishers and subscribers for multicast traffic (for example, the number of messages sent and received, or the number of messages lost) is periodically transmitted to the server from the client. When status is received, the COMMEV attribute of the COMMINFO object specifies whether or not the queue manager puts an event message on the SYSTEM.ADMIN.PUBSUB.EVENT. The event message contains the status information received. This information is an invaluable diagnostic aid in finding the source of a problem.

Use the MQSC command **DISPLAY CONN** to display connection information about the applications connected to the queue manager. For more information on the **DISPLAY CONN** command, see  **DISPLAY CONN** (*WebSphere MQ V7.1 Reference*).

Use the MQSC command **DISPLAY TPSTATUS** to display the status of your publishers and subscribers. For more information on the **DISPLAY TPSTATUS** command, see  **DISPLAY TPSTATUS** (*WebSphere MQ V7.1 Reference*).

COMMEV and the multicast message reliability indicator

The *reliability indicator*, used in conjunction with the **COMMEV** attribute of the COMMINFO object, is a key element in the monitoring of WebSphere MQ Multicast publishers and subscribers. The reliability indicator (the **MSGREL** field that is returned on the Publish or Subscribe status commands) is a WebSphere MQ indicator that illustrates the percentage of transmissions that have no errors. Sometimes messages have to be retransmitted due to a transmission error, which is reflected in the value of **MSGREL**. Potential causes of transmission errors include slow subscribers, busy networks, and network outages. **COMMEV** controls whether event messages are generated for multicast handles that are created using the COMMINFO object and is set to one of three possible values:

DISABLED

Event messages are not written.

ENABLED

Event messages are always written, with a frequency defined in the COMMINFO **MONINT** parameter.

EXCEPTION

Event messages are written if the message reliability is below the reliability threshold. A message reliability level of 90% or less indicates that there might be a problem with the network configuration, or that one or more of the Publish/Subscribe applications is running too slowly:

- A value of **MSGREL(100,100)** indicates that there have been no issues in either the short term, or the long-term time frame.
- A value of **MSGREL(80,60)** indicates that 20% of the messages are currently having issues, but that it is also an improvement on the long-term value of 60.

Clients might continue transmitting and receiving multicast traffic even when the unicast connection to the queue manager is broken, therefore the data might be out of date.

Multicast message reliability

Use this information to learn how to set the WebSphere MQ Multicast subscription and message history.

A key element of overcoming transmission failure with multicast is WebSphere MQ's buffering of transmitted data (a history of messages to be kept at the transmitting end of the link). This process means that no buffering of messages is required in the putting application process because WebSphere MQ provides the reliability. The size of this history is configured via the communication information (COMMINFO) object, as described in the following information. A bigger transmission buffer means that there is more transmission history to be retransmitted if needed, but due to the nature of multicast, 100% assured delivery cannot be supported.

The WebSphere MQ Multicast message history is controlled in the communication information (COMMINFO) object by the **MSGHIST** attribute:

MSGHIST

This value is the amount of message history in kilobytes that is kept by the system to handle retransmissions in the case of NACKs (negative acknowledgments).

A value of 0 gives the least level of reliability. The default value is 100 KB.

The WebSphere MQ Multicast new subscription history is controlled in the communication information (COMMINFO) object by the **NSUBHIST** attribute:

NSUBHIST

The new subscriber history controls whether a subscriber joining a publication stream receives as much data as is currently available, or receives only publications made from the time of the subscription.

NONE

A value of NONE causes the transmitter to transmit only publication made from the time of the subscription. NONE is the default value.

ALL

A value of ALL causes the transmitter to retransmit as much history of the topic as is known. In some circumstances, this situation can give a similar behavior to retained publications.

Note: Using the value of ALL might have a detrimental effect on performance if there is a large topic history because all the topic history is retransmitted.

Related Links



DEFINE COMMINFO (*WebSphere MQ V7.1 Reference*)



ALTER COMMINFO (*WebSphere MQ V7.1 Reference*)

Advanced multicast tasks

Use this information to learn about advanced WebSphere MQ Multicast administration tasks such as configuring .ini files and interoperability with WebSphere MQ LLM.

For considerations for security in a Multicast installation, see “Multicast security” on page 460.

Bridging between multicast and non-multicast publish/subscribe domains

Use this information to understand what happens when a non-multicast publisher publishes to a WebSphere MQ Multicast enabled topic.

If a non-multicast publisher publishes to a topic that is defined as **MCAST** enabled and **BRIDGE** enabled, the queue manager transmits the message out over multicast directly to any subscribers that might be listening. A multicast publisher cannot publish to topics that are not multicast enabled.

Existing topics can be multicast enabled by setting the **MCAST** and **COMMINFO** parameters of a topic object.

See  Initial multicast concepts (*WebSphere MQ V7.1 Product Overview Guide*) for more information about these parameters.

The **COMMINFO** object **BRIDGE** attribute controls publications from applications that are not using multicast. If **BRIDGE** is set to **ENABLED** and the **MCAST** parameter of the topic is also set to **ENABLED**, publications from applications that are not using multicast are bridged to applications that do. For more information on the **BRIDGE** parameter, see  **DEFINE COMMINFO** (*WebSphere MQ V7.1 Reference*).

Configuring the .ini files for Multicast

Use this information to understand the WebSphere MQ Multicast fields in the .ini files.

Additional WebSphere MQ Multicast configuration can be made in an ini file. The specific ini file that you must use is dependent on the type of applications:

- Client: Configure the `MQ_DATA_PATH/mqclient.ini` file.
- Queue manager: Configure the `MQ_DATA_PATH/qmgrs/QMNAME/qm.ini` file.

where `MQ_DATA_PATH` is the location of the WebSphere MQ data directory (`/var/mqm/mqclient.ini`), and `QMNAME` is the name of the queue manager to which the .ini file applies.

The .ini file contains fields used to fine-tune the behavior of WebSphere MQ Multicast:

```
Multicast:
  Protocol           = IP | UDP
  IPVersion          = IPV4 | IPV6 | ANY | BOTH
  LimitTransRate     = DISABLED | STATIC | DYNAMIC
  TransRateLimit     = 100000
  SocketTTL          = 1
  Batch              = NO
  Loop               = 1
  Interface          = <IPaddress>
  FeedbackMode       = ACK | NACK | WAIT1
  HeartbeatTimeout   = 20000
  HeartbeatInterval = 2000
```

Protocol

- UDP** In this mode, packets are sent using the UDP protocol. Network elements cannot provide assistance in the multicast distribution as they do in IP mode however. The packet format remains compatible with PGM. This is the default value.
- IP** In this mode, the transmitter sends raw IP packets. Network elements with PGM support assist in the reliable multicast packet distribution. This mode is fully compatible with the PGM standard.

IPVersion

- IPV4** Communicate using the IPv4 protocol only. This is the default value.
- IPV6** Communicate using the IPv6 protocol only.
- ANY** Communicate using IPv4, IPv6, or both, depending on which protocol is available.

BOTH Supports communication using both IPv4 and IPv6.

LimitTransRate

DISABLED

There is no transmission rate control. This is the default value.

STATIC

Implements static transmission rate control. The transmitter would not transmit at a rate exceeding the rate specified by the **TransRateLimit** parameter.

DYNAMIC

The transmitter adapts its transmission rate according to the feedback it gets from the receivers. In this case the transmission rate limit cannot be more than the value specified by the **TransRateLimit** parameter. The transmitter tries to reach an optimal transmission rate.

TransRateLimit

The transmission rate limit in Kbps.

SocketTTL

The value of **SocketTTL** determines if the multicast traffic can pass through a router, or the number of routers it can pass through.

Batch Controls whether messages are batched or sent immediately. There are 2 possible values:

- *NO* The messages are not batched, they are sent immediately.
- *YES* The messages are batched.

Loop Set the value to 1 to enable multicast loop. Multicast loop defines whether the data sent is looped back to the host or not.

Interface

The IP address of the interface on which multicast traffic flows. For more information and troubleshooting, see: "Testing multicast applications on a non-multicast network" on page 1223 and "Setting the appropriate network for multicast traffic" on page 1223

FeedbackMode

NACK

Feedback by negative acknowledgments. This is the default value.

ACK Feedback by positive acknowledgments.

WAIT1

Feedback by positive acknowledgments where the transmitter waits for only 1 ACK from any of the receivers.

HeartbeatTimeout

The heartbeat timeout in milliseconds. A value of 0 indicates that the heartbeat timeout events are not raised by the receiver or receivers of the topic. The default value is 20000.

HeartbeatInterval

The heartbeat interval in milliseconds. A value of 0 indicates that no heartbeats are sent. The heartbeat interval must be considerably smaller than the **HeartbeatTimeout** value to avoid false heartbeat timeout events. The default value is 2000.

Multicast interoperability with WebSphere MQ Low Latency Messaging

Use this information to understand the interoperability between WebSphere MQ Multicast and WebSphere MQ Low Latency Messaging (LLM).

Basic payload transfer is possible for an application using LLM, with another application using multicast to exchange messages in both directions. Although multicast uses LLM technology, the LLM product itself is not embedded. Therefore it is possible to install both LLM and WebSphere MQ Multicast, and operate and service the two products separately.

LLM applications that communicate with multicast might need to send and receive message properties. The WebSphere MQ message properties and MQMD fields are transmitted as LLM message properties with specific LLM message property codes as shown in the following table:

WebSphere MQ property	WebSphere MQ LLM property type	LLM property kind	LLM property code
MQMD.Report	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1001
MQMD.MsgType	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1002
MQMD.Expiry	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1003
MQMD.Feedback	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1004
MQMD.Encoding	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1005
MQMD.CodedCharSetId	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1006
MQMD.Format	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1007
MQMD.Priority	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1008
MQMD.Persistence	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1009
MQMD.MsgId	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_ByteArray	-1010
MQMD.BackoutCount	RMM_MSG_PROP_INT32	LLM_PROP_KIND_Int32	-1012
MQMD.ReplyToQ	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1013
MQMD.ReplyToQMger	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1014
MQMD.PutDate	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1020
MQMD.PutTime	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1021
MQMD.ApplOriginData	RMM_MSG_PROP_BYTES	LLM_PROP_KIND_String	-1022
MQPubOptions	RMM_MSG_PROP_INT32	LLM_PROP_KIND_int32	-1053

For more information about LLM, see the LLM product documentation: [WebSphere MQ Low Latency Messaging](#).

Administering HP Integrity NonStop Server

Use this information to learn about administration tasks for the IBM WebSphere MQ client for HP Integrity NonStop Server.

Two administration tasks are available to you:

1. Manually starting the TMF/Gateway from Pathway.
2. Stopping the TMF/Gateway from Pathway.

Manually starting the TMF/Gateway from Pathway

You can allow Pathway to automatically start the TMF/Gateway on the first enlistment request, or you can manually start the TMF/Gateway from Pathway.

Procedure

To manually start the TMF/Gateway from Pathway, enter the following PATHCOM command:

```
START SERVER <server_class_name>
```

If a client application makes an enlistment request before the TMF/Gateway completes recovery of in-doubt transactions, the request is held for up to 1 second. If recovery does not complete within that time, the enlistment is rejected. The client then receives an MQRC_UOW_ENLISTMENT_ERROR error from use of a transactional MQI.

Stopping the TMF/Gateway from Pathway

This task describes how to stop the TMF/Gateway from Pathway, and how to restart the TMF/Gateway after you stop it.

Procedure

1. To prevent any new enlistment requests being made to the TMF/Gateway, enter the following command:

```
FREEZE SERVER <server_class_name>
```
2. To trigger the TMF/Gateway to complete any in-flight operations and to end, enter the following command:

```
STOP SERVER <server_class_name>
```
3. To allow the TMF/Gateway to restart either automatically on first enlistment or manually, following steps 1 and 2, enter the following command:

```
THAW SERVER <server_class_name>
```

Applications are prevented from making new enlistment requests and it is not possible to issue the **START** command until you issue the **THAW** command.

Administering IBM i

Introduces the methods available to you to administer IBM WebSphere MQ on IBM i.

Administration tasks include creating, starting, altering, viewing, stopping, and deleting clusters, processes, and IBM WebSphere MQ objects (queue managers, queues, namelists, process definitions, channels, client connection channels, listeners, services, and authentication information objects).

See the following links for details of how to administer IBM WebSphere MQ for IBM i:

- “Managing WebSphere MQ for IBM i using CL commands” on page 169
- “Alternative ways of administering WebSphere MQ for IBM i” on page 183
- “Work management” on page 188

Related concepts:

“Availability, backup, recovery, and restart” on page 195

Setting up security on IBM i

Related reference:

Changing configuration information on IBM i (*WebSphere MQ V7.1 Installing Guide*)



Understanding WebSphere MQ for IBM i queue manager library names (*WebSphere MQ V7.1 Product Overview Guide*)

“Quiescing WebSphere MQ for IBM i” on page 234



The dead letter queue handler on IBM i (*WebSphere MQ V7.1 Installing Guide*)

“Determining problems with IBM WebSphere MQ for IBM i applications” on page 1167



Installable services and components on IBM i (*WebSphere MQ V7.1 Programming Guide*)



System and default objects on IBM i (*WebSphere MQ V7.1 Reference*)

Managing WebSphere MQ for IBM i using CL commands

Use this information to understand the WebSphere MQ IBM i commands.

Most groups of WebSphere MQ commands, including those associated with queue managers, queues, topics, channels, namelists, process definitions, and authentication information objects can be accessed using the relevant **WRK*** command.

The principal command in the set is **WRKMQM**. This command allows you, for example, to display a list of all the queue managers on the system, together with status information. Alternatively, you can process all queue-manager specific commands using various options against each entry.

From the **WRKMQM** command you can select specific areas of each queue manager, for example, working with channels, topics or queues, and from there select individual objects.

Recording WebSphere MQ application definitions

When you create or customize WebSphere MQ applications, it is useful to keep a record of all WebSphere MQ definitions created. This record can be used for:

- Recovery purposes
- Maintenance
- Rolling out WebSphere MQ applications

You can record WebSphere MQ application definitions in 1 of 2 ways:

1. Creating CL programs to generate your WebSphere MQ definitions for the server.
2. Creating MQSC text files as SRC members to generate your WebSphere MQ definitions using the cross-platform WebSphere MQ command language.

SupportPac MS03 "WebSphere MQ - Save Queue Manager object definitions using PCFs" saves all the objects, such as queues or channels, defined in a either local or remote queue manager to a file. The SupportPac can be used to generate the MQSC script file. You can download this SupportPac from the

 [WebSphere MQ SupportPac Web page](#).

For further details about defining queue objects, see “Script (MQSC) Commands” on page 77 and “Using Programmable Command Formats” on page 7.

Before you start using the WebSphere MQ for IBM i using CL commands

Use this information to start the WebSphere MQ subsystem and create a local queue manager.

Before you begin

Ensure that the WebSphere MQ subsystem is running (using the command STRSBS QMQM/QMQM), and that the job queue associated with that subsystem is not held. By default, the WebSphere MQ subsystem and job queue are both named QMQM in library QMQM.

About this task

Using the IBM i command line to start a queue manager

Procedure

1. Create a local queue manager by issuing the **CRTMQM** command from an IBM i command line. When you create a queue manager, you have the option of making that queue manager the default queue manager. The default queue manager (of which there can only be one) is the queue manager to which a CL command applies, if the queue manager name parameter (MQMNAME) is omitted.
2. Start a local queue manager by issuing the **STRMQM** command from an IBM i command line. If the queue manager startup takes more than a few seconds WebSphere MQ will show status messages intermittently detailing the start up progress. For more information on these messages see "Reason codes" on page 1379.

What to do next

You can stop a queue manager by issuing the **ENDMQM** command from the IBM i command line, and control a queue manager by issuing other WebSphere MQ commands from an IBM i command line.

Remote queue managers cannot be started remotely but must be created and started in their systems by local operators. An exception to this is where remote operating facilities (outside WebSphere MQ for IBM i) exist to enable such operations.

The local queue administrator cannot stop a remote queue manager.

Note: As part of quiescing a WebSphere MQ system, you have to quiesce the active queue managers. This is described in "Quiescing WebSphere MQ for IBM i" on page 234.

Creating WebSphere MQ for IBM i objects

Use this information to understand the methods for creating WebSphere MQ objects for IBM i.

Before you begin

The following tasks suggest various ways in which you can use WebSphere MQ for IBM i from the command line.

About this task

There are two online methods to create WebSphere MQ objects, which are:

Procedure

1. Using a Create command, for example: The **Create MQM Queue** command: **CRTMQMQ**
2. Using a Work with MQM object command, followed by F6, for example: The **Work with MQM Queues** command: **WRKMQMQ**

What to do next

For a list of all commands see  WebSphere MQ for IBM i CL commands (*WebSphere MQ V7.1 Reference*).

Note: All MQM commands can be submitted from the Message Queue Manager Commands menu. To display this menu, type GO CMDMQM on the command line and press the Enter key.

The system displays the prompt panel automatically when you select a command from this menu. To display the prompt panel for a command that you have typed directly on the command line, press F4 before pressing the Enter key.

Creating a local queue using the CRTMQMQ command:

Procedure

1. Type CHGMQM on the command line and press the F4 key.
2. On the Create MQM Queue panel, type the name of the queue that you want to create in the Queue name field. To specify a mixed case name, you enclose the name in apostrophes.
3. Type *LCL in the Queue type field.
4. Specify a queue manager name, unless you are using the default queue manager, and press the Enter key. You can overtype any of the values with a new value. Scroll forward to see further fields. The options used for clusters are at the end of the list of options.
5. When you have changed any values, press the Enter key to create the queue.

Creating a local queue using the WRKMQMQ command:

Procedure

1. Type WRKMQMQ on the command line.
2. Enter the name of a queue manager.
3. If you want to display the prompt panel, press F4. The prompt panel is useful to reduce the number of queues displayed, by specifying a generic queue name or queue type.
4. Press Enter and the Work with MQM Queues panel is displayed. You can overtype any of the values with a new value. Scroll forward to see further fields. The options used for clusters are at the end of the list of options.
5. Press F6 to create a new queue; this takes you to the CRTMQMQ panel. See “Creating a local queue using the CRTMQMQ command” for instructions on how to create the queue. When you have created the queue, the Work with MQM Queues panel is displayed again. The new queue is added to the list when you press F5=Refresh.

Altering queue manager attributes:

About this task

To alter the attributes of the queue manager specified on the **CHGMQM** command, specifying the attributes and values that you want to change. For example, use the following options to alter the attributes of `jupiter.queue.manager`:

Procedure

Type **CHGMQM** on the command line and press the F4 key.

Results

The command changes the dead-letter queue used, and enables inhibit events.

Working with local queues

This section contains examples of some of the commands that you can use to manage local queues. All the commands shown are also available using options from the WRKMQMQ command panel.

Defining a local queue

For an application, the local queue manager is the queue manager to which the application is connected. Queues that are managed by the local queue manager are said to be local to that queue manager.

Use the command **CRTMQMQ QTYPE *LCL** to create a definition of a local queue and also to create the data structure that is called a queue. You can also modify the queue characteristics from those of the default local queue.

In this example, the queue we define, `orange.local.queue`, is specified to have these characteristics:

- It is enabled for gets, disabled for puts, and operates on a first-in-first-out (FIFO) basis.
- It is an *ordinary* queue, that is, it is not an initiation queue or a transmission queue, and it does not generate trigger messages.
- The maximum queue depth is 1000 messages; the maximum message length is 2000 bytes.

The following command does this on the default queue manager:

```
CRTMQMQ QNAME('orange.local.queue') QTYPE(*LCL)
      TEXT('Queue for messages from other systems')
      PUTENBL(*NO)
      GETENBL(*YES)
      TRGENBL(*NO)
      MSGDLYSEQ(*FIFO)
      MAXDEPTH(1000)
      MAXMSGLEN(2000)
      USAGE(*NORMAL)
```

Note:

1. USAGE *NORMAL indicates that this queue is not a transmission queue.
2. If you already have a local queue with the name `orange.local.queue` on the same queue manager, then this command fails. Use the REPLACE *YES attribute if you want to overwrite the existing definition of a queue, but see also “Changing local queue attributes” on page 173.

Defining a dead-letter queue

Each queue manager must have a local queue to be used as a dead-letter queue so that messages that cannot be delivered to their correct destination can be stored for later retrieval. You must explicitly tell the queue manager about the dead-letter queue. You can do this by specifying a dead-letter queue on the **CRTMQM** command, or you can use the **CHGMQM** command to specify one later. You must also define the dead-letter queue before it can be used.

A sample dead-letter queue called `SYSTEM.DEAD.LETTER.QUEUE` is supplied with the product. This queue is automatically created when you create the queue manager. You can modify this definition if required. There is no need to rename it, although you can if you like.

A dead-letter queue has no special requirements except that:

- It must be a local queue.
- Its MAXMSGL (maximum message length) attribute must enable the queue to accommodate the largest messages that the queue manager has to handle **plus** the size of the dead-letter header (MQDLH).

WebSphere MQ provides a dead-letter queue handler that allows you to specify how messages found on a dead-letter queue are to be processed or removed. For further information, see  The WebSphere MQ for IBM i dead-letter queue handler (*WebSphere MQ V7.1 Installing Guide*).

Displaying default object attributes

When you define a WebSphere MQ object, it takes any attributes that you do not specify from the default object. For example, when you define a local queue, the queue inherits any attributes that you omit in the definition from the default local queue, which is called `SYSTEM.DEFAULT.LOCAL.QUEUE`. To see exactly what these attributes are, use the following command:

```
DSPMQMQ QNAME(SYSTEM.DEFAULT.LOCAL.QUEUE) MQMNAME(MYQUEUEMANAGER)
```

Copying a local queue definition

You can copy a queue definition using the **CPYMQMQ** command. For example:

```
CPYMQMQ FROMQ('orange.local.queue') TOQ('magenta.queue') MQMNAME(MYQUEUEMANAGER)
```

This command creates a queue with the same attributes as our original queue `orange.local.queue`, rather than those of the system default local queue.

You can also use the **CPYMQMQ** command to copy a queue definition, but substituting one or more changes to the attributes of the original. For example:

```
CPYMQMQ FROMQ('orange.local.queue') TOQ('third.queue') MQMNAME(MYQUEUEMANAGER)  
MAXMSGLEN(1024)
```

This command copies the attributes of the queue `orange.local.queue` to the queue `third.queue`, but specifies that the maximum message length on the new queue is to be 1024 bytes, rather than 2000.

Note: When you use the **CPYMQMQ** command, you copy the queue attributes only, not the messages on the queue.

Changing local queue attributes

You can change queue attributes in two ways, using either the **CHGMQMQ** command or the **CPYMQMQ** command with the **REPLACE *YES** attribute. In “Defining a local queue” on page 172, you defined the queue `orange.local.queue`. If, for example, you need to increase the maximum message length on this queue to 10,000 bytes.

- Using the **CHGMQMQ** command:

```
CHGMQMQ QNAME('orange.local.queue') MQMNAME(MYQUEUEMANAGER) MAXMSGLEN(10000)
```

This command changes a single attribute, that of the maximum message length; all the other attributes remain the same.

- Using the **CRTMQMQ** command with the **REPLACE *YES** option, for example:

```
CRTMQMQ QNAME('orange.local.queue') QTYPE(*LCL) MQMNAME(MYQUEUEMANAGER)  
MAXMSGLEN(10000) REPLACE(*YES)
```

This command changes not only the maximum message length, but all the other attributes, which are given their default values. The queue is now put enabled whereas previously it was put inhibited. Put enabled is the default, as specified by the queue `SYSTEM.DEFAULT.LOCAL.QUEUE`, unless you have changed it.

If you *decrease* the maximum message length on an existing queue, existing messages are not affected. Any new messages, however, must meet the new criteria.

Clearing a local queue

To delete all the messages from a local queue called `magenta.queue`, use the following command:

```
CLRMQM QNAME('magenta.queue') MQMNAME(MYQUEUEMANAGER)
```

You cannot clear a queue if:

- There are uncommitted messages that have been put on the queue under syncpoint.
- An application currently has the queue open.

Deleting a local queue

Use the command **DLTMQM** to delete a local queue.

A queue cannot be deleted if it has uncommitted messages on it, or if it is in use.

Enabling large queues

WebSphere MQ supports queues larger than 2 GB. See your operating system documentation for information on how to enable IBM i to support large files.

The IBM i product documentation can be found here: [IBM i](#)

Some utilities might not be able to cope with files greater than 2 GB. Before enabling large file support, check your operating system documentation for information on restrictions on such support.

Working with alias queues

This section contains examples of some of the commands that you can use to manage alias queues. All the commands shown are also available using options from the **WRKMQM** command panel.

An alias queue (sometimes known as a queue alias) provides a method of redirecting MQI calls. An alias queue is not a real queue but a definition that resolves to a real queue. The alias queue definition contains a target queue name, which is specified by the **TGTQNAME** attribute.

When an application specifies an alias queue in an MQI call, the queue manager resolves the real queue name at run time.

For example, an application has been developed to put messages on a queue called `my.alias.queue`. It specifies the name of this queue when it makes an **MQOPEN** request and, indirectly, if it puts a message on this queue. The application is not aware that the queue is an alias queue. For each MQI call using this alias, the queue manager resolves the real queue name, which could be either a local queue or a remote queue defined at this queue manager.

By changing the value of the **TGTQNAME** attribute, you can redirect MQI calls to another queue, possibly on another queue manager. This is useful for maintenance, migration, and load-balancing.

Defining an alias queue

The following command creates an alias queue:

```
CRTMQMQ QNAME('my.alias.queue') QTYPE(*ALS) TGTQNAME('yellow.queue')  
MQMNAME(MYQUEUEMANAGER)
```

This command redirects MQI calls that specify `my.alias.queue` to the queue `yellow.queue`. The command does not create the target queue; the MQI calls fail if the queue `yellow.queue` does not exist at run time.

If you change the alias definition, you can redirect the MQI calls to another queue. For example:

```
CHGMQM QNAME('my.alias.queue') TGTQNAME('magenta.queue') MQMNAME(MYQUEUEMANAGER)
```

This command redirects MQI calls to another queue, `magenta.queue`.

You can also use alias queues to make a single queue (the target queue) appear to have different attributes for different applications. You do this by defining two aliases, one for each application. Suppose there are two applications:

- Application ALPHA can put messages on `yellow.queue`, but is not allowed to get messages from it.
- Application BETA can get messages from `yellow.queue`, but is not allowed to put messages on it.

You can do this using the following commands:

```
/* This alias is put enabled and get disabled for application ALPHA */
```

```
CRTMQMQ QNAME('alphas.alias.queue') QTYPE(*ALS) TGTQNAME('yellow.queue')  
PUTENBL(*YES) GETENBL(*NO) MQMNAME(MYQUEUEMANAGER)
```

```
/* This alias is put disabled and get enabled for application BETA */
```

```
CRTMQMQ QNAME('betas.alias.queue') QTYPE(*ALS) TGTQNAME('yellow.queue')  
PUTENBL(*NO) GETENBL(*YES) MQMNAME(MYQUEUEMANAGER)
```

ALPHA uses the queue name `alphas.alias.queue` in its MQI calls; BETA uses the queue name `betas.alias.queue`. They both access the same queue, but in different ways.

You can use the `REPLACE *YES` attribute when you define alias queues, in the same way that you use these attributes with local queues.

Using other commands with alias queues

You can use the appropriate commands to display or change alias queue attributes. For example:

```
* Display the alias queue's attributes */
```

```
DSPMQMQ QNAME('alphas.alias.queue') MQMNAME(MYQUEUEMANAGER)
```

```
/* ALTER the base queue name, to which the alias resolves. */  
/* FORCE = Force the change even if the queue is open. */
```

```
CHQMOMQ QNAME('alphas.alias.queue') TGTQNAME('orange.local.queue') FORCE(*YES)  
MQMNAME(MYQUEUEMANAGER)
```

Working with model queues

This section contains examples of some of the commands that you can use to manage model queues. All the commands shown are also available using options from the `WRKMOMQ` command panel.

A queue manager creates a dynamic queue if it receives an MQI call from an application specifying a queue name that has been defined as a model queue. The name of the new dynamic queue is generated by the queue manager when the queue is created. A model queue is a template that specifies the attributes of any dynamic queues created from it.

Model queues provide a convenient method for applications to create queues as they are required.

Defining a model queue

You define a model queue with a set of attributes in the same way that you define a local queue. Model queues and local queues have the same set of attributes, except that on model queues you can specify whether the dynamic queues created are temporary or permanent. (Permanent queues are maintained across queue manager restarts, temporary ones are not). For example:

```
CRTMQM QNAME('green.model.queue') QTYPE(*MDL) DFNTYPE(*PERMDYN)
```

This command creates a model queue definition. From the DFNTYPE attribute, the actual queues created from this template are permanent dynamic queues. The attributes not specified are automatically copied from the `SYSTEM.DEFAULT.MODEL.QUEUE` default queue.

You can use the REPLACE *YES attribute when you define model queues, in the same way that you use them with local queues.

Using other commands with model queues

You can use the appropriate commands to display or alter a model queue's attributes. For example:

```
/* Display the model queue's attributes */
```

```
DSPMQMQ MQMNAME(MYQUEUEMANAGER) QNAME('green.model.queue')
```

```
/* ALTER the model queue to enable puts on any */
```

```
/* dynamic queue created from this model. */
```

```
CHGMQM MQMNAME(MYQUEUEMANAGER) QNAME('blue.model.queue') PUTENBL(*YES)
```

Working with triggering

Use this information to learn about triggering and process definitions.

WebSphere MQ provides a facility for starting an application automatically when certain conditions on a queue are met. One example of the conditions is when the number of messages on a queue reaches a specified number. This facility is called *triggering* and is described in detail in  [Triggering channels \(WebSphere MQ V7.1 Installing Guide\)](#).

What is triggering?

The queue manager defines certain conditions as constituting trigger events. If triggering is enabled for a queue and a trigger event occurs, the queue manager sends a trigger message to a queue called an initiation queue. The presence of the trigger message on the initiation queue indicates that a trigger event has occurred.

Trigger messages generated by the queue manager are not persistent. This has the effect of reducing logging (thereby improving performance), and minimizing duplicates during restart, so improving restart time.

What is the trigger monitor?

The program which processes the initiation queue is called a trigger-monitor application, and its function is to read the trigger message and take appropriate action, based on the information contained in the trigger message. Normally this action would be to start some other application to process the queue which caused the trigger message to be generated. From the point of view of the queue manager, there is nothing special about the trigger-monitor application - it is another application that reads messages from a queue (the initiation queue).

Altering the job submission attributes of the trigger monitor

The trigger monitor supplied as command **STRMQMTRM** submits a job for each trigger message using the system default job description, QDFTJOB. This has limitations in that the submitted jobs are always called QDFTJOB and have the attributes of the default job description including the library list, *SYSVAL. WebSphere MQ provides a method for overriding these attributes. For example, it is possible to customize the submitted jobs to have more meaningful job names as follows:

1. In the job description specify the description you want, for example logging values.
2. Specify the Environment Data of the process definition used in the triggering process:
`CHGMQMPCRC PRCNAME(MY_PROCESS) MQMNAME(MHA3) ENVDATA ('JOB(MYLIB/TRIGJOB)')`

The Trigger Monitor performs a SBMJOB using the specified description.

It is possible to override other attributes of the SBMJOB by specifying the appropriate keyword and value in the Environment Data of the process definition. The only exception to this is the CMD keyword because this attribute is filled by the trigger monitor. An example of the command to specify the Environment Data of the process definition where both the job name and description are to be altered follows:

```
CHGMQMPCRC PRCNAME(MY_PROCESS) MQMNAME(MHA3) ENVDATA ('JOB(MYLIB/TRIGJOB)
JOB(TRIGGER)')
```

Defining an application queue for triggering

An application queue is a local queue that is used by applications for messaging, through the MQI. Triggering requires a number of queue attributes to be defined on the application queue. Triggering itself is enabled by the TRGENBL attribute.

In this example, a trigger event is to be generated when there are 100 messages of priority 5 or higher on the local queue `motor.insurance.queue`, as follows:

```
CRTMQMQ MQMNAME(MYQUEUEMANAGER) QNAME('motor.insurance.queue') QTYPE(*LCL)
PRCNAME('motor.insurance.quote.process') MAXMSGLEN(2000)
DFTMSGPST(*YES) INITQNAME('motor.ins.init.queue')
TRGENBL(*YES) TRGTYPE(*DEPTH) TRGDEPTH(100) TRGMSGPTY(5)
```

where the parameters are:

MQMNAME(MYQUEUEMANAGER)

The name of the queue manager.

QNAME('motor.insurance.queue')

The name of the application queue being defined.

PRCNAME('motor.insurance.quote.process')

The name of the application to be started by a trigger monitor program.

MAXMSGLEN(2000)

The maximum length of messages on the queue.

DFTMSGPST(*YES)

Messages on this queue are persistent by default.

INITQNAME('motor.ins.init.queue')

The name of the initiation queue on which the queue manager is to put the trigger message.

TRGENBL(*YES)

The trigger attribute value.

TRGTYPE(*DEPTH)

A trigger event is generated when the number of messages of the required priority (**TRGMSGPTY**) reaches the number specified in **TRGDEPTH**.

TRGDEPTH(100)

The number of messages required to generate a trigger event.

TRGMSGPTY(5)

The priority of messages that are to be counted by the queue manager in deciding whether to generate a trigger event. Only messages with priority 5 or higher are counted.

Defining an initiation queue

When a trigger event occurs, the queue manager puts a trigger message on the initiation queue specified in the application queue definition. Initiation queues have no special settings, but you can use the following definition of the local queue `motor.ins.init.queue` for guidance:

```
CRTMQMQ MQMNAME(MYQUEUEMANAGER) QNAME('motor.ins.init.queue') QTYPE(*LCL)
      GETENBL(*YES) SHARE(*NO) TRGTYPE(*NONE)
      MAXMSGL(2000)
      MAXDEPTH(1000)
```

Creating a process definition

Use the **CRTMQMPC** command to create a process definition. A process definition associates an application queue with the application that is to process messages from the queue. This is done through the **PRCNAME** attribute on the application queue `motor.insurance.quote`. The following command creates the required process, `motor.insurance.quote.process`, identified in this example:

```
CRTMQMPC MQMNAME(MYQUEUEMANAGER) PRCNAME('motor.insurance.quote.process')
      TEXT('Insurance request message processing')
      APPTYPE(*OS400) APPID(MQTEST/TESTPROG)
      USRDATA('open, cclose, 235')
```

where the parameters are:

MQMNAME(MYQUEUEMANAGER)

The name of the queue manager.

PRCNAME('motor.insurance.quote.process')

The name of the process definition.

TEXT('Insurance request message processing')

A description of the application program to which this definition relates. This text is displayed when you use the **DSPMQMPC** command. This can help you to identify what the process does. If you use spaces in the string, you must enclose the string in single quotation marks.

APPTYPE(*OS400)

The type of application to be started.

APPID(MQTEST/TESTPROG)

The name of the application executable file, specified as a fully qualified file name.

USRDATA('open, cclose, 235')

User-defined data, which can be used by the application.

Displaying your process definition

Use the **DSPMQMPC** command to examine the results of your definition. For example:

```
MQMNAME(MYQUEUEMANAGER) DSPMQMPC('motor.insurance.quote.process')
```

You can also use the **CHGMQMPC** command to alter an existing process definition, and the **DLTMQMPC** command to delete a process definition.

Displaying your process definition

Use the **DSPMQMPC** command to examine the results of your definition. For example:

```
MQMNAME(MYQUEUEMANAGER) DSPMQMPC('motor.insurance.quote.process')
```

You can also use the **CHGMQMPC** command to alter an existing process definition, and the **DLTMQMPC** command to delete a process definition.

Communicating between two systems

The following example illustrates how to set up two WebSphere MQ for IBM i systems, using CL commands, so that they can communicate with one another.

The systems are called SYSTEMA and SYSTEMB, and the communications protocol used is TCP/IP.

Carry out the following procedure:

1. Create a queue manager on SYSTEMA, calling it QMGRA1.

```
CRTMQM      MQMNAME(QMGRA1) TEXT('System A - Queue +
                Manager 1') UDLMSGQ(SYSTEM.DEAD.LETTER.QUEUE)
```

2. Start this queue manager.

```
STRMQM      MQMNAME(QMGRA1)
```

3. Define the WebSphere MQ objects on SYSTEMA that you need to send messages to a queue manager on SYSTEMB.

```
/* Transmission queue */
CRTMQMQ      QNAME(XMITQ.TO.QMGRB1) QTYPE(*LCL) +
                MQMNAME(QMGRA1) TEXT('Transmission Queue +
                to QMGRB1') MAXDEPTH(5000) USAGE(*TMQ)

/* Remote queue that points to a queue called TARGETB           */
/* TARGETB belongs to queue manager QMGRB1 on SYSTEMB          */
CRTMQMQ      QNAME(TARGETB.ON.QMGRB1) QTYPE(*RMT) +
                MQMNAME(QMGRA1) TEXT('Remote Q pointing +
                at Q TARGETB on QMGRB1 on Remote System +
                SYSTEMB') RMTQNAME(TARGETB) +
                RMTMQMNAME(QMGRB1) TMQNAME(XMITQ.TO.QMGRB1)

/* TCP/IP sender channel to send messages to the queue manager on SYSTEMB*/
CRTMQMCHL    CHLNAME(QMGRA1.TO.QMGRB1) CHLTYPE(*SDR) +
                MQMNAME(QMGRA1) TRPTYPE(*TCP) +
                TEXT('Sender Channel From QMGRA1 on +
                SYSTEMA to QMGRB1 on SYSTEMB') +
                CONNAME(SYSTEMB) TMQNAME(XMITQ.TO.QMGRB1)
```

4. Create a queue manager on SYSTEMB, calling it QMGRB1.

```
CRTMQM      MQMNAME(QMGRB1) TEXT('System B - Queue +
                Manager 1') UDLMSGQ(SYSTEM.DEAD.LETTER.QUEUE)
```

5. Start the queue manager on SYSTEMB.

```
STRMQM      MQMNAME(QMGRB1)
```

6. Define the WebSphere MQ objects that you need to receive messages from the queue manager on SYSTEMA.

```
/* Local queue to receive messages on */
CRTMQMQ      QNAME(TARGETB) QTYPE(*LCL) MQMNAME(QMGRB1) +
                TEXT('Sample Local Queue for QMGRB1')

/* Receiver channel of the same name as the sender channel on SYSTEMA */
CRTMQMCHL    CHLNAME(QMGRA1.TO.QMGRB1) CHLTYPE(*RCVR) +
                MQMNAME(QMGRB1) TRPTYPE(*TCP) +
                TEXT('Receiver Channel from QMGRA1 to +
                QMGRB1')
```

7. Finally, start a TCP/IP listener on SYSTEMB so that the channel can be started. This example uses the default port of 1414.

```
STRMQMLSR    MQMNAME(QMGRB1)
```

You are now ready to send test messages between SYSTEMA and SYSTEMB. Using one of the supplied samples, put a series of messages to your remote queue on SYSTEMA.

Start the channel on SYSTEMA, either by using the command **STRMQMCHL**, or by using the command **WRKMQMCHL** and entering a start request (Option 14) against the sender channel.

The channel should go to **RUNNING** status and the messages are sent to queue **TARGETB** on SYSTEMB.

Check your messages by issuing the command:

```
WRKMQMMSG QNAME(TARGETB) MQMNAME(QMGRB1).
```

Sample resource definitions

This sample contains the AMQSAMP4 sample IBM i CL program.

```
/* **** */
/*
/* Program name: AMQSAMP4
/*
/* Description: Sample CL program defining MQM queues
/* to use with the sample programs
/* Can be run, with changes as needed, after
/* starting the MQM
/*
/* <N_OCO_COPYRIGHT>
/* Licensed Materials - Property of IBM
/*
/* 63H9336
/* (c) Copyright IBM Corp. 1993, 2019 All Rights Reserved.
/*
/* US Government Users Restricted Rights - Use, duplication or
/* disclosure restricted by GSA ADP Schedule Contract with
/* IBM Corp.
/* <NOC_COPYRIGHT>
/*
/* **** */
/*
/* Function:
/*
/* AMQSAMP4 is a sample CL program to create or reset the
/* MQI resources to use with the sample programs.
/*
/* This program, or a similar one, can be run when the MQM
/* is started - it creates the objects if missing, or resets
/* their attributes to the prescribed values.
/*
/*
/*
/* Exceptions signaled: none
/* Exceptions monitored: none
/*
/* AMQSAMP4 takes a single parameter, the Queue Manager name
/*
/* **** */
/*
/* **** */
/* Queue Manager Name Parameter
/*
/* **** */
```

QSYS/DCL VAR(&QMGRNAME) TYPE(*CHAR)

```

/*****
/*      EXAMPLES OF DIFFERENT QUEUE TYPES      */
/*      */
/*      Create local, alias and remote queues  */
/*      */
/*      Uses system defaults for most attributes */
/*      */
/*****
/*      Create a local queue                    */
/*      CRTMQMQ      QNAME('SYSTEM.SAMPLE.LOCAL')      +
/*                   MQMNAME(&QMGRNAME)                +
/*                   QTYPE(*LCL)  REPLACE(*YES)         +
/*                   */
/*                   TEXT('Sample local queue') /* description */+
/*                   SHARE(*YES)           /* Shareable */+
/*                   DFTMSGPST(*YES) /* Persistent messages OK */+
/*
/*      Create an alias queue                    */
/*      CRTMQMQ      QNAME('SYSTEM.SAMPLE.ALIAS')      +
/*                   MQMNAME(&QMGRNAME)                +
/*                   QTYPE(*ALS)  REPLACE(*YES)         +
/*                   */
/*                   TEXT('Sample alias queue')         +
/*                   DFTMSGPST(*YES) /* Persistent messages OK */+
/*                   TGTQNAME('SYSTEM.SAMPLE.LOCAL')
/*
/*      Create a remote queue - in this case, an indirect reference */
/*      is made to the sample local queue on OTHER queue manager */
/*      CRTMQMQ      QNAME('SYSTEM.SAMPLE.REMOTE')      +
/*                   MQMNAME(&QMGRNAME)                +
/*                   QTYPE(*RMT)  REPLACE(*YES)         +
/*                   */
/*                   TEXT('Sample remote queue')/* description */+
/*                   DFTMSGPST(*YES) /* Persistent messages OK */+
/*                   RMTQNAME('SYSTEM.SAMPLE.LOCAL')      +
/*                   RMTMQMNAME(OTHER) /* Queue is on OTHER */+
/*
/*      Create a transmission queue for messages to queues at OTHER */
/*      By default, use remote node name */
/*      CRTMQMQ      QNAME('OTHER') /* transmission queue name */+
/*                   MQMNAME(&QMGRNAME)                +
/*                   QTYPE(*LCL)  REPLACE(*YES)         +
/*                   TEXT('Transmission queue to OTHER') +
/*                   USAGE(*TMQ) /* transmission queue */+
/*
/*****
/*      SPECIFIC QUEUES AND PROCESS USED BY SAMPLE PROGRAMS      */
/*      */
/*      Create local queues used by sample programs                */
/*      Create MQI process associated with sample initiation queue */
/*      */
/*****
/*      General reply queue                                        */
/*      CRTMQMQ      QNAME('SYSTEM.SAMPLE.REPLY')      +
/*                   MQMNAME(&QMGRNAME)                +
/*                   QTYPE(*LCL)  REPLACE(*YES)         +
/*                   */
/*                   TEXT('General reply queue')         +

```

```

                DFTMSGPST(*NO) /* Not Persistent          */
/* Queue used by AMQSINQ4                                */
   CRTMQMQ      QNAME('SYSTEM.SAMPLE.INQ')              +
                MQMNAME(&QMGRNAME)                    +
                QTYPE(*LCL) REPLACE(*YES)               +
                TEXT('Queue for AMQSINQ4')              +
                SHARE(*YES) /* Shareable                */+
                DFTMSGPST(*NO) /* Not Persistent        */+
                TRGENBL(*YES) /* Trigger control on     */+
                TRGTYPE(*FIRST)/* Trigger on first message*/+
                PRCNAME('SYSTEM.SAMPLE.INQPROCESS')     +
                INITQNAME('SYSTEM.SAMPLE.TRIGGER')
/* Queue used by AMQSSET4                                */
   CRTMQMQ      QNAME('SYSTEM.SAMPLE.SET')              +
                MQMNAME(&QMGRNAME)                    +
                QTYPE(*LCL) REPLACE(*YES)               +
                TEXT('Queue for AMQSSET4')              +
                SHARE(*YES) /* Shareable                */+
                DFTMSGPST(*NO)/* Not Persistent         */+
                TRGENBL(*YES) /* Trigger control on     */+
                TRGTYPE(*FIRST)/* Trigger on first message*/+
                PRCNAME('SYSTEM.SAMPLE.SETPROCESS')     +
                INITQNAME('SYSTEM.SAMPLE.TRIGGER')
/* Queue used by AMQSECH4                                */
   CRTMQMQ      QNAME('SYSTEM.SAMPLE.ECHO')             +
                MQMNAME(&QMGRNAME)                    +
                QTYPE(*LCL) REPLACE(*YES)               +
                TEXT('Queue for AMQSECH4')              +
                SHARE(*YES) /* Shareable                */+
                DFTMSGPST(*NO)/* Not Persistent         */+
                TRGENBL(*YES) /* Trigger control on     */+
                TRGTYPE(*FIRST)/* Trigger on first message*/+
                PRCNAME('SYSTEM.SAMPLE.ECHOPROCESS')    +
                INITQNAME('SYSTEM.SAMPLE.TRIGGER')
/* Initiation Queue used by AMQSTRG4, sample trigger process */
   CRTMQMQ      QNAME('SYSTEM.SAMPLE.TRIGGER') +
                MQMNAME(&QMGRNAME)                    +
                QTYPE(*LCL) REPLACE(*YES) +
                TEXT('Trigger queue for sample programs')
/* MQI Processes associated with triggered sample programs */
/*
/***** Note - there are versions of the triggered samples *****/
/***** in different languages - set APPID for these *****/
/***** process to the variation you want to trigger *****/
/*
   CRTMQMPRC   PRCNAME('SYSTEM.SAMPLE.INQPROCESS')     +
                MQMNAME(&QMGRNAME)                    +
                REPLACE(*YES)                          +

```

```

                TEXT('Trigger process for AMQSINQ4')          +
                ENVDATA('JOBPTY(3)') /* Submit parameter */ +
/** Select the triggered program here                **/      +
                APPID('QMQM/AMQSINQ4') /* C                  +
                /* APPID('QMQM/AMQ0INQ4') /* COBOL */        +
                /* APPID('QMQM/AMQ3INQ4') /* RPG - ILE */    +

                CRTMQMPCRC  PRCNAME('SYSTEM.SAMPLE.SETPROCESS') +
                            MQMNAME(&QMGRNAME)                +
                            REPLACE(*YES)                    +
                            +
                TEXT('Trigger process for AMQSSET4')          +
                ENVDATA('JOBPTY(3)') /* Submit parameter */ +
/** Select the triggered program here                **/      +
                APPID('QMQM/AMQSSET4') /* C                  +
                /* APPID('QMQM/AMQ0SET4') /* COBOL */        +
                /* APPID('QMQM/AMQ3SET4') /* RPG - ILE */    +

                CRTMQMPCRC  PRCNAME('SYSTEM.SAMPLE.ECHOPROCESS') +
                            MQMNAME(&QMGRNAME)                +
                            REPLACE(*YES)                    +
                            +
                TEXT('Trigger process for AMQSECH4')          +
                ENVDATA('JOBPTY(3)') /* Submit parameter */ +
/** Select the triggered program here                **/      +
                APPID('QMQM/AMQSECH4') /* C                  +
                /* APPID('QMQM/AMQ0ECH4') /* COBOL */        +
                /* APPID('QMQM/AMQ3ECH4') /* RPG - ILE */    +

                /*****/
                /*                                          */
                /* Normal return.                          */
                /*                                          */
                /*****/
                SNDPGMMSG  MSG('AMQSAMP4 Completed creating sample +
                                objects for ' *CAT &QMGRNAME)
                RETURN
                ENDPGM

                /*****/
                /*                                          */
                /* END OF AMQSAMP4                          */
                /*                                          */
                /*****/

```

Alternative ways of administering WebSphere MQ for IBM i

Use this information to learn about WebSphere MQ for IBM i, MQSC commands, PCF commands, and remote administration.

You can use WebSphere MQ instrumentation events to monitor the operation of queue managers. See “Instrumentation events” on page 830 for information about WebSphere MQ instrumentation events and how to use them.

You normally use IBM i CL commands to administer WebSphere MQ for IBM i. See “Managing WebSphere MQ for IBM i using CL commands” on page 169 for an overview of these commands.

Using CL commands is the preferred method of administering the system. However, you can use various other methods. This section gives an overview of those methods and includes the following topics:

Local and remote administration

You administer WebSphere MQ for IBM i objects locally or remotely.

Local administration means carrying out administration tasks on any queue managers that you have defined on your local system. In WebSphere MQ, you can consider this as local administration because no WebSphere MQ channels are involved, that is, the communication is managed by the operating system. To perform this type of task, you must either log onto the remote system and issue the commands from there, or create a process that can issue the commands for you.

WebSphere MQ supports administration from a single point through what is known as *remote administration*. Remote administration consists of sending programmable command format (PCF) control messages to the SYSTEM.ADMIN.COMMAND.QUEUE on the target queue manager.

There are a number of ways of generating PCF messages. These are:

1. Writing a program using PCF messages. See “Administration using PCF commands” on page 185.
2. Writing a program using the MQAI, which sends out PCF messages. See “Using the MQAI to simplify the use of PCFs” on page 18.
3. Using the WebSphere MQ Explorer, available with WebSphere MQ for Windows, which allows you to use a graphical user interface (GUI) and generates the correct PCF messages. See “Using the WebSphere MQ Explorer with WebSphere MQ for IBM i” on page 186.
4. Use **STRMQMQSC** to send commands indirectly to a remote queue manager. See “Administration using MQSC commands.”

For example, you can issue a remote command to change a queue definition on a remote queue manager.

Some commands cannot be issued in this way, in particular, creating or starting queue managers and starting command servers. To perform this type of task, you must either log onto the remote system and issue the commands from there or create a process that can issue the commands for you.

Administration using MQSC commands

Use this information to learn about MQSC commands, and how to use them to administer WebSphere MQ for IBM i.

WebSphere MQ script (MQSC) commands are written in human-readable form, that is, in EBCDIC text. You use MQSC commands to manage queue manager objects, including the queue manager itself, queues, process definitions, namelists, channels, client connection channels, listeners, services, topics, and authentication information objects.

You issue MQSC commands to a queue manager using the **STRMQMQSC** WebSphere MQ CL command. This method is a batch method only, taking its input from a source physical file in the server library system. The default name for this source physical file is QMQSC.

WebSphere MQ for IBM i does not supply a source file called QMQSC. To process MQSC commands you must create the QMQSC source file in a library of your choice, by issuing the following command:

```
CRTSRCPF FILE(MYLIB/QMQSC) RCDLEN(240) TEXT('WebSphere MQ - MQSC Source')
```

MQSC source is held in members within this source file. To work with the members enter the following command:

```
WRKMBRPDM MYLIB/QMQSC
```

You can now add new members and maintain existing ones

You can also enter MQSC commands interactively, by issuing RUNMQSC or:

1. Typing in the queue manager name and pressing the Enter key to access the **WRKMQM** results panel.

2. Selecting F23=More options on this panel.
3. Selecting option 26 against an active queue manager on the panel shown in Figure 42.

To end such an MQSC session, type **end**.

Figure 42 is an extract from an MQSC command file showing an MQSC command (DEFINE QLOCAL) with its attributes.

```

:
:
DEFINE QLOCAL(ORANGE.LOCAL.QUEUE) REPLACE +
DESCR(' ') +
PUT(ENABLED) +
DEFPRTY(0) +
DEFPSIST(NO) +
GET(ENABLED) +
MAXDEPTH(5000) +
MAXMSGL(1024) +
DEFSOPT(SHARED) +
NOHARDENBO +
USAGE(NORMAL) +
NOTRIGGER;
:
:

```

Figure 42. Extract from the MQSC command file, myprog.in

For portability among WebSphere MQ environments, limit the line length in MQSC command files to 72 characters. The plus sign indicates that the command is continued on the next line.

Object attributes specified in MQSC are shown in this section in uppercase (for example, RQMNAME), although they are not case-sensitive.

Note:

1. The format of an MQSC file does not depend on its location in the file system.
2. MQSC attribute names are limited to eight characters.
3. MQSC commands are available on other platforms, including z/OS.

For a description of each MQSC command and its syntax, see “Script (MQSC) Commands” on page 77.

Administration using PCF commands

The purpose of WebSphere MQ programmable command format (PCF) commands is to allow administration tasks to be programmed into an administration program. In this way you can create queues and process definitions, and change queue managers, from a program.

PCF commands cover the same range of functions provided by MQSC commands. However, unlike MQSC commands, PCF commands and their replies are not in a text format that you can read.

You can write a program to issue PCF commands to any queue manager in the network from a single node. In this way, you can both centralize and automate administration tasks.

Each PCF command is a data structure that is embedded in the application data part of a WebSphere MQ message. Each command is sent to the target queue manager using the MQI function MQPUT in the same way as any other message. The command server on the queue manager receiving the message interprets it as a command message and runs the command. To get the replies, the application issues an MQGET call and the reply data is returned in another data structure. The application can then process the reply and act accordingly.

Briefly, these are some of the things the application programmer must specify to create a PCF command message:

Message descriptor

This is a standard WebSphere MQ message descriptor, in which:

- Message type (*MsgType*) is MQMT_REQUEST.
- Message format (*Format*) is MQFMT_ADMIN.

Application data

Contains the PCF message including the PCF header, in which:

- The PCF message type (*Type*) specifies MQCFT_COMMAND.
- The command identifier specifies the command, for example, *Change Queue* (MQCMD_CHANGE_Q).

Escape PCFs are PCF commands that contain MQSC commands within the message text. You can use PCFs to send commands to a remote queue manager. See “Using the MQAI to simplify the use of PCFs” on page 18 for further information.

For a complete description of the PCF data structures and how to implement them, see  Structures for commands and responses (*WebSphere MQ V7.1 Reference*).

Using the WebSphere MQ Explorer with WebSphere MQ for IBM i

Use this information to administer WebSphere MQ for IBM i using the WebSphere MQ Explorer.

WebSphere MQ for Windows (x86 platform), and WebSphere MQ for Linux (x86 and x86-64 platforms) provide an administration interface called the WebSphere MQ Explorer to perform administration tasks as an alternative to using CL, control, or MQSC commands.

The WebSphere MQ Explorer allows you to perform local or remote administration of your network from a computer running Windows (x86 platform), or Linux (x86 and x86-64 platforms), by pointing the WebSphere MQ Explorer at the queue managers and clusters you are interested in. The platforms and levels of WebSphere MQ that can be administered using the WebSphere MQ Explorer are described in Remote queue managers.

With the WebSphere MQ Explorer, you can:

- Start and stop a queue manager (on your local machine only).
- Define, display, and alter the definitions of WebSphere MQ objects such as queues, topics, and channels.
- Browse the messages on a queue.
- Start and stop a channel.
- View status information about a channel.
- View queue managers in a cluster.
- Check to see which applications, users, or channels have a particular queue open.
- Create a new queue manager cluster using the Create New Cluster wizard.
- Add a queue manager to a cluster using the Add Queue Manager to Cluster wizard.
- Manage the authentication information object, used with Secure Sockets Layer (SSL) channel security.

Using the online guidance, you can:

- Define and control various resources including queue managers, queues, channels, process definitions, client connection channels, listeners, topics, services, namelists, and clusters.
- Start or stop a queue manager and its associated processes.
- View queue managers and their associated objects on your workstation or from other workstations.

- Check the status of queue managers, clusters, and channels.

Ensure that you have satisfied the following requirements before attempting to use the WebSphere MQ Explorer to manage WebSphere MQ on a server machine. Check that:

1. A command server is running for **any** queue manager being administered, started on the server by the CL command **STRMQCSVR**.
2. A suitable TCP/IP listener exists for every remote queue manager. This is the WebSphere MQ listener started by the **STRMQLSR** command.
3. The server connection channel, called `SYSTEM.ADMIN.SVRCONN`, exists on every remote queue manager. You **must** create this channel yourself. It is mandatory for every remote queue manager being administered. Without it, remote administration is not possible.
4. Verify that the `SYSTEM.MQEXPLORER.REPLY.MODEL` queue exists.

Managing the command server for remote administration

Use this information to learn about the remote administration of WebSphere MQ IBM i command server.

Each queue manager can have a command server associated with it. A command server processes any incoming commands from remote queue managers, or PCF commands from applications. It presents the commands to the queue manager for processing and returns a completion code or operator message depending on the origin of the command.

A command server is mandatory for all administration involving PCFs, the MQAI, and also for remote administration.

Note: For remote administration, you must ensure that the target queue manager is running. Otherwise, the messages containing commands cannot leave the queue manager from which they are issued. Instead, these messages are queued in the local transmission queue that serves the remote queue manager. Avoid this situation if at all possible.

There are separate control commands for starting and stopping the command server. You can perform the operations described in the following sections using the WebSphere MQ Explorer.

Starting and stopping the command server

To start the command server, use this CL command:

```
STRMQCSVR MQMNAME('saturn.queue.manager')
```

where `saturn.queue.manager` is the queue manager for which the command server is being started.

To stop the command server, use one of the following CL commands:

1.

```
ENDMQCSVR MQMNAME('saturn.queue.manager') OPTION(*CNTRLD)
```

to perform a controlled stop, where `saturn.queue.manager` is the queue manager for which the command server is being stopped. This is the default option, which means that the `OPTION(*CNTRLD)` can be omitted.

2.

```
ENDMQCSVR MQMNAME('saturn.queue.manager') OPTION(*IMMED)
```

to perform an immediate stop, where `saturn.queue.manager` is the queue manager for which the command server is being stopped.

Displaying the status of the command server

For remote administration, ensure that the command server on the target queue manager is running. If it is not running, remote commands cannot be processed. Any messages containing commands are queued in the target queue manager's command queue `SYSTEM.ADMIN.COMMAND.QUEUE`.

To display the status of the command server for a queue manager, called here `saturn.queue.manager`, the CL command is:

```
DSPMQMCSVR MQMNAME('saturn.queue.manager')
```

Issue this command on the target machine. If the command server is running, the panel shown in Figure 43 appears:

```
Display MQM Command Server (DSPMQMCSVR)

Queue manager name . . . . . > saturn.queue.manager

MQM Command Server Status. . . . > RUNNING

F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display
F24=More keys
```

Figure 43. Display MQM Command Server panel

Work management

This information describes the way in which WebSphere MQ handles work requests, and details the options available for prioritizing and controlling the jobs associated with WebSphere MQ.

Warning

Do **not** alter WebSphere MQ work management objects unless you fully understand the concepts of IBM i and WebSphere MQ work management. Additional information regarding subsystems and job descriptions can be found under *Work Management* in the IBM i product documentation, see [IBM i](#). Pay particular attention to the sections on "Job Starting and Routing" and "Batch Jobs".

WebSphere MQ for IBM i incorporates the IBM i UNIX environment and IBM i threads. Do **not** make any changes to the objects in the Integrated File System (IFS).

During normal operations, a WebSphere MQ queue manager starts a number of batch jobs to perform different tasks. By default these batch jobs run in the `QMQM` subsystem that is created when WebSphere MQ is installed.

Work management refers to the process of tailoring WebSphere MQ tasks to obtain the optimum performance from your system, or to make administration simpler.

For example, you can:

- Change the run-priority of jobs to make one queue manager more responsive than another.
- Redirect the output of a number of jobs to a particular output queue.
- Make all jobs of a certain type run in a specific subsystem.

- Isolate errors to a subsystem.

Work management is carried out by creating or changing the job descriptions associated with the WebSphere MQ jobs. You can configure work management for:

- An entire WebSphere MQ installation.
- Individual queue managers.
- Individual jobs for individual queue managers.

Description of WebSphere MQ tasks

This is a table of the WebSphere MQ jobs and a brief description of each.

When a queue manager is running, you see some or all of the following batch jobs running under the QMQM user profile in the WebSphere MQ subsystem. The jobs are described briefly in Table 4.

You can view all jobs connected to a queue manager using option 22 on the Work with Queue Manager (WRKMQM) panel. You can view listeners using the WRKMQMLSR command.

Table 4. WebSphere MQ tasks.

Job name	Function
AMQALMPX	The checkpoint processor that periodically takes journal checkpoints.
AMQZMUC0	Utility manager. This job executes critical queue manager utilities, for example the journal chain manager.
AMQZXMA0	The execution controller that is the first job started by the queue manager. It handles MQCONN requests, and starts agent processes to process WebSphere MQ API calls.
AMQZFUMA	Object authority manager (OAM).
AMQZLAA0	Queue manager agents that perform most of the work for applications that connect to the queue manager using MQCNO_STANDARD_BINDING.
AMQZLSA0	Queue manager agent.
AMQZMUF0	Utility Manager
AMQZMGR0	Process controller. This job is used to start up and manage listeners and services.
AMQZMUR0	Utility manager. This job executes critical queue manager utilities, for example the journal chain manager.
AMQFQPUB	Queued publish/subscribe daemon.
AMQFCXBA	Broker worker job.
RUNMQBRK	Broker control job.
AMQRMPPA	Channel process pooling job.
AMQCRSTA	TCP/IP-invoked channel responder.
AMQCRS6B	LU62 receiver channel and client connection (see note).
AMQRRMFA	Repository manager for clusters.
AMQCLMAA	Non-threaded TCP/IP listener.
AMQPCSEA	PCF command processor that handles PCF and remote administration requests.
RUNMQTRM	Trigger monitor.
RUNMQDLQ	Dead letter queue handler.
RUNMQCHI	The channel initiator.
RUNMQCHL	Sender channel job that is started for each sender channel.
RUNMQLSR	Threaded TCP/IP listener.
AMQRCMLA	Channel MQSC and PCF command processor.

Table 4. WebSphere MQ tasks. (continued)

Job name	Function
<p>Note: The LU62 receiver job runs in the communications subsystem and takes its runtime properties from the routing and communications entries that are used to start the job. See  Initiated end (Receiver) (<i>WebSphere MQ V7.1 Installing Guide</i>) for more information.</p>	

WebSphere MQ for IBM i work management objects

When WebSphere MQ is installed, various objects are supplied in the QMQM library to assist with work management. These objects are the ones necessary for WebSphere MQ jobs to run in their own subsystem.

Sample job descriptions are provided for two of the WebSphere MQ batch jobs. If no specific job description is provided for a WebSphere MQ job, it runs with the default job description QMQMJOB.

The work management objects that are supplied when you install WebSphere MQ are listed in Table 5 and the objects created for a queue manager are listed in Table 6

Note: The work management objects can be found in the QMQM library and the queue manager objects can be found in the queue manager library.

Table 5. Work management objects

Name	Type	Description
AMQALMPX	*JOB	The job description that is used by the checkpoint process
AMQZLAA0	*JOB	The job description that is used by the WebSphere MQ agent processes
AMQZLSA0	*JOB	The isolated bindings queue manager agent
AMQZXMA0	*JOB	The job description that is used by WebSphere MQ execution controllers
QMQM	*SBS	The subsystem in which all WebSphere MQ jobs run
QMQM	*JOBQ	The job queue attached to the supplied subsystem
QMQMJOB	*JOB	The default WebSphere MQ job description, used if there is not a specific job description for a job
QMQMMSG	*MSGQ	The default message queue for WebSphere MQ jobs.
QMQMRUN20	*CLS	A class description for high priority WebSphere MQ jobs
QMQMRUN35	*CLS	A class description for medium priority WebSphere MQ jobs
QMQMRUN50	*CLS	A class description for low priority WebSphere MQ jobs

Table 6. Work management objects created for a queue manager

Name	Type	Description
AMQA000000	*JRNRCV	Local journal receiver
AMQAJRN	*JRN	Local journal
AMQJRNINF	*USRSPC	User space that is updated with the latest journal receivers required for startup and media recovery of a queue manager. This user space can be queried by an application to determine which journal receivers require archiving and which can be safely deleted.
AMQAJRNMSG	*MSGQ	Local journal message queue
AMQCRC6B	*PGM	Program to start the LU6.2 connection
AMQRFOLD	*FILE	Migrated queue manager channel definition file

Table 6. Work management objects created for a queue manager (continued)

Name	Type	Description
QMQMMSG	*MSGQ	Queue manager message queue

How WebSphere MQ uses the work management objects

This information describes the way in which WebSphere MQ uses the work management objects, and provides configuration examples.

Attention: Do not alter the job queue entry settings in the QMQM subsystem to limit the number of jobs allowed in the subsystem by priority. If you attempt to do this, you can stop essential WebSphere MQ jobs from running after they are submitted and cause the queue manager startup to fail.

To understand how to configure work management, you must first understand how WebSphere MQ uses job descriptions.

The job description used to start the job controls many attributes of the job. For example:

- The job queue on which the job is queued and on which subsystem the job runs.
- The routing data used to start the job and class that the job uses for its runtime parameters.
- The output queue that the job uses for print files.

The process of starting a WebSphere MQ job can be considered in three steps:

1. WebSphere MQ selects a job description.

WebSphere MQ uses the following technique to determine which job description to use for a batch job:

- a. Look in the queue manager library for a job description with the same name as the job. See

 Understanding WebSphere MQ for IBM i queue manager library names (*WebSphere MQ V7.1 Product Overview Guide*) for further details about the queue manager library.

- b. Look in the queue manager library for the default job description QMQMJOB.
- c. Look in the QMQM library for a job description with the same name as the job.
- d. Use the default job description, QMQMJOB, in the QMQM library.

2. The job is submitted to the job queue.

Job descriptions supplied with WebSphere MQ have been set up, by default, to put jobs on to job queue QMQM in library QMQM. The QMQM job queue is attached to the supplied QMQM subsystem, so by default the jobs start running in the QMQM subsystem.

3. The job enters the subsystem and goes through the routing steps.

When the job enters the subsystem, the routing data specified on the job description is used to find routing entries for the job.

The routing data must match one of the routing entries defined in the QMQM subsystem, and this defines which of the supplied classes (QMQRUN20, QMQRUN35, or QMQRUN50) is used by the job.

Note: If WebSphere MQ jobs do not appear to be starting, make sure that the subsystem is running and the job queue is not held,

If you have modified the WebSphere MQ work management objects, make sure everything is associated correctly. For example, if you specify a job queue other than QMQM/QMQM on the job description, make sure that an ADDJOBQE is performed for the subsystem, that is, QMQM.

You can create a job description for each job documented in Table 4 on page 189 using the following worksheet as an example:

What is the queue manager library name? _____

Does job description AMQZXMA0 exist in the queue manager library?	Yes	No
Does job description QMQMJOB0 exist in the queue manager library?	Yes	No
Does job description AMQZXMA0 exist in the QMQM library?	Yes	No
Does job description QMQMJOB0 exist in the QMQM library?	Yes	No

If you answer No to all these questions, create a global job description QMQMJOB0 in the QMQM library.

The WebSphere MQ message queue

A WebSphere MQ message queue, QMQMMSG, is created in each queue manager library. Operating system messages are sent to this queue when queue manager jobs end and WebSphere MQ sends messages to the queue. For example, to report which journal receivers are needed at startup. Keep the number of messages in this message queue at a manageable size to make it easier to monitor.

Default system examples

These examples show how an unmodified WebSphere MQ installation works when some of the standard jobs are submitted at queue manager startup time.

First, the AMQZXMA0 execution controller job starts.

1. Issue the **STRMQM** command for queue manager TESTQM.
2. WebSphere MQ searches the queue manager library QMTESTQM, firstly for job description AMQZXMA0, and then job description QMQMJOB0.
Neither of these job descriptions exist, so WebSphere MQ looks for job description AMQZXMA0 in the product library QMQM. This job description exists, so it is used to submit the job.
3. The job description uses the WebSphere MQ default job queue, so the job is submitted to job queue QMQM/QMQM.
4. The routing data on the AMQZXMA0 job description is QMQMRUN20, so the system searches the subsystem routing entries for one that matches that data.
By default, the routing entry with sequence number 9900 has comparison data that matches QMQMRUN20, so the job is started with the class defined on that routing entry, which is also called QMQMRUN20.
5. The QMQM/QMQMRUN20 class has run priority set to 20, so the AMQZXMA0 job runs in subsystem QMQM with the same priority as most interactive jobs on the system.

Next, the AMQALMPX checkpoint process job starts.

1. WebSphere MQ searches the queue manager library QMTESTQM, firstly for job description AMQALPMX, and then job description QMQMJOB0.
Neither of these job descriptions exist, so WebSphere MQ looks for job descriptions AMQALMPX and QMQMJOB0 in the product library QMQM.
Job description AMQALMPX does not exist but QMQMJOB0 does, so QMQMJOB0 is used to submit the job.

Note: The QMQMJOB0 job description is always used for WebSphere MQ jobs that do not have their own job description.
2. The job description uses the WebSphere MQ default job queue, so the job is submitted to job queue QMQM/QMQM.
3. The routing data on the QMQMJOB0 job description is QMQMRUN35, so the system searches the subsystem routing entries for one that matches that data.
By default, the routing entry with sequence number 9910 has comparison data that matches QMQMRUN35, so the job is started with the class defined on that routing entry, which is also called QMQMRUN35.

4. The QMQM/QMQMRUN35 class has run priority set to 35, so the AMQALMPX job runs in subsystem QMQM with a lower priority than most interactive jobs on the system, but higher priority than most batch jobs.

Configuring work management examples

Use this information to learn how you can change and create WebSphere MQ job descriptions to change the runtime attributes of WebSphere MQ jobs.

The key to the flexibility of WebSphere MQ work management lies in the two-tier way that WebSphere MQ searches for job descriptions:

- If you create or change job descriptions in a queue manager library, those changes override the global job descriptions in QMQM, but the changes are *local* and affect that particular queue manager alone.
 - If you create or change global job descriptions in the QMQM library, those job descriptions affect all queue managers on the system, unless overridden locally for individual queue managers.
1. The following example increases the priority of channel control jobs for an individual queue manager.

To make the repository manager and channel initiator jobs, AMQRRMFA and RUNMQCHI, run as quickly as possible for queue manager TESTQM, carry out the following steps:

- a. Create local duplicates of the QMQM/QMQMJOBDD job description with the names of the WebSphere MQ processes that you want to control in the queue manager library. For example:

```
CRTDUPOBJ OBJ(QMQMJOBDD) FROMLIB(QMQM) OBJTYPE(*JOBDD) TOLIB(QMTESTQM)
NEWOBJ(RUNMQCHI)
CRTDUPOBJ OBJ(QMQMJOBDD) FROMLIB(QMQM) OBJTYPE(*JOBDD) TOLIB(QMTESTQM)
NEWOBJ(AMQRRMFA)
```

- b. Change the routing data parameter on the job description to ensure that the jobs use the QMQMRUN20 class.

```
CHGJOBDD JOBDD(QMTESTQM/RUNMQCHI) RTGDTA('QMQMRUN20')
CHGJOBDD JOBDD(QMTESTQM/AMQRRMFA) RTGDTA('QMQMRUN20')
```

The AMQRRMFA and RUNMQCHI jobs for queue manager TESTQM now:

- Use the new local job descriptions in the queue manager library
- Run with priority 20, because the QMQMRUN20 class is used when the jobs enter the subsystem.

2. The following example defines a new run priority class for the QMQM subsystem.

- a. Create a duplicate class in the QMQM library, to allow other queue managers to access the class, by issuing the following command:

```
CRTDUPOBJ OBJ(QMQMRUN20) FROMLIB(QMQM) OBJTYPE(*CLS) TOLIB(QMQM)
NEWOBJ(QMQMRUN10)
```

- b. Change the class to have the new run priority by issuing the following command:

```
CHGCLS CLS(QMQM/QMQMRUN10) RUNPTY(10)
```

- c. Add the new class definition to the subsystem by issuing the following command:

```
ADDRTGE SBSD(QMQM/QMQM) SEQNBR(8999) CMPVAL('QMQMRUN10') PGM(QSYS/QCMD)
CLS(QMQM/QMQMRUN10)
```

Note: You can specify any numeric value for the routing sequence number, but the values must be in sequential order. This sequence number tells the subsystem the order in which routing entries are to be searched for a routing data match.

- d. Change the local or global job description to use the new priority class by issuing the following command:

```
CHGJOBDD JOBDD(QMQMlibname/QMQMJOBDD) RTGDTA('QMQMRUN10')
```

Now all the queue manager jobs associated with the QMLibraryname use a run priority of 10.

3. The following example runs a queue manager in its own subsystem

To make all the jobs for queue manager TESTQM run in the QBATCH subsystem, carry out the following steps:

- a. Create a local duplicate of the QMQM/QMQMJOB job description in the queue manager library with the command


```
CRTDUPOBJ OBJ(QMQMJOB) FROMLIB(QMQM) OBJTYPE(*JOB) TOLIB(QMTESTQM)
```
- b. Change the job queue parameter on the job description to ensure that the jobs use the QBATCH job queue.


```
CHGJOB JOB(QMTESTQM/QMQMJOB) JOBQ(*LIBL/QBATCH)
```

Note: The job queue is associated with the subsystem description. If you find that the jobs are staying on the job queue, verify that the job queue definition is defined on the SBS. Use the DSPSBS command for the subsystem and take option 6, "Job queue entries".

All jobs for queue manager TESTQM now:

- Use the new local default job description in the queue manager library
- Are submitted to job queue QBATCH.

To ensure that jobs are routed and prioritized correctly:

- Either create routing entries for the WebSphere MQ jobs in subsystem QBATCH, or
- Rely on a catch-all routing entry that calls QCMD, irrespective of what routing data is used. This option works only if the maximum active jobs option for job queue QBATCH is set to *NOMAX. The system default is 1.

4. The following example creates another WebSphere MQ subsystem
 - a. Create a duplicate subsystem in the QMQM library by issuing the following command:


```
CRTDUPOBJ OBJ(QMQM) FROMLIB(QMQM) OBJTYPE(*SBS) TOLIB(QMQM) NEWOBJ(QMQM2)
```
 - b. Remove the QMQM job queue by issuing the following command:


```
RMVJOBQE SBS(QMQM/QMQM2) JOBQ(QMQM/QMQM2)
```
 - c. Create a new job queue for the subsystem by issuing the following command:


```
CRTJOBQ JOBQ(QMQM/QMQM2) TEXT('Job queue for MQSeries Queue Manager')
```
 - d. Add a job queue entry to the subsystem by issuing the following command:


```
ADDJOBQE SBS(QMQM/QMQM2) JOBQ(QMQM/QMQM2) MAXACT(*NOMAX)
```
 - e. Create a duplicate QMQMJOB in the queue manager library by issuing the following command:


```
CRTDUPOBJ OBJ(QMQMJOB) FROMLIB(QMQM) OBJTYPE(*JOB) TOLIB(QMlibraryname)
```
 - f. Change the job description to use the new job queue by issuing the following command:


```
CHGJOB JOB(QMlibraryname/QMQMJOB) JOBQ(QMQM/QMQM2)
```
 - g. Start the subsystem by issuing the following command:


```
STRSBS SBS(QMQM/QMQM2)
```

Note:

- a. You can specify the subsystem in any library. If for any reason the product is reinstalled, or the QMQM library is replaced, any changes you made are removed.
- b. All the queue manager jobs associated with the QMlibraryname now run under subsystem QMQM2.

5. The following example collects all output for a job type.

To collect all the checkpoint process, AMQALMPX, job logs for multiple queue managers onto a single output queue, carry out the following steps:

- a. Create an output queue, for example


```
CRTOUTQ OUTQ(MYLIB/CHCKPTLOGS)
```
- b. Create a global duplicate of the QMQM/QMQMJOB job description, using the name of the WebSphere MQ process that you want to control, for example


```
CRTDUPOBJ OBJ(QMQMJOB) FROMLIB(QMQM) OBJTYPE(*JOB) NEWOBJ(AMQALMPX)
```

- c. Change the output queue parameter on the job description to point to your new output queue, and change the job logging level so that all messages are written to the job log.

```
CHGJOB JOB(QMQM/AMQALMPX) OUTQ(MYLIB/CHKPTLOGS) LOG(4 00 *SECLVL)
```

All WebSphere MQ AMQALMPX jobs, for all queue managers, use the new global AMQALMPX job description, provided that there are no local overriding job descriptions in the local queue manager library.

All job log spool files for these jobs are now written to output queue CHKPTLOGS in library MYLIB.

Note:

- a. The preceding example works only if the QPJOBLOG, or any print file, has a value of *JOB for its output queue parameter. In the preceding example, the QSYS/QPDJOBLOG file needs OUTQ set to *JOB.
- b. To change a system print file, use the CHGPRTF command. For example:

```
CHGPRTF PRTF(QJOBLOG) OUTQ(*JOB)
```

The *JOB option indicates that your job descriptions must be used.

- c. You can send any spool files associated with the WebSphere MQ jobs to a particular output queue. However, verify that the print file being used has the appropriate value for the OUTQ parameter.

Availability, backup, recovery, and restart

Use this information to understand how WebSphere MQ for IBM i uses the IBM i journaling support to help its backup and restore strategy.

You must be familiar with standard IBM i backup and recovery methods, and with the use of journals and their associated journal receivers on IBM i, before reading this section. For information on these topics, see *IBM i Backup and Recovery*.

To understand the backup and recovery strategy, you first need to understand how WebSphere MQ for IBM i organizes its data in the IBM i file system and the integrated file system (IFS).

WebSphere MQ for IBM i holds its data in an individual library for each queue manager instance, and in stream files in the IFS file system.

The queue manager specific libraries contain journals, journal receivers, and objects required to control the work management of the queue manager. The IFS directories and files contain WebSphere MQ configuration files, the descriptions of WebSphere MQ objects, and the data they contain.

Every change to these objects, that is recoverable across a system failure, is recorded in a journal *before* it is applied to the appropriate object. This has the effect that such changes can be recovered by replaying the information recorded in the journal.

You can configure WebSphere MQ for IBM i to use multiple queue manager instances on different servers to provide increased queue manager availability and speed up recovery in the case of a server or queue manager failure.

WebSphere MQ for IBM i journals

Use this information to understand how WebSphere MQ for IBM i uses journals in its operation to control updates to local objects.

Each queue manager library contains a journal for that queue manager, and the journal has the name *QMGRLIB/AMQAJRN*, where *QMGRLIB* is the name of the queue manager library, and *A* is a letter, *A* in the case of a single instance queue manager, that is unique to the queue manager instance.

QMGRLIB takes the name *QM*, followed by the name of the queue manager in a unique form. For example, a queue manager named *TEST* has a queue manager library named *QMTEST*. The queue manager library can be specified when creating a queue manager using the **CRTMQM** command.

Journals have associated journal receivers that contain the information being journaled. The receivers are objects to which information can only be appended and will fill up eventually.

Journal receivers use up valuable disk space with out-of-date information. However, you can place the information in permanent storage to minimize this problem. One journal receiver is attached to the journal at any particular time. If the journal receiver reaches its predetermined threshold size, it is detached and replaced by a new journal receiver. You can specify the threshold of journal receivers when you create a queue manager using **CRTMQM** and the **THRESHOLD** parameter.

The journal receivers associated with the local WebSphere MQ for IBM i journal exist in each queue manager library, and adopt a naming convention as follows:

AMQArnnnnn

where

A is a letter A-Z. It is *A* for single instance queue managers. It varies for different instances of a multi-instance queue manager.

nnnnn is decimal 00000 to 99999 that is incremented by 1 for the next journal in the sequence.

r is decimal 0 to 9, that is incremented by 1 each time a receiver is restored.

The sequence of the journals is based on date. However, the naming of the next journal is based on the following rules:

1. *AMQArnnnnn* goes to *AMQAr(nnnnn+1)*, and *nnnnn* wraps when it reaches 99999. For example, *AMQA099999* goes to *AMQA000000*, and *AMQA999999* goes to *AMQA900000*.
2. If a journal with a name generated by rule 1 already exists, the message *CPI70E3* is sent to the *QSYSOPR* message queue and automatic receiver switching stops.
The currently-attached receiver continues to be used until you investigate the problem and manually attach a new receiver.
3. If no new name is available in the sequence (that is, all possible journal names are on the system) you need to do both of the following:
 - a. Delete journals no longer needed (see "Journal management" on page 201).
 - b. Record the journal changes into the latest journal receiver using (**RCDMQMIMG**) and then repeat the previous step. This allows the old journal receiver names to be reused.

The *AMQAJRN* journal uses the *MNGRCV(*SYSTEM)* option to enable the operating system to automatically change journal receivers when the threshold is reached. For more information on how the system manages receivers, see *IBM i Backup and Recovery*.

The journal receiver's default threshold value is 100,000 KB. You can set this to a larger value when you create the queue manager. The initial value of the *LogReceiverSize* attribute is written to the *LogDefaults* stanza of the *mqs.ini* file.

When a journal receiver extends beyond its specified threshold, the receiver is detached and a new journal receiver is created, inheriting attributes from the previous receiver. Changes to the LogReceiverSize or LogASP attributes after a queue manager has been created are ignored when the system automatically attaches a new journal receiver

See  Changing configuration information on IBM i (*WebSphere MQ V7.1 Installing Guide*) for further details on configuring the system.

If you need to change the size of journal receivers after the queue manager has been created, create a new journal receiver and set its owner to QMQM using the following commands:

```
CRTJRNRCV JRNRCV(QMGRLIB/AMQArnnnnn) THRESHOLD(xxxxxx) +  
  TEXT('MQM LOCAL JOURNAL RECEIVER')  
CHGOBJOWN OBJ(QMGRLIB/AMQArnnnnn) OBJTYPE(*JRNRCV) NEWOWN(QMQM)
```

where

QMGRLIB

Is the name of your queue manager library

A Is the instance identifier (usually A).

nnnnnn

Is the next journal receiver in the naming sequence described previously

xxxxxx Is the new receiver threshold (in KB)

Note: The maximum size of the receiver is governed by the operating system. To check this value look at the THRESHOLD keyword on the **CRTJRNRCV** command.

Now attach the new receiver to the AMQAJRN journal with the command:

```
CHGJRN JRN(QMGRLIB/AMQAJRN) JRNRCV(QMGRLIB/AMQAnnnnnn)
```

See “Journal management” on page 201 for details on how to manage these journal receivers.

WebSphere MQ for IBM i journal usage:

Use this information to understand how WebSphere MQ for IBM i uses journals in its operation to control updates to local objects.

Persistent updates to message queues happen in two stages. The records representing the update are first written to the journal, then the queue file is updated.

The journal receivers can therefore become more up to date than the queue files. To ensure that restart processing begins from a consistent point, WebSphere MQ uses checkpoints.

A checkpoint is a point in time when the record described in the journal is the same as the record in the queue. The checkpoint itself consists of the series of journal records needed to restart the queue manager. For example, the state of all transactions (that is, units of work) active at the time of the checkpoint.

Checkpoints are generated automatically by WebSphere MQ. They are taken when the queue manager starts and shuts down, and after a certain number of operations are logged.

You can force a queue manager to take a checkpoint by issuing the RCDQMIMG command against all objects on a queue manager and displaying the results, as follows:

```
RCDQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(<Q_MGR_NAME>) DSPJRNTA(*YES)
```

As the queues handle further messages, the checkpoint record becomes inconsistent with the current state of the queues.

When WebSphere MQ is restarted, it locates the latest checkpoint record in the log. This information is held in the checkpoint file that is updated at the end of every checkpoint. The checkpoint record represents the most recent point of consistency between the log and the data. The data from this checkpoint is used to rebuild the queues as they existed at the checkpoint time. When the queues are re-created, the log is then played forward to bring the queues back to the state they were in before system failure or close down.

To understand how WebSphere MQ uses the journal, consider the case of a local queue called TESTQ in the queue manager TEST. This is represented by the IFS file:

/QIBM/UserData/mqm/qmgrs/TEST/queues

If a specified message is put on this queue, and then retrieved from the queue, the actions that take place are shown in Figure Figure 44.

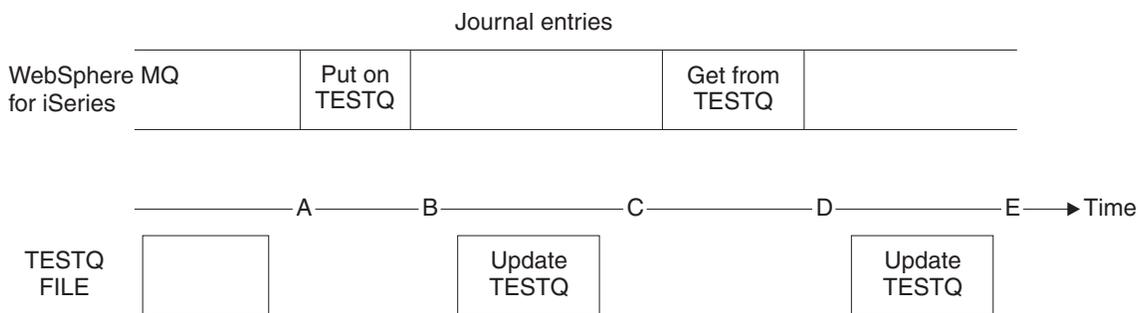


Figure 44. Sequence of events when updating MQM objects

The five points, A through E, shown in the diagram represent points in time that define the following states:

- A** The IFS file representation of the queue is consistent with the information contained in the journal.
- B** A journal entry is written to the journal defining a Put operation on the queue.
- C** The appropriate update is made to the queue.
- D** A journal entry is written to the journal defining a Get operation from the queue.
- E** The appropriate update is made to the queue.

The key to the recovery capabilities of WebSphere MQ for IBM i is that the user can save the IFS file representation of TESTQ as at time A, and subsequently recover the IFS file representation of TESTQ as at time E, by restoring the saved object and replaying the entries in the journal from time A onwards.

This strategy is used by WebSphere MQ for IBM i to recover persistent messages after system failure. WebSphere MQ remembers a particular entry in the journal receivers, and ensures that on startup it replays the entries in the journals from this point onwards. This startup entry is periodically recalculated so that WebSphere MQ only has to perform the minimum necessary replay on the next startup.

WebSphere MQ provides individual recovery of objects. All persistent information relating to an object is recorded in the local WebSphere MQ for IBM i journals. Any WebSphere MQ object that becomes damaged or corrupt can be completely rebuilt from the information held in the journal.

For more information on how the system manages receivers, see “Availability, backup, recovery, and restart” on page 195.

Media images:

Use this information to understand media images, and recovery from media images.

A WebSphere MQ object of long duration can represent a large number of journal entries, going back to the point at which it was created. To avoid this, WebSphere MQ for IBM i has the concept of a *media image* of an object.

This media image is a complete copy of the WebSphere MQ object recorded in the journal. If an image of an object is taken, the object can be rebuilt by replaying journal entries from this image onwards. The entry in the journal that represents the replay point for each WebSphere MQ object is referred to as its *media recovery entry*. WebSphere MQ keeps track of the:

- Media recovery entry for each queue manager object.
- Oldest entry from within this set (see error message AMQ7462 in “Journal management” on page 201 for details).

Images of the *CTLG object and the *MQM object are taken regularly because these objects are crucial to queue manager restart.

Images of other objects are taken when convenient. By default, images of **all** objects are taken when a queue manager is shut down using the **ENDMQM** command with parameter ENDCCTJOB(*YES). This operation can take a considerable amount of time for very large queue managers. If you need to shut down quickly, specify parameter RCDMQMIMG(*NO) with ENDCCTJOB(*YES). In such cases, you are recommended to record a complete media image in the journals after the queue manager has been restarted, using the following command:

```
RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(<Q_MGR_NAME>)
```

WebSphere MQ automatically records an image of an object, if it finds a convenient point at which an object can be compactly described by a small entry in the journal. However, this might never happen for some objects, for example, queues that consistently contain large numbers of messages.

Rather than allow the date of the oldest media recovery entry to continue for an unnecessarily long period, use the WebSphere MQ command RCDMQMIMG, which enables you to take an image of selected objects manually.

Recovery from media images

WebSphere MQ automatically recovers some objects from their media image if it is found that they are corrupt or damaged. In particular, this applies to the special *MQM and *CTLG objects as part of the normal queue manager startup. If any syncpoint transaction was incomplete at the time of the last shutdown of the queue manager, any queue affected is also recovered automatically, in order to complete the startup operation.

You must recover other objects manually, using the WebSphere MQ command RCRMQOBJ. This command replays the entries in the journal to re-create the WebSphere MQ object. Should a WebSphere MQ object become damaged, the only valid actions are to delete it or re-create it by this method. Note, however, that nonpersistent messages cannot be recovered in this fashion.

Checkpoints:

Checkpoints are taken at various times to provide a known consistent start point for recovery.

The checkpoint process AMQALMPX is responsible for taking the checkpoint at the following points:

- Queue manager startup (STRMQM).
- Queue manager shutdown (ENDMQM).
- After a period of time has elapsed since the last checkpoint (the default period is 30 minutes) and a minimum number of log records have been written since the previous checkpoint (the default value is 100).
- After a number of log records have been written. The default value is 10 000.
- After the journal threshold size has been exceeded and a new journal receiver has been automatically created.
- When a full media image is taken with:

```
RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(<Q_MGR_NAME>) DSPJRNDTA(*YES)
```

Backups of WebSphere MQ for IBM i data

Use this information to understand the two types of WebSphere MQ backup for each queue manager.

For each queue manager, there are two types of WebSphere MQ backup to consider:

- Data and journal backup.
To ensure that both sets of data are consistent, do this only after shutting down the queue manager.
- Journal backup.
You can do this while the queue manager is active.

For both methods, you need to find the names of the queue manager IFS directory and the queue manager library. You can find these in the WebSphere MQ configuration file (mq.ini). For more information, see  The QueueManager stanza.

Use the following procedures to do both types of backup:

Data and journal backup of a particular queue manager

Note: Do not use a save-while-active request when the queue manager is running. Such a request cannot complete unless all commitment definitions with pending changes are committed or rolled back. If this command is used when the queue manager is active, the channel connections might not end normally. Always use the following procedure.

1. Create an empty journal receiver, using the command:

```
CHGJRN JRN(QMTEST/AMQAJRN) JRNRCV(*GEN)
```
2. Use the **RCDMQMIMG** command to record an MQM image for all WebSphere MQ objects, and then force a checkpoint using the command:

```
RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) DSPJRNDTA(*YES) MQMNAME(TEST)
```
3. End channels and ensure that the queue manager is not running. If your queue manager is running, stop it with the **ENDMQM** command.
4. Backup the queue manager library by issuing the following command:

```
SAVLIB LIB(QMTEST)
```
5. Back up the queue manager IFS directories by issuing the following command:

```
SAV DEV(...) OBJ((' /QIBM/UserData/mqm/qmgrs/test '))
```

Journal backup of a particular queue manager

Because all relevant information is held in the journals, as long as you perform a full save at

some time, partial backups can be performed by saving the journal receivers. These record all changes since the time of the full backup and are performed by issuing the following commands:

1. Create an empty journal receiver, using the command:
`CHGJRN JRN(QMTEST/AMQAJRN) JRNRCV(*GEN)`
2. Use the **RCDMQMIMG** command to record an MQM image for all WebSphere MQ objects, and then force a checkpoint using the command:
`RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) DSPJRNDTA(*YES) MQMNAME(TEST)`
3. Save the journal receivers using the command:
`SAVOBJ OBJ(AMQ*) LIB(QMTEST) OBJTYPE(*JRNRCV)`

A simple backup strategy is to perform a full backup of the WebSphere MQ libraries every week, and perform a daily journal backup. This, of course, depends on how you have set up your backup strategy for your enterprise.

Journal management:

As part of your backup strategy, take care of your journal receivers. It is useful to remove journal receivers from the WebSphere MQ libraries for various reasons:

- To release space; this applies to all journal receivers
- To improve the performance when starting (STRMQM)
- To improve the performance of recreating objects (RCRMQMOBJ)

Before deleting a journal receiver, you must take care that you have a backup copy and that you no longer need the journal receiver.

Journal receivers can be removed from the queue manager library *after* they have been detached from the journals and saved, provided that they are available for restoration if needed for a recovery operation.

The concept of journal management is shown in Figure 45 on page 202.

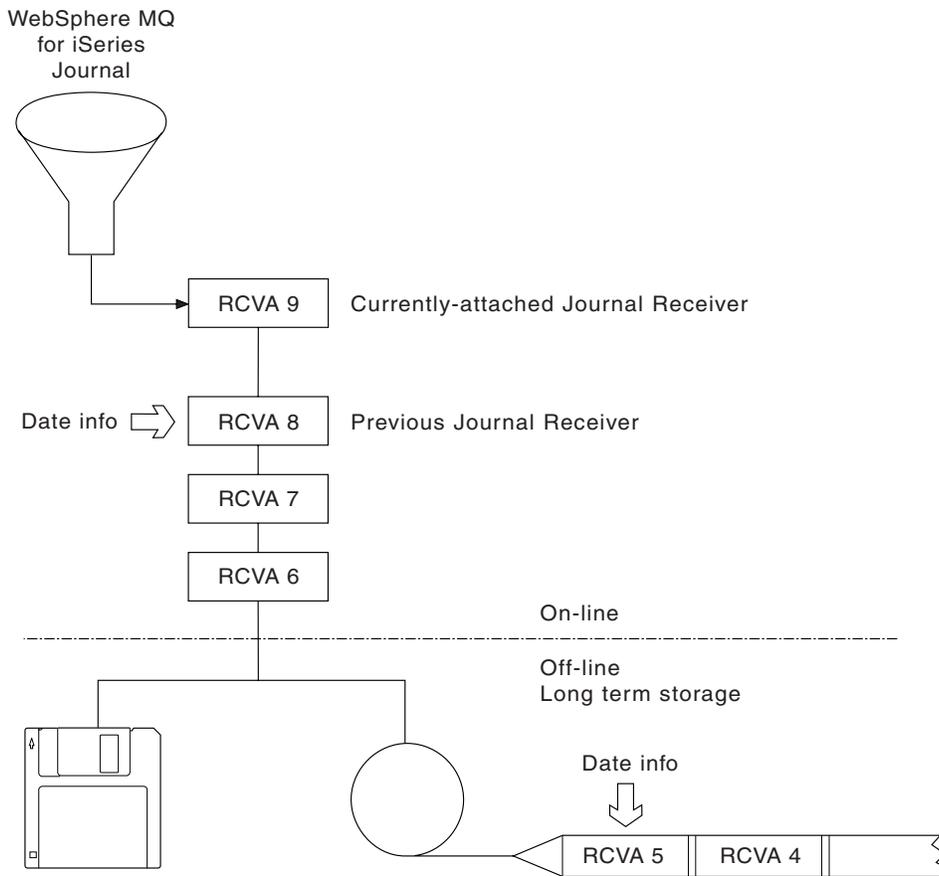


Figure 45. WebSphere MQ for IBM i journaling

It is important to know how far back in the journals WebSphere MQ is likely to need to go, in order to determine when a journal receiver that has been backed up can be removed from the queue manager library, and when the backup itself can be discarded.

WebSphere MQ issues two messages to the queue manager message queue (QMQMMSG in the queue manager library) to help determine this time. These messages are issued when it starts, when it changes a local journal receiver, and you use RCDMQIMG to force a checkpoint. The two messages are:

AMQ7460

Startup recovery point. This message defines the date and time of the startup entry from which WebSphere MQ replays the journal in the event of a startup recovery pass. If the journal receiver that contains this record is available in the WebSphere MQ libraries, this message also contains the name of the journal receiver containing the record.

AMQ7462

Oldest media recovery entry. This message defines the date and time of the oldest entry to use to re-create an object from its media image.

The journal receiver identified is the oldest one required. Any other WebSphere MQ journal receivers with older creation dates are no longer needed. If only stars are displayed, you need to restore backups from the date indicated to determine which is the oldest journal receiver.

When these messages are logged, WebSphere MQ also writes a user space object to the queue manager library that contains only one entry: the name of the oldest journal receiver that needs to be kept on the system. This user space is called AMQJRNINF, and the data is written in the format:

JJJJJJJJLLLLLLLLLLLLYYYYMMDDHHMMSSmmm

where:

JJJJJJJJJ

Is the oldest receiver name that WebSphere MQ still needs.

LLLLLLLLL

Is the journal receiver library name.

YYYY Is the year of the oldest journal entry that WebSphere MQ needs.

MM Is the month of the oldest journal entry that WebSphere MQ needs.

DD Is the day of the oldest journal entry that WebSphere MQ needs.

HH Is the hour of the oldest journal entry that WebSphere MQ needs.

SS Is the seconds of the oldest journal entry that WebSphere MQ needs.

mmm Is the milliseconds of the oldest journal entry that WebSphere MQ needs.

When the oldest journal receiver has been deleted from the system, this user space contains asterisks (*) for the journal receiver name.

Note: Periodically performing `RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) DSPJRNDTA(*YES)` can save startup time for WebSphere MQ and reduce the number of local journal receivers you need to save and restore for recovery.

WebSphere MQ for IBM i does not refer to the journal receivers unless it is performing a recovery pass either for startup, or for recreating an object. If it finds that a journal it requires is not present, it issues message AMQ7432 to the queue manager message queue (QMQMMSG), reporting the time and date of the journal entry it requires to complete the recovery pass.

If this happens, restore all journal receivers that were detached after this date from the backup, to allow the recovery pass to succeed.

Keep the journal receiver that contains the startup entry, and any subsequent journal receivers, available in the queue manager library.

Keep the journal receiver containing the oldest Media Recovery Entry, and any subsequent journal receivers, available at all times, and either present in the queue manager library or backed-up.

When you force a checkpoint:

- If the journal receiver named in AMQ7460 is not advanced, this indicates that there is an incomplete unit of work that needs to be committed or rolled back.
- If the journal receiver named in AMQ7462 is not advanced, this indicates that there are one or more damaged objects.

Restoring a complete queue manager (data and journals):

Use this information to restore one or more queue managers from a backup or from a remote machine.

If you need to recover one or more WebSphere MQ queue managers from a backup, perform the following steps.

1. Quiesce the WebSphere MQ queue managers.
2. Locate your latest backup set, consisting of your most recent full backup and subsequently backed up journal receivers.
3. Perform a RSTLIB operation, from the full backup, to restore the WebSphere MQ data libraries to their state at the time of the full backup, by issuing the following commands:

```
RSTLIB LIB(QMQLIB1) .....  
RSTLIB LIB(QMQLIB2) .....
```

If a journal receiver was partially saved in one journal backup, and fully saved in a subsequent backup, restore only the fully saved one. Restore journals individually, in chronological order.

4. Perform an RST operation to restore the WebSphere MQ IFS directories to the IFS file system, using the following command:

```
RST DEV(...) OBJ('/QIBM/UserData/mqm/qmgrs/testqm') ...
```

5. Start the message queue manager. This replays all journal records written since the full backup and restores all the WebSphere MQ objects to the consistent state at the time of the journal backup.

If you want to restore a complete queue manager on a different machine, use the following procedure to restore everything from the queue manager library. (We use TEST as the sample queue manager name.)

1. CRTMQM TEST
2. DLTLIB LIB(QMTEST)
3. RSTLIB SAVLIB(QMTEST) DEV(*SAVF) SAVF(QMGRLIBSAV)
4. Delete the following IFS files:
 - /QIBM/UserData/mqm/qmgrs/TEST/QMQMCHKPT
 - /QIBM/UserData/mqm/qmgrs/TEST/qmanager/QMQMOBJCAT
 - /QIBM/UserData/mqm/qmgrs/TEST/qmanager/QMANAGER
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.AUTH.DATA.QUEUE/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.CHANNEL.INITQ/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.CLUSTER.COMMAND.QUEUE/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.CLUSTER.REPOSITORY.QUEUE/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.CLUSTER.TRANSMIT.QUEUE/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.PENDING.DATA.QUEUE/q
 - /QIBM/UserData/mqm/qmgrs/TEST/queues/SYSTEM.ADMIN.COMMAND.QUEUE/q
5. STRMQM TEST
6. RCRMQMOBJ OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(TEST)

Restoring journal receivers for a particular queue manager:

Use this information to understand the different ways to restore journal receivers.

The most common action is to restore a backed-up journal receiver to a queue manager library, if a receiver that has been removed is needed again for a subsequent recovery function.

This is a simple task, and requires the journal receivers to be restored using the standard IBM i RSTOBJ command:

```
RSTOBJ OBJ(QMQMDATA/AMQA000005) OBJTYPE(*JRNRV) .....
```

A series of journal receivers might need to be restored, rather than a single receiver. For example, AMQA000007 is the oldest receiver in the WebSphere MQ libraries, and both AMQA000005 and AMQA000006 need to be restored.

In this case, restore the receivers individually in reverse chronological order. This is not always necessary, but is good practice. In severe situations, you might need to use the IBM i command WRKJRNA to associate the restored journal receivers with the journal.

When restoring journals, the system automatically creates an attached journal receiver with a new name in the journal receiver sequence. However, the new name generated might be the same as a journal receiver you need to restore. Manual intervention is needed to overcome this problem; to create a new name journal receiver in sequence, and new journal before restoring the journal receiver.

For example, consider the problem with saved journal AMQAJRN and the following journal receivers:

- AMQA000000
- AMQA100000
- AMQA200000
- AMQA300000
- AMQA400000
- AMQA500000
- AMQA600000
- AMQA700000
- AMQA800000
- AMQA900000

When restoring journal AMQAJRN to a queue manager library, the system automatically creates journal receiver AMQA000000. This automatically generated receiver conflicts with one of the existing journal receivers (AMQA000000) you want to restore, which you cannot restore.

The solution is:

1. Manually create the next journal receiver (see “WebSphere MQ for IBM i journals” on page 196):
`CRTJRNRCV JRNRCV(QMQLIB/AMQA9000001) THRESHOLD(XXXXX)`
2. Manually create the journal with the journal receiver:
`CRTJRN JRN(QMGLIB/AMQAJRN) MNGRCV(*SYSTEM) +
JRNRCV(QMGLIB/AMQA9000001) MSGQ(QMGLIB/AMQAJRNMSG)`
3. Restore the local journal receivers AMQA000000 to AMQA900000.

Multi-instance queue managers

Multi-instance queue managers improve availability by automatically switching to a standby server if the active server fails. The active and standby servers are multiple instances of the same queue manager; they share the same queue manager data. If the active instance fails you need to transfer its journal to the standby that takes over so that the queue manager can rebuild its queues.

Configure the IBM i systems you are running multi-instance queue managers on so that, if the active queue manager instance fails, the journal it is using is available to the standby instance that takes over. You can design your own configuration and administration tasks to make the journal from the active instance available to the instance that takes over. If you do not want to lose messages, your design must ensure the standby journal is consistent with the active journal at the point of failure. You can adapt your design from one of the two configurations that are described with examples in subsequent topics that do maintain consistency.

1. Mirror the journal from the system that is running the active queue manager instance to the systems that are running standby instances.
2. Place the journal in an Independent Auxiliary Storage Pool (IASP) that is transferable from the system running the active instance to a standby instance.

The first solution requires no additional hardware or software as it uses basic ASPs. The second solution requires switchable IASPs which need IBM i clustering support that is available as a separately priced IBM i License Product 5761-SS1 Option 41.

Reliability and availability:

Multi-instance queue managers aim to improve the availability of applications. Technological and physical constraints mean you need different solutions to meet the demands of disaster recovery, backing up queue managers and continuous operation.

In configuring for reliability and availability you trade off a large number of factors, resulting in four distinct design points:

Disaster recovery

Optimized for recovery after a major disaster that destroys all your local assets.

Disaster recovery on IBM i is often based on geographic mirroring of IASP.

Backup

Optimized for recovery after a localized failure, commonly a human error or some unforeseen technical problem.

WebSphere MQ provides backup queue managers to back up queue managers periodically. You could also use asynchronous replication of queue manager journals to improve the currency of the backup.

Availability

Optimized for restoring operations quickly giving the appearance of a nearly uninterrupted service following foreseeable technical failures such as a server or disk failure.

Recovery is typically measured in minutes, with detection sometimes taking longer than the recovery process. A multi-instance queue manager assists you in configuring for *availability*.

Continuous operation

Optimized for providing an uninterrupted service.

Continuous operation solutions have to solve the detection problem, and nearly always involve submitting the same work through more than one system and either using the first result, or if correctness is a major consideration, comparing at least two outcomes.

A multi-instance queue manager assists you in configuring for *availability*. One instance of the queue manager is active at a time. Switching over to a standby instance takes from a little more than ten seconds to a fifteen minutes or more, depending on how the system is configured, loaded and tuned.

A multi-instance queue manager can give the appearance of a nearly uninterrupted service if used with reconnectable WebSphere MQ MQI clients, which are able to continue processing without the application program necessarily being aware of a queue manager outage; see the topic  Automated client reconnection (*WebSphere MQ V7.1 Installing Guide*).

Components of a high availability solution:

Construct a high availability solution using multi-instance queue managers by providing robust networked storage for queue manager data, journal replication or robust IASP storage for queue manager journals, and using reconnectable clients, of applications configured as restartable queue manager services.

A multi-instance queue manager reacts to the detection of queue manager failure by resuming the startup of another queue manager instance on another server. To complete its startup, the instance needs access to the shared queue manager data in networked storage, and to its copy of the local queue manager journal.

To create a high availability solution, you need to manage the availability of the queue manager data, the currency of the local queue manager journal, and either build reconnectable client applications, or deploy

your applications as queue manager services to restart automatically when the queue manager resumes. Automatic client reconnect is not supported by WebSphere MQ classes for Java.

Queue manager data

Place queue manager data onto networked storage that is shared, highly available and reliable, possibly by using RAID level 1 disks or greater. The file system needs to meet the requirements for a shared file system for multi-instance queue managers; for more information about the requirements for shared file systems, see  Requirements for shared file systems (*WebSphere MQ V7.1 Installing Guide*). Network File System Version 4 (NFS4) is a protocol that meets these requirements.

Queue manager journals

You also need to configure the IBM i journals used by the queue manager instances so that the standby instance is able to restore its queue manager data to a consistent state. For uninterrupted service, this means you must restore the journals to their state when the active instance failed. Unlike backup or disaster recovery solutions, restoring journals to an earlier checkpoint is not sufficient.

You cannot physically share journals between multiple IBM i systems on networked storage. To restore queue manager journals to the consistent state at the point of failure, you either need to transfer the physical journal that was local to the active queue manager instance at the time of failure to the new instance that has been activated, or maintain mirrors of the journal on running standby instances. The mirrored journal is a remote journal replica that has been kept exactly in sync with the local journal belonging to the failed instance.

Three configurations are starting points for designing how you manage the journals for a multi-instance queue manager,

1. Using synchronized journal replication (journal mirroring) from the active instance ASP, to the standby instances ASPs.
2. Transferring an IASP you have configured to hold the queue manager journal from the active instance to the standby instance that is taking over as the active instance.
3. Using synchronized secondary IASP mirrors.

See the ASP options in the WebSphere MQ  CRTMQM command, for more information on putting queue manager data onto an iASP.

Also, see  Administrator > High Availability.

Applications

To build a client to automatically reconnect to the queue manager when the standby queue manager resumes, connect your application to the queue manager using MQCONN and specify

MQCNO_RECONNECT_Q_MGR in the MQCNO Options field. See,  High availability sample programs (*WebSphere MQ V7.1 Programming Guide*) for three sample programs using reconnectable clients, and  Application recovery (*WebSphere MQ V7.1 Installing Guide*) for information about designing client applications for recovery.

Creating a network share for queue manager data using NetServer:

Create a network share on an IBM i server for storing queue manager data. Set up connections from two servers, which are going to host queue manager instances, to access the network share.

Before you begin

- You require 3 IBM i servers for this task. The network share is defined on one of the servers, GAMMA. The other two servers, ALPHA and BETA, are to connect to GAMMA.
- The task is most easily performed, if you install WebSphere MQ on all three servers.
- Install the System i[®] Navigator; see  System i Navigator.

About this task

- Create the queue manager directory on GAMMA and set the correct ownership and permissions for the user profiles QMQM and QMQMADM. The directory and permission are easily created by installing WebSphere MQ on GAMMA.
- Use System i Navigator to create a share to the queue manager data directory on GAMMA.
- Create directories on ALPHA and BETA that point to the share.

Procedure

1. On GAMMA create the directory to host the queue manager data with the QMQM user profile as the owner, and QMQMADM as the primary group.

Tip:

A quick and reliable way to create the directory with the right permissions is to install WebSphere MQ on GAMMA.

Later, if you do not want to run WebSphere MQ on GAMMA, uninstall WebSphere MQ. After uninstallation, the directory /QIBM/UserData/mqm/qmgrs remains on GAMMA with the owner QMQM user profile, and QMQMADM the primary group.

The task uses the /QIBM/UserData/mqm/qmgrs directory on GAMMA for the share.

2. Start the System i Navigator **Add connection** wizard and connect to the GAMMA system.
 - a. Double-click the **System i Navigator** icon on your Windows desktop.
 - b. Click **Yes** to create a connection.
 - c. Follow the instructions in the **Add Connection** wizard and create a connection from the IBM i system to GAMMA.

The connection to GAMMA is added to **My Connections**.
3. Add a new file share on GAMMA.
 - a. In the System i Navigator window, click the File Shares folder in My Connections/GAMMA/File Systems.
 - b. In the My Tasks window, click **Manage i5/OS NetServer shares**.

A new window, i5/OS NetServer - GAMMA, opens on your desktop and shows shared objects.
 - c. Right-click the Shared Objects folder > **File** > **New** > **File**.

A new window, i5/OS NetServer File Share - GAMMA, opens.
 - d. Give the share a name, WMQ for example.
 - e. Set the access control to Read/Write.
 - f. Select the **Path name** by browsing to the /QIBM/UserData/mqm/qmgrs directory you created earlier, and click **OK**.

The i5/OS NetServer File Share - GAMMA window closes, and WMQ is listed in the shared objects window.

4. Right click **WMQ** in the shared objects window. Click **File > Permissions**.

A window opens, Qmgrs Permissions - GAMMA, for the object /QIBM/UserData/mqm/qmgrs.

- a. Check the following permissions for QMQM, if they are not already set:

- Read
- Write
- Execute
- Management
- Existence
- Alter
- Reference

- b. Check the following permissions for QMQMADM, if they are not already set:

- Read
- Write
- Execute
- Reference

- c. Add other user profiles that you want to give permissions to /QIBM/UserData/mqm/qmgrs.

For example, you might give the default user profile (Public) Read and Execute permissions to /QIBM/UserData/mqm/qmgrs.

5. Check that all the user profiles that are granted access to /QIBM/UserData/mqm/qmgrs on GAMMA have the same password as they do on the servers that access GAMMA.

In particular, ensure that the QMQM user profiles on other servers, which are going to access the share, have the same password as the QMQM user profile on GAMMA.

Tip: Click the My Connections/GAMMA/Users and Groups folder in the System i Navigator to set the passwords. Alternatively, use the **CHFUSRPRF** and **CHGPWD** commands.

Results

Check you can access GAMMA from other servers using the share. If you are doing the other tasks, check you can access GAMMA from ALPHA and BETA using the path /QNTC/GAMMA/WMQ. If the /QNTC/GAMMA directory does not exist on ALPHA or BETA then you must create the directory. Depending on the NetServer domain, you might have to IPL ALPHA or BETA before creating the directory.

```
CRTDIR DIR('/QNTC/GAMMA')
```

When you have checked that you have access to /QNTC/GAMMA/WMQ from ALPHA or BETA, issuing the command, `CRTMQM MQMNAME('QM1') MQMDIRP('/QNTC/GAMMA/WMQ')` creates /QIBM/UserData/mqm/qmgrs/QM1 on GAMMA.

What to do next

Create a multi-instance queue manager by following the steps in either of the tasks, “Creating a multi-instance queue manager using journal mirroring and NetServer” on page 216 or “Converting a single instance queue manager to a multi-instance queue manager using NetServer and journal mirroring” on page 220.

Fail over performance:

The time it takes to detect a queue manager instance has failed, and then to resume processing on a standby can vary between tens of seconds to fifteen minutes or more depending on the configuration. Performance needs to be a major consideration in designing and testing a high availability solution.

There are advantages and disadvantages to weigh up in deciding whether to configure a multi-instance queue manager to use journal replication, or to use an IASP. Mirroring requires the queue manager to write synchronously to a remote journal. From a hardware point of view, this need not affect performance, but from a software perspective there is a greater pathlength involved in writing to a remote journal than just to a local journal, and this might be expected to reduce the performance of a running queue manager to some extent. However, when the standby queue manager takes over, the delay in synchronizing its local journal from the remote journal maintained by the active instance before it failed, is typically small in comparison to the time it takes for IBM i to detect and transfer the IASP to the server running the standby instance of the queue manager. IASP transfer times can be as much as ten to fifteen minutes rather than being completed in seconds. The IASP transfer time depends on the number of objects that need to be *varied-on* when the IASP is transferred to the standby system and the size of the access paths, or indexes, that need to be merged.

When the standby queue manager takes over, the delay in synchronizing its local journal from the remote journal maintained by the active instance before it failed, is typically small in comparison to the time it takes for IBM i to detect and transfer the independent ASP to the server running the standby instance of the queue manager. Independent ASP transfer times can be as much as ten to fifteen minutes rather than being completed in seconds. The independent ASP transfer time depends on the number of objects that need to be *varied-on* when the independent ASP is transferred to the standby system and the size of the access paths, or indices, that need to be merged.

However, transferring the journal is not the only factor influencing the time it takes for the standby instance to fully resume. You also need to consider the time it takes for the network file system to release the lock on queue manager data that signals to the standby instance to try to continue with its start-up, and also the time it takes to recover queues from the journal so that the instance is able to start processing messages again. These other sources of delay all add to the time it takes to start a standby instance. The total time to switch over consists of the following components,

Failure detection time

The time it takes for NFS to release the lock on the queue manager data, and the standby instance to continue its startup process.

Transfer time

In the case of an HA cluster, the time it takes IBM i to transfer the IASP from the system hosting the active instance to the standby instance, and in the case of journal replication, the time it takes to update the local journal at the standby with the data from the remote replica.

Restart time

The time it takes for the newly active queue manager instance to rebuild its queues from the latest checkpoint in its restored journal and to resume processing messages.

Note:

If the standby instance that has taken over is configured to synchronously replicate to the previously active instance, the startup could be delayed. The new activated instance might be unable to replicate to its remote journal, if the remote journal is on the server that hosted the previously active instance, and the server has failed.

The default time to wait for a synchronous response is one minute. You can configure the maximum delay before the replication times out. Alternatively, you can configure standby instances to start using asynchronous replication to the failed active instance. Later you switch

the to synchronous replication, when the failed instance is running on standby again. The same consideration applies to using synchronous independent ASP mirrors.

You can make separate baseline measurements for these components to help you assess the overall time to failover, and to factor into your decision which configuration approach to use. In making the best configuration decision you also need to consider how other applications on the same server will failover, and whether there are backup or disaster recovery processes that already use IASP.

IASP transfer times can be shortened by tuning your cluster configuration:

1. User profiles across systems in the cluster should have the same GID and UID to eliminate the need for the vary-on process to change UIDs and GIDs.
2. Minimize the number of database objects in the system and basic user disk pools, as these need to be merged to create the cross-reference table for the disk-pool group.
3. Further performance tips can be found in the IBM Redbook, *Implementing PowerHA for IBM i, SG24-7405*.

A configuration using basic ASPs, journal mirroring, and a small configuration should switch over in the order of tens of seconds.

Mirrored journal configuration for ASP:

Configure a robust multi-instance queue manager using synchronous replication between mirrored journals.

A mirrored queue manager configuration uses journals that are created in basic or independent auxiliary storage pools (ASP).

On IBM i, queue manager data is written to journals and to a file system. Journals contain the master copy of queue manager data. Journals are shared between systems using either synchronous or asynchronous journal replication. A mix of local and remote journals are required to restart a queue manager instance. Queue manager restart reads journal records from the mix of local and remote journals on the server, and the queue manager data on the shared network file system. The data in the file system speeds up restarting the queue manager. Checkpoints are stored in the file system, marking points of synchronization between the file system and the journals. Journal records stored before the checkpoint are not required for typical queue manager restarts. However, the data in the file system might not be up to date, and journal records after the checkpoint are used to complete the queue manager restart. The data in the journals attached to the instance are kept up to date so that the restart can complete successfully.

But even the journal records might not be up to date, if the remote journal on the standby server was being asynchronously replicated, and the failure occurred before it was synchronized. In the event that you decide to restart a queue manager using a remote journal that is not synchronized, the standby queue manager instance might either reprocess messages that were deleted before the active instance failed, or not process messages that were received before the active instance failed.

Another, rare possibility, is that the file system contains the most recent checkpoint record, and an unsynchronized remote journal on the standby does not. In this case the queue manager does not restart automatically. You have a choice of waiting until the remote journal is synchronized, or cold starting the standby queue manager from the file system. Even though, in this case, the file system contains a more recent checkpoint of the queue manager data than the remote journal, it might not contain all the messages that were processed before the active instance failed. Some messages might be reprocessed, and some not processed, after a cold restart that is out of synchronization with the journals.

With a multi-instance queue manager, the file system is also used to control which instance of a queue manager is active, and which is the standby. The active instance acquires a lock to the queue manager data. The standby waits to acquire the lock, and when it does, it becomes the active instance. The lock is

released by the active instance, if it ends normally. The lock is released by the file system if the file system detects the active instance has failed, or cannot access the file system. The file system must meet the requirements for detecting failure; see  Requirements for shared file systems (*WebSphere MQ V7.1 Installing Guide*).

The architecture of multi-instance queue managers on IBM i provides automatic restart following server or queue manager failure. It also supports restoration of queue manager data following failure of the file system where the queue manager data is stored.

In Figure 46, if ALPHA fails, you can manually restart QM1 on beta, using the mirrored journal. By adding the multi-instance queue manager capability to QM1, the standby instance of QM1 resumes automatically on BETA if the active instance on ALPHA fails. QM1 can also resume automatically if it is the server ALPHA that fails, not just the active instance of QM1. Once BETA becomes the host of the active queue manager instance, the standby instance can be started on ALPHA.

Figure 46 shows a configuration that mirrors journals between two instances of a queue manager using NetServer to store queue manager data. You might expand the pattern to include more journals, and hence more instances. Follow the journal naming rules explained in the topic, “WebSphere MQ for IBM i journals” on page 196. Currently the number of running instances of a queue manager is limited to two, one is active and one is in standby.

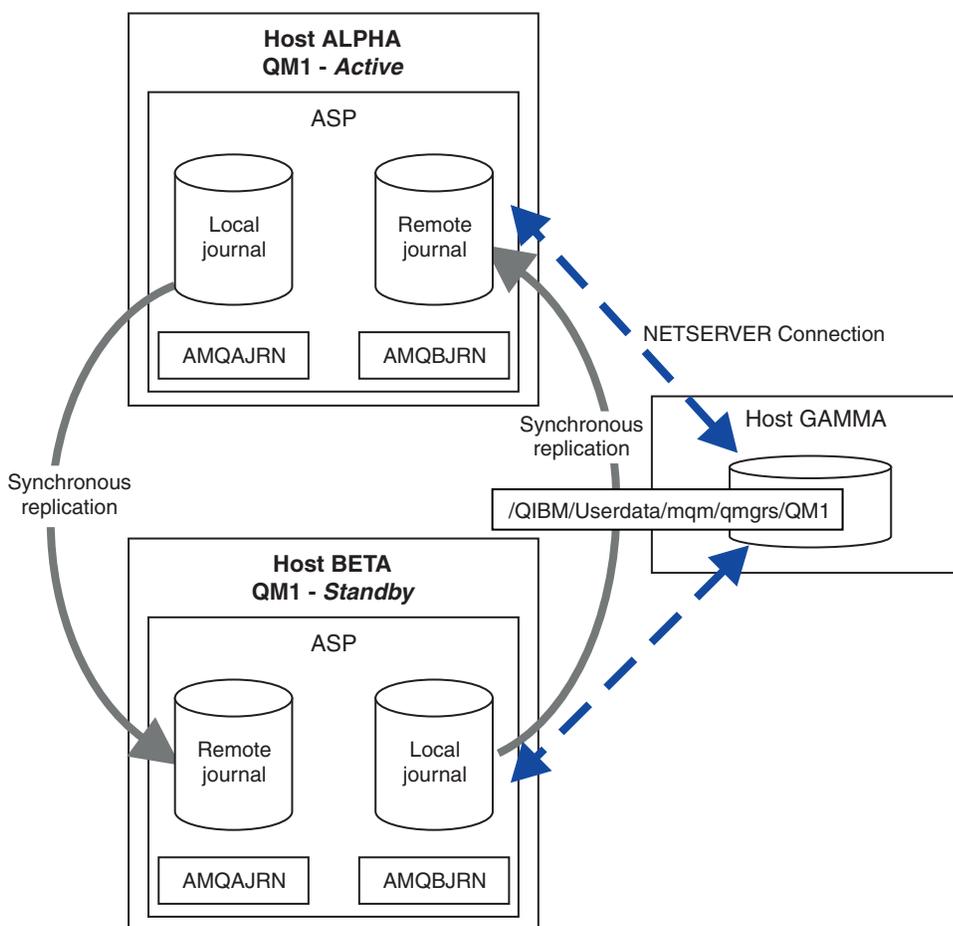


Figure 46. Mirror a queue manager journal

The local journal for QM1 on host ALPHA is called AMQAJRN (or more fully, QMQM1/AMQAJRN) and on BETA the journal is QMQM1/AMQBJRN. Each local journal replicates to remote journals on all other instances of the queue manager. If the queue manager is configured with two instances, a local journal is replicated to one remote journal.

***SYNC or *ASYNCR remote journal replication**

IBM i journals are mirrored using either synchronous (*SYNC) or asynchronous (*ASYNCR) journaling; see

Remote journal management.

The replication mode in Figure 46 on page 212 is *SYNC, not *ASYNCR. *ASYNCR is faster, but if a failure occurs when the remote journal state is *ASYNCPEND, the local and remote journal are not consistent. The remote journal must catch up with the local journal. If you choose *SYNC, then the local system waits for the remote journal before returning from a call that requires a completed write. The local and remote journals generally remain consistent with one another. Only if the *SYNC operation takes longer than a designated time¹, and remote journaling is deactivated, do the journals get out of synchronization. An error is logged to the journal message queue and to QSYSOPR. The queue manager detects this message, writes an error to the queue manager error log, and deactivates remote replication of the queue manager journal. The active queue manager instance resumes without remote journaling to this journal. When the remote server is available again, you must manually reactivate synchronous remote journal replication. The journals are then resynchronized.

A problem with the *SYNC/*ASYNCR configuration illustrated in Figure 46 on page 212 is how the standby queue manager instance on BETA takes control. As soon as the queue manager instance on BETA writes its first persistent message, it attempts to update the remote journal on ALPHA. If the cause of control passing from ALPHA to BETA was the failure of ALPHA, and ALPHA is still down, remote journaling to ALPHA fails. BETA waits for ALPHA to respond, and then deactivates remote journaling and resumes processing messages with only local journaling. BETA has to wait a while to detect that ALPHA is down, causing a period of inactivity.

The choice between setting remote journaling to *SYNC or *ASYNCR is a trade-off. Table 7 summarizes the trade-offs between using *SYNC and *ASYNCR journaling between a pair of queue managers:

Table 7. Remote journaling options

	Standby	*SYNC	*ASYNCR
Active			
*SYNC		<ol style="list-style-type: none"> 1. Consistent switchover and failover 2. The standby instance does not resume immediately after failover. 3. Remote journaling must be available all the time 4. Queue manager performance depends on remote journaling 	<ol style="list-style-type: none"> 1. Consistent switchover and failover 2. Remote journaling must be switched to *SYNC when standby server available 3. Remote journaling must remain available after it has been restarted 4. Queue manager performance depends on remote journaling
*ASYNCR		<ol style="list-style-type: none"> 1. Not a sensible combination 	<ol style="list-style-type: none"> 1. Some messages might be lost or duplicated after a failover or switchover 2. Standby instance need not be available all the time for the active instance to continue without delay. 3. Performance does not depend on remote journaling

1. The designated time is 60 seconds on i5/OS Version 5 and in the range 1 - 3600 seconds on IBM i 6.1 onwards.

***SYNC/*SYNC**

The active queue manager instance uses *SYNC journaling, and when the standby queue manager instance starts, it immediately tries to use *SYNC journaling.

1. The remote journal is transactionally consistent with the local journal of the active queue manager. If the queue manager is switched over to the standby instance, it can resume immediately. The standby instance normally resumes without any loss or duplication of messages. Messages are only lost or duplicated if remote journaling failed since the last checkpoint, and the previously active queue manager cannot be restarted.
2. If the queue manager fails over to the standby instance, it might not be able to start immediately. The standby queue manager instance is activated with *SYNC journaling. The cause of the failover might prevent remote journaling to the server hosting the standby instance. The queue manager waits until the problem is detected before processing any persistent messages. An error is logged to the journal message queue and to QSYSOPR. The queue manager detects this message, writes an error to the queue manager error log, and deactivates remote replication of the queue manager journal. The active queue manager instance resumes without remote journaling to this journal. When the remote server is available again, you must manually reactivate synchronous remote journal replication. The journals are then resynchronized.
3. The server to which the remote journal is replicated must always be available to maintain the remote journal. The remote journal is typically replicated to the same server that hosts the standby queue manager. The server might become unavailable. An error is logged to the journal message queue and to QSYSOPR. The queue manager detects this message, writes an error to the queue manager error log, and deactivates remote replication of the queue manager journal. The active queue manager instance resumes without remote journaling to this journal. When the remote server is available again, you must manually reactivate synchronous remote journal replication. The journals are then resynchronized.
4. Remote journaling is slower than local journaling, and substantially slower if the servers are separated by a large distance. The queue manager must wait for remote journaling, which reduces queue manager performance.

The *SYNC/*SYNC configuration between a pair of servers has the disadvantage of a delay in resuming the standby instance after failover. The *SYNC/*ASYNCR configuration does not have this problem.

*SYNC/*SYNC does guarantee no message loss after switchover or failover, as long as a remote journal is available. If you want to reduce the risk of message loss after failover or switchover you have two choices. Either stop the active instance if the remote journal becomes inactive, or create remote journals on more than one server.

***SYNC/*ASYNCR**

The active queue manager instance uses *SYNC journaling, and when the standby queue manager instance starts, it uses *ASYNCR journaling. Shortly after the server hosting the new standby instance becomes available, the system operator must switch the remote journal on the active instance to *SYNC. When the operator switches remote journaling from *ASYNCR to *SYNC the active instance pauses if the status of the remote journal is *ASYNCRPEND. The active queue manager instance waits until remaining journal entries are transferred to the remote journal. When the remote journal has synchronized with the local journal, the new standby is transactionally consistent again with the new active instance. From the perspective of the management of multi-instance queue managers, in an *SYNC/*ASYNCR configuration the IBM i system operator has an additional task. The operator must switch remote journaling to *SYNC in addition to restarting the failed queue manager instance.

1. The remote journal is transactionally consistent with the local journal of the active queue manager. If the active queue manager instance is switched over, or fails over to the standby instance, the standby instance can then resume immediately. The standby instance normally

resumes without any loss or duplication of messages. Messages are only lost or duplicated if remote journaling failed since the last checkpoint, and the previously active queue manager cannot be restarted.

2. The system operator must switch remote journal from *ASync to *Sync shortly after the system hosting the active instance becomes available again. The operator might wait for the remote journal to catch up before switching the remote journal to *Sync. Alternatively the operator might switch the remote instance to *Sync immediately, and force the active instance to wait until the standby instance journal has caught up. When remote journaling is set to *Sync, the standby instance is generally transactionally consistent with the active instance. Messages are only lost or duplicated if remote journaling failed since the last checkpoint, and the previously active queue manager cannot be restarted.
3. When the configuration has been restored from a switchover or failover, the server on which the remote journal is hosted must be available all the time.

Choose *Sync/*ASync when you want the standby queue manager to resume quickly after a failover. You must restore the remote journal setting to *Sync on the new active instance manually. The *Sync/*ASync configuration matches the normal pattern of administering a pair of multi-instance queue managers. After one instance has failed, there is a time before the standby instance is restarted, during which the active instance cannot fail over.

***ASync/*ASync**

Both the servers hosting the active and standby queue managers are configured to use *ASync remote journaling.

1. When switchover or failover take place, the queue manager continues with the journal on the new server. The journal might not be synchronized when the switchover or failover takes place. Consequently messages might be lost or duplicated.
2. The active instance runs, even if the server hosting the standby queue manager is not be available. The local journal is replicated asynchronously with the standby server when it is available.
3. The performance of the local queue manager is unaffected by remote journaling.

Choose *ASync/*ASync if performance is your principal requirement, and you are prepared to lose or duplicate some messages after failover or switchover.

***ASync/*Sync**

There is no reason to use this combination of options.

Queue manager activation from a remote journal

Journals are either replicated synchronously or asynchronously. The remote journal might not be active, or it might be catching up with the local journal. The remote journal might be catching up, even if it is synchronously replicated, because it might have been recently activated. The rules that the queue manager applies to the state of the remote journal it uses during start-up are as follows.

1. Standby startup fails if it must replay from the remote journal on the standby and the journal status is *FAILED or *INACTPEND.
2. When activation of the standby begins, the remote journal status on the standby must be either *ACTIVE or *INACTIVE. If the state is *INACTIVE, it is possible for activation to fail, if not all the journal data has been replicated.

The failure occurs if the queue manager data on the network file system has a more recent checkpoint record than present in the remote journal. The failure is unlikely to happen, as long as the remote journal is activated well within the default 30 minute maximum interval between checkpoints. If the standby queue manager does read a more recent checkpoint record from the file system, it does not start.

You have a choice: Wait until the local journal on the active server can be restored, or cold start the standby queue manager. If you choose to cold start, the queue manager starts with no journal data, and relies on the consistency and completeness of the queue manager data in the file system.

Note: If you cold start a queue manager, you run the risk of losing or duplicating messages after the last checkpoint. The message transactions were written to the journal, but some of the transactions might not have been written to the queue manager data in the file system. When you cold start a queue manager, a fresh journal is started, and transactions not written to the queue manager data in the file system are lost.

3. The standby queue manager activation waits for the remote journal status on the standby to change from *ASYNCPEND or *SYNCPEND to *ASYNCR or *SYNCR. Messages are written to the job log of the execution controller periodically.

Note: In this case activation is waiting on the remote journal local to the standby queue manager that is being activated. The queue manager also waits for a time before continuing without a remote journal. It waits when it tries to write synchronously to its remote journal (or journals) and the journal is not available.

4. Activation stops if the journal status changes to *FAILED or *INACTPEND.

The names and states of the local and remote journals to be used in the activation are written to the queue manager error log.

Creating a multi-instance queue manager using journal mirroring and NetServer:

Create a multi-instance queue manager to run on two IBM i servers. The queue manager data is stored on a third IBM i server using NetServer. The queue manager journal is mirrored between the two servers using remote journaling. The **ADDQMJRN** command is used to simplify creating the remote journals.

Before you begin

1. The task requires three IBM i servers. Install WebSphere MQ on two of them, ALPHA and BETA in the example. WebSphere MQ must be at least at version 7.0.1.1.
2. The third server is an IBM i server, connected by NetServer to ALPHA and BETA. It is used to share the queue manager data. It does not have to have a WebSphere MQ installation. It is useful to install WebSphere MQ on the server as a temporary step, to set up the queue manager directories and permissions.
3. Make sure that the QMQM user profile has the same password on all three servers.
4. Install IBM i NetServer; see [IBM i NetServer](#).

About this task

Perform the following steps to create the configuration shown in Figure 47 on page 219. The queue manager data is connected using IBM i NetServer.

- Create connections from ALPHA and BETA to the directory share on GAMMA that is to store the queue manager data. The task also sets up the necessary permissions, user profiles and passwords.
- Add Relational Database Entries (RDBE) to the IBM i systems that are going to run queue manager instances. The RDBE entries are used to connect to the IBM i systems used for remote journaling.
- Create the queue manager QM1 on the IBM i server, ALPHA.
- Add the queue manager control information for QM1 on the other IBM i server, BETA.
- Create remote journals on both the IBM i servers for both queue manager instances. Each queue manager writes to the local journal. The local journal is replicated to the remote journal. The command, **ADDQMJRN** simplifies adding the journals and the connections.
- Start the queue manager, permitting a standby instance.

Procedure

1. Do the task, "Creating a network share for queue manager data using NetServer" on page 208.

As a result, ALPHA and BETA have a share, /QNTC/GAMMA/WMQ, that points to /QIBM/UserData/mqm/qmgrs on GAMMA. The user profiles QMQM and QMQMADM have the necessary permissions, and QMQM has matching passwords on all three systems.

2. Add Relational Database Entries (RDBE) to the IBM i systems that are going to host queue manager instances.

- a. On ALPHA create the connection to BETA.

```
ADDRDBDIRE RDB(BETA) RMTLOCNAME(BETA *IP) RMTAUTMTH(*USRIDPWD)
```

- b. On BETA create the connections to ALPHA.

```
ADDRDBDIRE RDB(ALPHA) RMTLOCNAME(ALPHA *IP) RMTAUTMTH(*USRIDPWD)
```

3. Create the queue manager QM1 on ALPHA, saving the queue manager data on GAMMA.

```
CRTMQM MQMNAME(QM1) UDLMSGQ(SYSTEM.DEAD.LETTER.QUEUE)
MQMDIRP('/QNTC/GAMMA/WMQ')
```

The path, /QNTC/GAMMA/WMQ, uses NetServer to create the queue manager data in /QIBM/UserData/mqm/qmgrs.

4. Run **ADDQMJRN** on ALPHA. The command adds a remote journal on BETA for QM1.

```
ADDQMJRN MQMNAME(QM1) RMTJRN RDB(BETA)
```

QM1 creates journal entries in its local journal on ALPHA when the active instance of QM1 is on ALPHA. The local journal on ALPHA is replicated to the remote journal on BETA.

5. Use the command, **DSPF**, to inspect the WebSphere MQ configuration data created by **CRTMQM** for QM1 on ALPHA.

The information is needed in the next step.

In this example, the following configuration is created in /QIBM/UserData/mqm/mqs.ini on ALPHA for QM1:

```
Name=QM1
Prefix=/QIBM/UserData/mqm
Library=QMQM1
Directory=QM1
DataPath=/QNTC/GAMMA/WMQ/QM1
```

6. Create a queue manager instance of QM1 on BETA using the **ADDQMINF** command. Run the following command on BETA to modify the queue manager control information in /QIBM/UserData/mqm/mqs.ini on BETA.

```
ADDQMINF MQMNAME(QM1)
PREFIX('/QIBM/UserData/mqm')
MQMDIR(QM1)
MQMLIB(QMQM1)
DATAPATH('/QNTC/GAMMA/WMQ/QM1')
```

Tip: Copy and paste the configuration information. The queue manager stanza is the same on ALPHA and BETA.

7. Run **ADDQMJRN** on BETA. The command adds a local journal on BETA and a remote journal on ALPHA for QM1.

```
ADDQMJRN MQMNAME(QM1) RMTJRN RDB(ALPHA)
```

QM1 creates journal entries in its local journal on BETA when the active instance of QM1 is on BETA. The local journal on BETA is replicated to the remote journal on ALPHA.

Note: As an alternative, you might want to set up remote journaling from BETA to ALPHA using asynchronous journaling.

Use this command to set up asynchronous journaling from BETA to ALPHA, instead of the command in step 7.

```
ADDQMJRN MQMNAME(QM1) RMTJRN RDB(ALPHA) RMTJRNDLV(*ASYNCL)
```

If the server or journaling on ALPHA is the source of the failure, BETA starts without waiting for new journal entries to be replicated to ALPHA.

Switch the replication mode to *SYNC, using the **CHGMQMJRN** command, when ALPHA is online again.

Use the information in “Mirrored journal configuration for ASP” on page 211 to decide whether to mirror the journals synchronously, asynchronously, or a mixture of both. The default is to replicate synchronously, with a 60 second wait period for a response from the remote journal.

8. Verify that the journals on ALPHA and BETA are enabled and the status of remote journal replication is *ACTIVE.
 - a. On ALPHA:
WRKMQMJRN MQMNAME(QM1)
 - b. On BETA:
WRKMQMJRN MQMNAME(QM1)
9. Start the queue manager instances on ALPHA and BETA.
 - a. Start the first instance on ALPHA, making it the active instance. Enabling switching over to a standby instance.
STRMQM MQMNAME(QM1) STANDBY(*YES)
 - b. Start the second instance on BETA, making it the standby instance.
STRMQM MQMNAME(QM1) STANDBY(*YES)

Results

Use **WRKMQM** to check queue manager status:

1. The status of the queue manager instance on ALPHA should be *ACTIVE.
2. The status of the queue manager instance on BETA should be *STANDBY.

Example

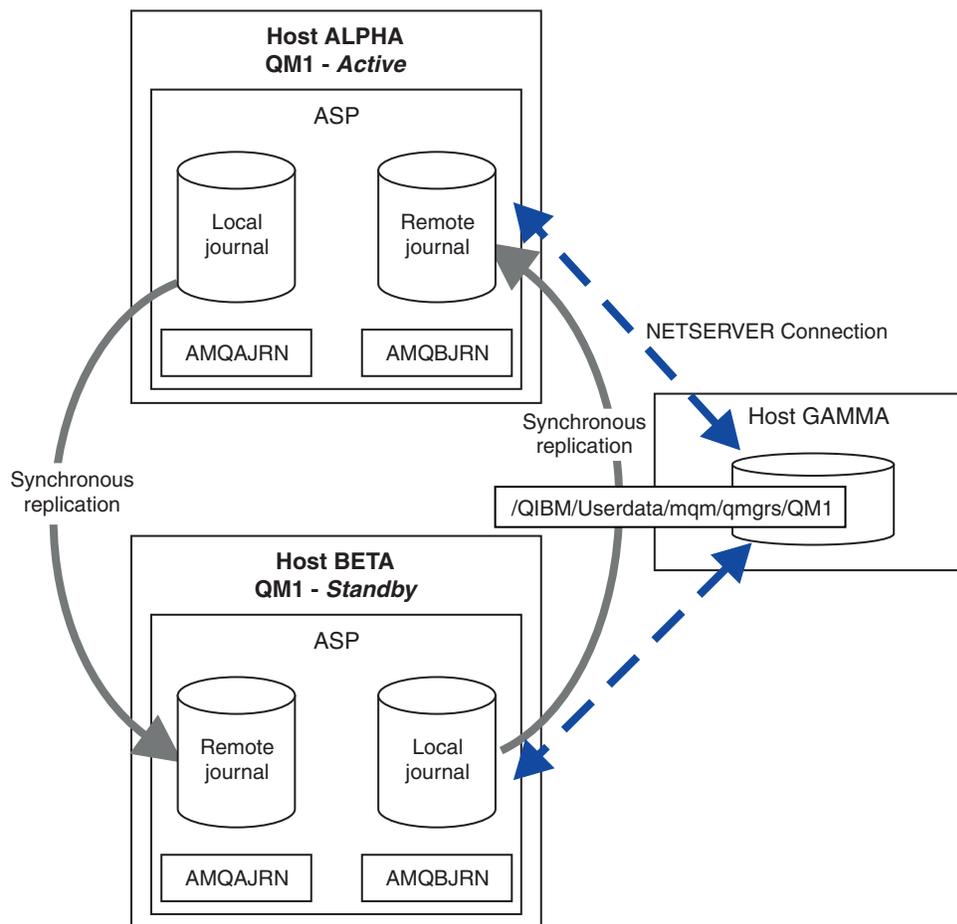


Figure 47. Mirrored journal configuration

What to do next

- Verify that the active and standby instances switch over automatically. You can run the sample high availability sample programs to test the switch over; see  High availability sample programs (*WebSphere MQ V7.1 Programming Guide*). The sample programs are 'C' clients. You can run them from a Windows or Unix platform.
 1. Start the high availability sample programs.
 2. On ALPHA, end the queue manager requesting switch over:
`ENDMQM MQMNAME(QM1) OPTION(*IMMED) ALSWITCH(*YES)`
 3. Check that the instance of QM1 on BETA is active.
 4. Restart QM1 on ALPHA
`STRMQM MQMNAME(QM1) STANDBY(*YES)`
- Look at alternative high availability configurations:
 1. Use NetServer to place the queue manager data on a Windows server.
 2. Instead of using remote journaling to mirror the queue manager journal, store the journal on an independent ASP. Use IBM i clustering to transfer the independent ASP from ALPHA to BETA.

Converting a single instance queue manager to a multi-instance queue manager using NetServer and journal mirroring:

Convert a single instance queue manager to a multi-instance queue manager. Move the queue manager data to a network share connected by NetServer. Mirror the queue manager journal to a second IBM i server using remote journaling.

Before you begin

1. The task requires three IBM i servers. The existing WebSphere MQ installation, on the server ALPHA in the example, must be at least at version 7.0.1.1. ALPHA is running a queue manager called QM1 in the example.
2. Install WebSphere MQ on the second IBM i server, BETA in the example.
3. The third server is an IBM i server, connected by NetServer to ALPHA and BETA. It is used to share the queue manager data. It does not have to have a WebSphere MQ installation. It is useful to install WebSphere MQ on the server as a temporary step, to set up the queue manager directories and permissions.
4. Make sure that the QMQM user profile has the same password on all three servers.
5. Install IBM i NetServer; see  IBM i NetServer.

About this task

Perform the following steps to convert a single instance queue manager to the multi-instance queue manager shown in Figure 48 on page 223. The single instance queue manager is deleted in the task, and then recreated, storing the queue manager data on the network share connected by NetServer. This procedure is more reliable than moving the queue manager directories and files to the network share using the **CPY** command.

- Create connections from ALPHA and BETA to the directory share on GAMMA that is to store the queue manager data. The task also sets up the necessary permissions, user profiles and passwords.
- Add Relational Database Entries (RDBE) to the IBM i systems that are going to run queue manager instances. The RDBE entries are used to connect to the IBM i systems used for remote journaling.
- Save the queue manager logs and definitions, stop the queue manager, and delete it.
- Recreate the queue manager, storing the queue manager data on the network share on GAMMA.
- Add the second instance of the queue manager to the other server.
- Create remote journals on both the IBM i servers for both queue manager instances. Each queue manager writes to the local journal. The local journal is replicated to the remote journal. The command, **ADDQMJRN** simplifies adding the journals and the connections.
- Start the queue manager, permitting a standby instance.

Note:

In step 5 on page 221 of the task, you delete the single instance queue manager, QM1. Deleting the queue manager deletes all the persistent messages on queues. For this reason, complete processing all the messages stored by the queue manager, before converting the queue manager. If processing all the messages is not possible, back up the queue manager library before step 5 on page 221. Restore the queue manager library after step 6 on page 221.

Note:

In step 6 on page 221 of the task, you recreate QM1. Although the queue manager has the same name, it has a different queue manager identifier. Queue manager clustering uses the queue manager identifier. To delete and recreate a queue manager in a cluster, you must first remove the queue manager from the cluster; see  Removing a queue manager from a cluster: Alternative method (*WebSphere MQ V7.1*

Installing Guide) or  Removing a queue manager from a cluster (*WebSphere MQ V7.1 Installing Guide*). When you have recreated the queue manager, add it to the cluster. Although it has the same name as before, it appears to be a new queue manager to the other queue managers in the cluster.

Procedure

1. Do the task, “Creating a network share for queue manager data using NetServer” on page 208.
As a result, ALPHA and BETA have a share, /QNTC/GAMMA/WMQ, that points to /QIBM/UserData/mqm/qmgrs on GAMMA. The user profiles QMQM and QMQMADM have the necessary permissions, and QMQM has matching passwords on all three systems.
2. Add Relational Database Entries (RDBE) to the IBM i systems that are going to host queue manager instances.
 - a. On ALPHA create the connection to BETA.
ADDRDBDIRE RDB(BETA) RMTLOCNAME(BETA *IP) RMTAUTMTH(*USRIDPWD)
 - b. On BETA create the connections to ALPHA.
ADDRDBDIRE RDB(ALPHA) RMTLOCNAME(ALPHA *IP) RMTAUTMTH(*USRIDPWD)
3. Download and install the SupportPac,  MS03, WebSphere MQ - Save Queue Manager object definitions using PCFs (**SAVEQMGR**).
The SupportPac installs the command **QSAVEQMGR**.
4. Create the scripts that recreate the queue manager objects.
QSAVEQMGR LCLQMGRNAM(QM1) FILENAME(' *CURLIB/QMQSC(QM1) ')
OUTPUT(*REPLACE) MAKEAUTH(*YES) AUTHFN(' *CURLIB/QMAUT(QM1) ')
5. Stop the queue manager and delete it.
ENDMQM MQMNAME(QM1) OPTION(*IMMED) ENDCCTJOB(*YES) RCDMQMIMG(*YES) TIMEOUT(15)
DLTMQM MQMNAME(QM1)
6. Create the queue manager QM1 on ALPHA, saving the queue manager data on GAMMA.
CRTMQM MQMNAME(QM1) UDLMSGQ(SYSTEM.DEAD.LETTER.QUEUE)
MQMDIRP('/QNTC/GAMMA/WMQ')
The path, /QNTC/GAMMA/WMQ, uses NetServer to create the queue manager data in /QIBM/UserData/mqm/qmgrs.
7. Recreate the queue manager objects for QM1 from the saved definitions.
STRMQMMQSC SRCMBR(QM1) SRCFILE(*CURLIB/QMQSC) MQMNAME(QM1)
8. Apply the authorizations from the saved information.
 - a. Compile the saved authorization program.
CRTCLPGM PGM(*CURLIB/QM1) SRCFILE(*CURLIB/QMAUT)
SRCMBR(QM1) REPLACE(*YES)
 - b. Run the program to apply the authorizations.
CALL PGM(*CURLIB/QM1)
 - c. Refresh the security information for QM1.
RFRMQMAUT MQMNAME(QM1)
9. Run **ADDQMJRN** on ALPHA. The command adds a remote journal on BETA for QM1.
ADDQMJRN MQMNAME(QM1) RMTJRNRDB(BETA)
QM1 creates journal entries in its local journal on ALPHA when the active instance of QM1 is on ALPHA. The local journal on ALPHA is replicated to the remote journal on BETA.
10. Use the command, **DSPF**, to inspect the WebSphere MQ configuration data created by **CRTMQM** for QM1 on ALPHA.
The information is needed in the next step.
In this example, the following configuration is created in /QIBM/UserData/mqm/mqs.ini on ALPHA for QM1:

```
Name=QM1
Prefix=/QIBM/UserData/mqm
Library=QMQM1
Directory=QM1
DataPath=/QNTC/GAMMA/WMQ/QM1
```

11. Create a queue manager instance of QM1 on BETA using the **ADDQMINF** command. Run the following command on BETA to modify the queue manager control information in /QIBM/UserData/mqm/mqs.ini on BETA.

```
ADDQMINF MQMNAME(QM1)
PREFIX('/QIBM/UserData/mqm')
MQMDIR(QM1)
MQMLIB(QMQM1)
DATAPATH('/QNTC/GAMMA/WMQ/QM1')
```

Tip: Copy and paste the configuration information. The queue manager stanza is the same on ALPHA and BETA.

12. Run **ADDQMJRN** on BETA. The command adds a local journal on BETA and a remote journal on ALPHA for QM1.

```
ADDQMJRN MQMNAME(QM1) RMTJRNDRDB(ALPHA)
```

QM1 creates journal entries in its local journal on BETA when the active instance of QM1 is on BETA. The local journal on BETA is replicated to the remote journal on ALPHA.

Note: As an alternative, you might want to set up remote journaling from BETA to ALPHA using asynchronous journaling.

Use this command to set up asynchronous journaling from BETA to ALPHA, instead of the command in step 7 on page 217.

```
ADDQMJRN MQMNAME(QM1) RMTJRNDRDB(ALPHA) RMTJRNDRDLV(*ASYNCR)
```

If the server or journaling on ALPHA is the source of the failure, BETA starts without waiting for new journal entries to be replicated to ALPHA.

Switch the replication mode to *SYNCR, using the **CHGMQMJRN** command, when ALPHA is online again.

Use the information in “Mirrored journal configuration for ASP” on page 211 to decide whether to mirror the journals synchronously, asynchronously, or a mixture of both. The default is to replicate synchronously, with a 60 second wait period for a response from the remote journal.

13. Verify that the journals on ALPHA and BETA are enabled and the status of remote journal replication is *ACTIVE.
 - a. On ALPHA:

```
WRKMQMJRN MQMNAME(QM1)
```
 - b. On BETA:

```
WRKMQMJRN MQMNAME(QM1)
```
14. Start the queue manager instances on ALPHA and BETA.
 - a. Start the first instance on ALPHA, making it the active instance. Enabling switching over to a standby instance.

```
STRMQM MQMNAME(QM1) STANDBY(*YES)
```
 - b. Start the second instance on BETA, making it the standby instance.

```
STRMQM MQMNAME(QM1) STANDBY(*YES)
```

Results

Use **WRKMQM** to check queue manager status:

1. The status of the queue manager instance on ALPHA should be *ACTIVE.
2. The status of the queue manager instance on BETA should be *STANDBY.

Example

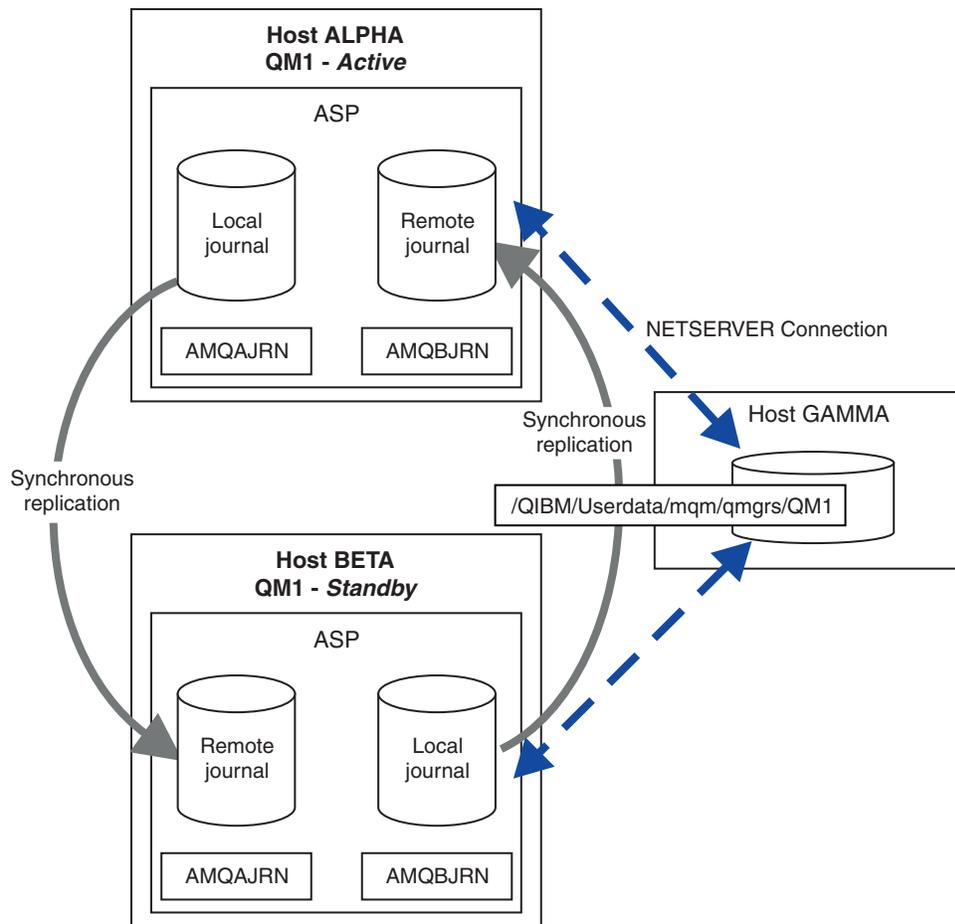


Figure 48. Mirrored journal configuration

What to do next

- Verify that the active and standby instances switch over automatically. You can run the sample high availability sample programs to test the switch over; see  High availability sample programs (*WebSphere MQ V7.1 Programming Guide*). The sample programs are 'C' clients. You can run them from a Windows or Unix platform.
 1. Start the high availability sample programs.
 2. On ALPHA, end the queue manager requesting switch over:
`ENDMQM MQMNAME(QM1) OPTION(*IMMED) ALSWITCH(*YES)`
 3. Check that the instance of QM1 on BETA is active.
 4. Restart QM1 on ALPHA
`STRMQM MQMNAME(QM1) STANDBY(*YES)`

- Look at alternative high availability configurations:
 1. Use NetServer to place the queue manager data on a Windows server.
 2. Instead of using remote journaling to mirror the queue manager journal, store the journal on an independent ASP. Use IBM i clustering to transfer the independent ASP from ALPHA to BETA.

Switched independent ASP journal configuration:

You do not need to replicate an independent ASP journal to create a multi-instance queue manager configuration. You do need to automate a means to transfer the independent ASP from the active queue manager to the standby queue manager. There are alternative high availability solutions possible using an independent ASP, not all of which require using a multi-instance queue manager.

When using an independent ASP you do not need to mirror the queue manager journal. If you have installed cluster management, and the servers hosting the queue manager instances are in the same cluster resource group, then the queue manager journal can be transferred automatically to another server within a short distance of the active server, if the host running the active instance fails. You can also transfer the journal manually, as part of a planned switch, or you can write a command procedure to transfer the independent ASP programmatically.

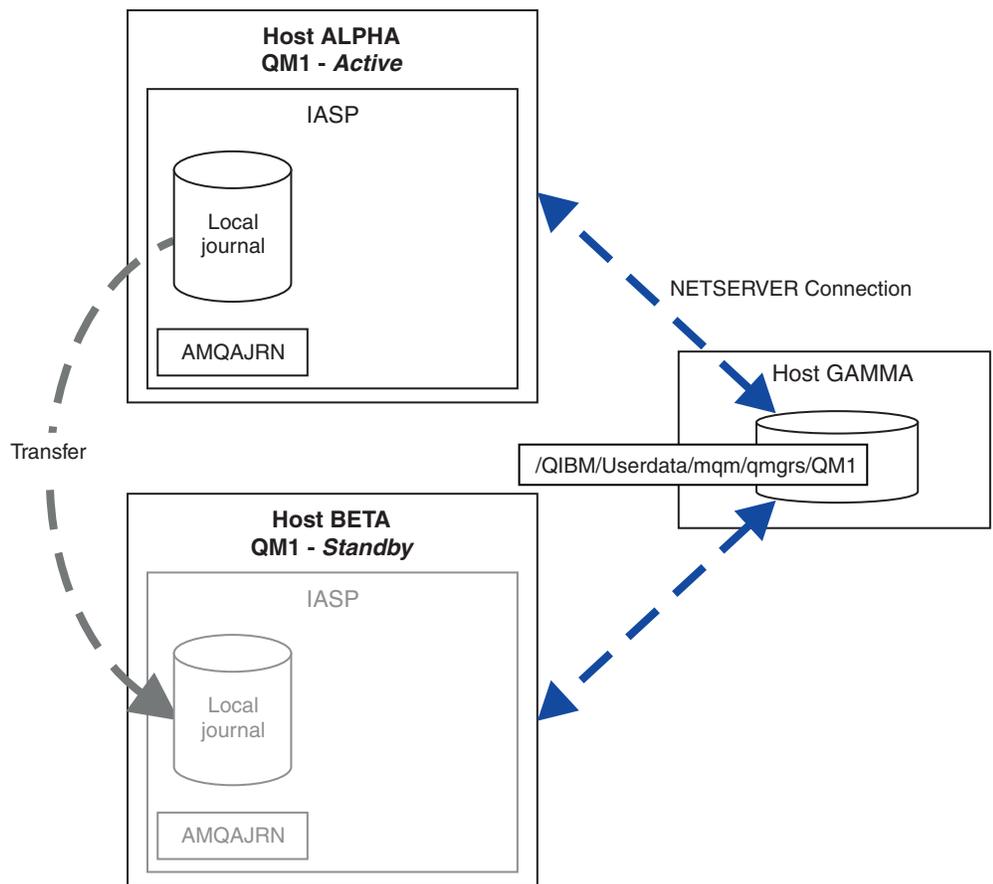


Figure 49. Transfer a queue manager journal using an independent ASP

For multi-instance queue manager operation, queue manager data must be stored on a shared file system. The file system can be hosted on a variety of different platforms. You cannot store multi-instance queue manager data on an ASP or independent ASP.

The shared file system performs two roles in the configuration: The same queue manager data is shared between all instances of the queue manager. The file system must have a robust locking protocol that ensures only one instance of the queue manager has access to queue manager data once it has started. If the queue manager fails, or the communications to the file server breaks, then the file system must release the lock to the queue manager data held by the active instance that is no longer communicating with the file system. The standby queue manager instance can then gain read/write access to the queue manager data. The file system protocol must conform to a set of rules to work correctly with multi-instance queue managers; see “Components of a high availability solution” on page 206.

The locking mechanism serializes the start queue manager command and controls which instance of the queue manager is active. Once a queue manager becomes active, it rebuilds its queues from the local journal that you, or the HA cluster, has transferred to the standby server. Reconnectable clients that are waiting for reconnection to the same queue manager get reconnected, and any inflight transactions are backed out. Applications that are configured to start as queue manager services are started.

You need to ensure that the local journal from the failed active queue manager instance on the independent ASP is transferred to the server that hosts the newly activated standby queue manager instance, either by configuring the cluster resource manager, or transferring the independent ASP manually. Using independent ASPs does not preclude configuring remote journals and mirroring, if you decide to use independent ASP for backup and disaster recovery, and use remote journal mirroring for multi-instance queue manager configuration.

If you have chosen to use an independent ASP, there are alternative highly available configurations you might consider. The background to these solutions are described in “Independent ASPs and high availability” on page 228.

1. Rather than use multi-instance queue managers, install and configure a single instance queue manager entirely on an independent ASP, and use IBM i high availability services to fail the queue manager over. You would probably need to augment the solution with a queue manager monitor to detect whether the queue manager has failed independently of the server. This is the basis of the solution provided in, *Supportpac MC41: Configuring WebSphere MQ for iSeries for High Availability*.
2. Use independent ASPs and cross site mirroring (XSM) to mirror the independent ASP rather than switching the independent ASP on the local bus. This extends the geographic range of the independent ASP solution to as far as the time taken to write log records over a long distance allows.

Creating a multi-instance queue manager using an independent ASP and NetServer:

Create a multi-instance queue manager to run on two IBM i servers. The queue manager data is stored on an IBM i server using NetServer. The queue manager journal is stored on an independent ASP. Use IBM i clustering or a manual procedure to transfer the independent ASP containing the queue manager journal to the other IBM i server.

Before you begin

1. The task requires three IBM i servers. Install WebSphere MQ on two of them, ALPHA and BETA in the example. WebSphere MQ must be at least at version 7.0.1.1.
2. The third server is an IBM i server, connected by NetServer to ALPHA and BETA. It is used to share the queue manager data. It does not have to have a WebSphere MQ installation. It is useful to install WebSphere MQ on the server as a temporary step, to set up the queue manager directories and permissions.
3. Make sure that the QMQM user profile has the same password on all three servers.
4. Install IBM i NetServer; see  IBM i NetServer.
5. Create procedures to transfer the independent ASP from the failed queue manager to the standby that is taking over. You might find some of the techniques in *SupportPac MC41: Configuring WebSphere MQ for iSeries for High Availability* helpful in designing your independent ASP transfer procedures.

About this task

Perform the following steps to create the configuration shown in Figure 50 on page 227. The queue manager data is connected using IBM i NetServer.

- Create connections from ALPHA and BETA to the directory share on GAMMA that is to store the queue manager data. The task also sets up the necessary permissions, user profiles and passwords.
- Create the queue manager QM1 on the IBM i server, ALPHA.
- Add the queue manager control information for QM1 on the other IBM i server, BETA.
- Start the queue manager, permitting a standby instance.

Procedure

1. Do the task, “Creating a network share for queue manager data using NetServer” on page 208.

As a result, ALPHA and BETA have a share, /QNTC/GAMMA/WMQ, that points to /QIBM/UserData/mqm/qmgrs on GAMMA. The user profiles QMQM and QMQMADM have the necessary permissions, and QMQM has matching passwords on all three systems.

2. Create the queue manager QM1 on ALPHA, saving the queue manager data on GAMMA.

```
CRTMQM MQMNAME(QM1) UDLMSGQ(SYSTEM.DEAD.LETTER.QUEUE)
      MQMDIRP('/QNTC/GAMMA/WMQ')
```

The path, /QNTC/GAMMA/WMQ, uses NetServer to create the queue manager data in /QIBM/UserData/mqm/qmgrs.

3. Use the command, **DSPF**, to inspect the WebSphere MQ configuration data created by **CRTMQM** for QM1 on ALPHA.

The information is needed in the next step.

In this example, the following configuration is created in /QIBM/UserData/mqm/mqs.ini on ALPHA for QM1:

```
Name=QM1
Prefix=/QIBM/UserData/mqm
Library=QMQM1
Directory=QM1
DataPath=/QNTC/GAMMA/WMQ/QM1
```

4. Create a queue manager instance of QM1 on BETA using the **ADDQMINF** command. Run the following command on BETA to modify the queue manager control information in /QIBM/UserData/mqm/mqs.ini on BETA.

```
ADDQMINF MQMNAME(QM1)
PREFIX('/QIBM/UserData/mqm')
MQMDIR(QM1)
MQMLIB(QMQM1)
DATAPATH('/QNTC/GAMMA/WMQ/QM1')
```

Tip: Copy and paste the configuration information. The queue manager stanza is the same on ALPHA and BETA.

5. Start the queue manager instances on ALPHA and BETA.
 - a. Start the first instance on ALPHA, making it the active instance. Enabling switching over to a standby instance.

```
STRMQM MQMNAME(QM1) STANDBY(*YES)
```
 - b. Start the second instance on BETA, making it the standby instance.

```
STRMQM MQMNAME(QM1) STANDBY(*YES)
```

Results

Use **WRKMQM** to check queue manager status:

1. The status of the queue manager instance on ALPHA should be *ACTIVE.
2. The status of the queue manager instance on BETA should be *STANDBY.

Example

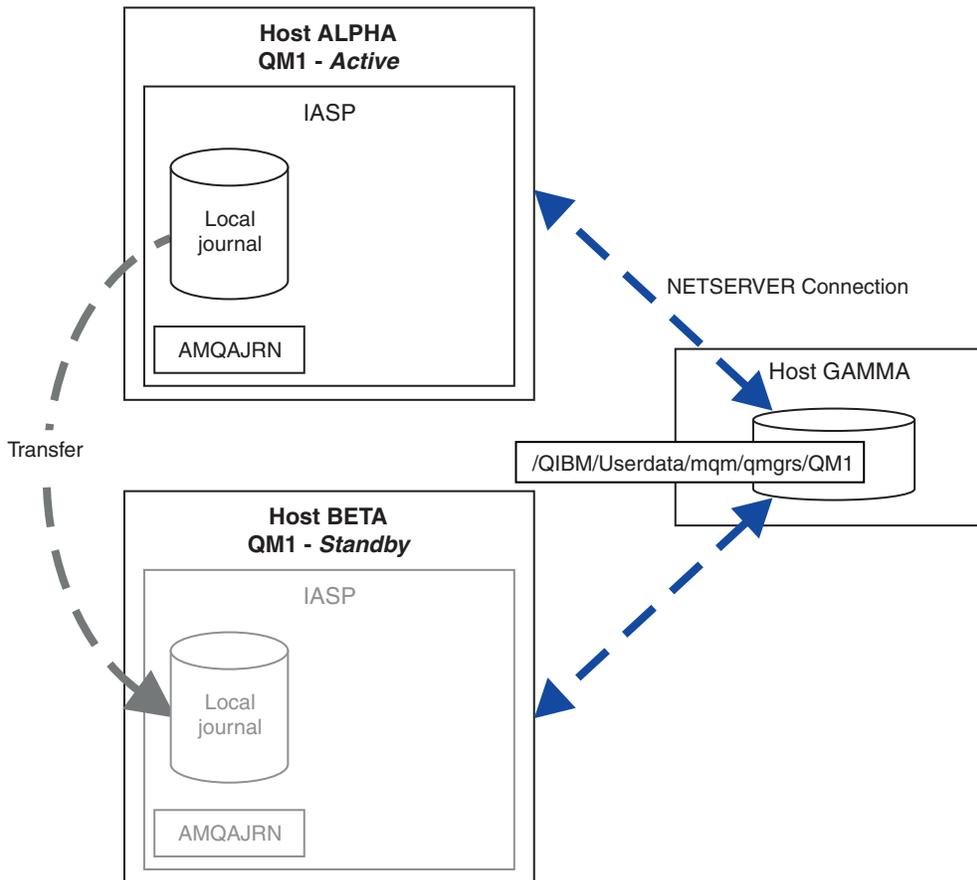


Figure 50. Transfer a queue manager journal using an independent ASP

What to do next

- Verify that the active and standby instances switch over automatically. You can run the sample high availability sample programs to test the switch over; see  High availability sample programs (*WebSphere MQ V7.1 Programming Guide*). The sample programs are 'C' clients. You can run them from a Windows or Unix platform.
 1. Start the high availability sample programs.
 2. On ALPHA, end the queue manager requesting switch over:
`ENDMQM MQMNAME(QM1) OPTION(*IMMED) ALSWITCH(*YES)`
 3. Check that the instance of QM1 on BETA is active.
 4. Restart QM1 on ALPHA
`STRMQM MQMNAME(QM1) STANDBY(*YES)`
- Look at alternative high availability configurations:
 1. Use NetServer to place the queue manager data on an IBM i server.

2. Instead of using an independent ASP to transfer the queue manager journal to the standby server, use remote journaling to mirror the journal onto the standby server.

Independent ASPs and high availability:

Independent ASPs enable applications and data to be moved between servers. The flexibility of independent ASPs means they are the basis for some IBM i high availability solutions. In considering whether to use an ASP or independent ASP for the queue manager journal, you should consider other high availability configuration based on independent ASPs.

Auxiliary storage pools (ASPs) are a building block of IBM i architecture. Disk units are grouped together to form a single ASP. By placing objects in different ASPs you can protect data in one ASP from being affected by disk failures in another ASP.

Every IBM i server has at least one *basic* ASP, known as the system ASP. It is designated as ASP1, and sometimes known as *SYSBAS. You can configure up to 31 additional basic *user* ASPs that are indistinguishable from the system ASP from the application's point of view, because they share the same namespace. By using multiple basic ASPs to distribute applications over many disks you can improve performance and reduce recovery time. Using multiple basic ASPs can also provide some degree of isolation against disk failure, but it does not improve reliability overall.

Independent ASPs are a special type of ASP. They are often called independent disk pools. Independent disk pools are key component of IBM i high availability. You can store data and applications that regard themselves as independent from the current system to which they are connected on independent disk storage units. You can configure switchable or non-switchable independent ASPs. From an availability perspective you are generally only concerned with switchable independent ASPs, which can be transferred automatically from server to server. As a result you can move the applications and data on the independent ASP from server to server.

Unlike basic user ASPs, independent ASPs do not share the same namespace as the system ASP. Applications that work with user ASPs require changes to work with an independent ASP. You need to verify your software, and third-party software you use, works in an independent ASP environment.

When the independent ASP is attached to a different server the namespace of the independent ASP has to be combined with the namespace of the system ASP. This process is called *varying-on* the independent ASP. You can vary-on an independent ASP without IPLing the server. Clustering support is required to transfer independent ASPs automatically from one server to another.

Building reliable solutions with independent ASPs

Journaling to an independent ASP, rather than journaling to an ASP and using journal replication, provides an alternative means to provide the standby queue manager with a copy of the local journal from the failed queue manager instance. To automatically transfer the independent ASP to another server you need to have installed and configured clustering support. There are a number of high-availability solutions for independent ASPs based on the cluster support, and low level disk mirroring, that you can combine with, or substitute for, using multi-instance queue managers.

The following list describes the components that are needed to build a reliable solution based on independent ASPs.

Journaling

Queue managers, and other applications, use local journals to write persistent data safely to disk to protect against loss of data in memory due to server failure. This is sometimes termed point-in-time consistency. It does not guarantee the consistency of multiple updates that take place over a period of time.

Commitment control

By using global transactions, you can coordinate updates to messages and databases so that the data written to the journal is consistent. It gives consistency over a period of time by using a two-phase commit protocol.

Switched disk

Switched disks are managed by the device cluster resource group (CRG) in an HA cluster. CRG switches independent ASPs automatically to a new server in the case of an unplanned outage. CRGs are geographically limited to the extent of the local IO bus.

By configuring your local journal on a switchable independent ASP, you can transfer the journal to a different server, and resume processing messages. No changes to persistent messages made without syncpoint control, or committed with syncpoint control, are lost, unless the independent ASP fails.

If you use both journaling and commitment control on switchable independent ASPs, you can transfer database journals and queue manager journals to a different server and resume processing transactions with no loss of consistency or committed transactions.

Cross-site mirroring (XSM)

XSM mirrors the primary independent ASP to a geographically remote secondary independent ASP across a TCP/IP network, and transfers control automatically in case of a failure. You have a choice of configuring a synchronous or asynchronous mirror. Synchronous mirroring reduces the performance of the queue manager because data is mirrored before the write operations on the production system complete, but it does guarantee the secondary independent ASP is up to date. Whereas if you use asynchronous mirroring you cannot guarantee that the secondary independent ASP is up to date. Asynchronous mirroring does maintain the consistency of the secondary independent ASP.

There are three XSM technologies.

Geographic mirroring

Geographic mirroring is an extension of clustering, enabling you to switch independent ASPs across a wide area. It has both synchronous and asynchronous modes. You can guarantee high availability only in synchronous mode, but the separation of independent ASPs might impact performance too much. You can combine geographic mirroring with switched disk to provide high availability locally and disaster recovery remotely.

Metro mirroring

Metro mirroring is a device level service that provides fast local synchronous mirroring over longer distances than the local bus. You can combine it with a multi-instance queue manager to give you high availability of the queue manager, and by having two copies of the independent ASP, high availability of the queue manager journal.

Global mirroring

Global mirroring is device level service that provides asynchronous mirroring, and is suitable for backing up and disaster recovery over longer distances, but is not a normal choice for high availability, because it only maintains point in time consistency rather than currency.

The key decision points you should consider are,

ASP or independent ASP?

You do not need to run a IBM i HA cluster to use multi-instance queue managers. You might choose independent ASPs, if you are already using independent ASPs, or you have availability requirements for other applications that require independent ASPs. It might be worth combining independent ASPs with multi-instance queue managers to replace queue manager monitoring as a means of detecting queue manager failure.

Availability?

What is the recovery time objective (RTO)? If you require the appearance of near uninterrupted behavior, then which solution has the quickest recovery time?

Journal availability?

How do you eliminate the journal as a single point of failure. You might adopt a hardware solution, using RAID 1 devices or better, or you might combine or use a software solution using replica journals or disk mirroring.

Distance?

How far apart are the active and standby queue manager instances. Can your users tolerate the performance degradation of replicating synchronously over distances greater than about 250 meters?

Skills?

There is work to be done to automate the administrative tasks involved in maintaining and exercising the solution regularly. The skills required to do the automation are different for the solutions based on ASPs and independent ASPs.

Deleting a multi-instance queue manager:

Before you delete a multi-instance queue manager, stop remote journaling, and remove queue manager instances.

Before you begin

1. In this example, two instances of the QM1 queue manager are defined on the servers ALPHA and BETA. ALPHA is the active instance and BETA is the standby. The queue manager data associated with the queue manager QM1 is stored on the IBM i server GAMMA, using NetServer. See “Creating a multi-instance queue manager using journal mirroring and NetServer” on page 216.
2. ALPHA and BETA must be connected so that any remote journals that are defined can be deleted by WebSphere MQ.
3. Verify that the /QNTC directory and server directory file share can be accessed, using the system commands **EDTF** or **WRKLNK**

About this task

Before you delete a multi-instance queue manager from a server using the **DLTMQM** command, remove any queue manager instances on other servers using the **RMVMQMINF** command.

When you remove a queue manager instance using the **RMVMQMINF** command, local and remote journals prefixed with AMQ, and associated with the instance, are deleted. Configuration information about the queue manager instance, local to the server, is also deleted.

Do not run the **RMVMQMINF** command on the server holding the remaining instance of the queue manager. Doing so prevents **DLTMQM** from working correctly.

Delete the queue manager using the **DLTMQM** command. Queue manager data is removed from the network share. Local and remote journals prefixed with AMQ and associated with the instance are deleted. **DLTMQM** also deletes configuration information about the queue manager instance, local to the server.

In the example, there are only two queue manager instances. WebSphere MQ supports a running multi-instance configuration that has one active queue manager instance and one standby instance. If you have created additional queue manager instances to use in running configurations, remove them, using the **RMVMQMINF** command, before deleting the remaining instance.

Procedure

1. Run the **CHGMQMJRN RMTJRNSTS(*INACTIVE)** command on each server to make remote journaling between the queue manager instances inactive.
 - a. On ALPHA:

```
CHGMQMJRN MQMNAME('QM1')
RMTJRN RDB('BETA') RMTJRNSTS(*INACTIVE)
```
 - b. On BETA:

```
CHGMQMJRN MQMNAME('QM1')
RMTJRN RDB('ALPHA') RMTJRNSTS(*INACTIVE)
```
2. Run the **ENDMQM** command on ALPHA, the active queue manager instance, to stop both instances of QM1.

```
ENDMQM MQMNAME(QM1) OPTION(*IMMED) INSTANCE(*ALL) ENDCCTJOB(*YES)
```
3. Run the **RMVMQMINF** command on ALPHA to remove the queue manager resources for the instance from ALPHA and BETA.

```
RMVMQMINF MQMNAME(QM1)
```

RMVMQMINF removes the queue manager configuration information for QM1 from ALPHA. If the journal name is prefixed by AMQ, it deletes the local journal associated with QM1 from ALPHA. If the journal name is prefixed by AMQ and a remote journal has been created, it also removes the remote journal from BETA.
4. Run the **DLTMQM** command on BETA to delete QM1.

```
DLTMQM MQMNAME(QM1)
```

DLTMQM deletes the queue manager data from the network share on GAMMA. It removes the queue manager configuration information for QM1 from BETA. If the journal name is prefixed by AMQ, it deletes the local journal associated with QM1 from BETA. If the journal name is prefixed by AMQ and a remote journal has been created, it also removes the remote journal from ALPHA.

Results

DLTMQM and **RMVMQMINF** delete the local and remote journals created by **CRTMQM** and **ADDMQJRN**. The commands also delete the journal receivers. The journals and journal receivers must follow the naming convention of having names starting with AMQ. **DLTMQM** and **RMVMQMINF** remove the queue manager objects, queue manager data, and the queue manager configuration information from `mq5.ini`.

What to do next

An alternative approach is to issue the following commands after deactivating journaling in step 1 and before ending the queue manager instances. Or, if you have not followed the naming convention, you must delete the journals and journal receivers by name.

1. On ALPHA:

```
RMVMQMJRN MQMNAME('QM1') RMTJRN RDB('BETA')
```
2. On BETA:

```
RMVMQMJRN MQMNAME('QM1') RMTJRN RDB('ALPHA')
```

After deleting the journals, continue with the rest of the steps.

Backing up a multi-instance queue manager:

The procedure shows you how to back up queue manager objects on the local server and the queue manager data on the network file server. Adapt the example to back up data for other queue managers.

Before you begin

In this example, the queue manager data associated with the queue manager QM1 is stored on the IBM i server called GAMMA, using NetServer. See “Creating a multi-instance queue manager using journal mirroring and NetServer” on page 216. WebSphere MQ is installed on the servers, ALPHA and BETA. The queue manager, QM1, is configured on ALPHA and BETA.

About this task

IBM i does not support saving data from a remote directory. Save the queue manager data on a remote file system using the backup procedures local to the file system server. In this task, the network file system is on an IBM i server, GAMMA. The queue manager data is backed up in a save file on GAMMA.

If the network file system was on Windows or Linux, you might store the queue manager data in a compressed file, and then save it. If you have a back-up system, such as Tivoli® Storage Manager, use it to back up the queue manager data.

Procedure

1. Create a save file on ALPHA for the queue manager library associated with QM1.
Use the queue manager library name to name the save file.
`CRTSAVF FILE(QGPL/QMQM1)`
2. Save the queue manager library in the save file on ALPHA.
`SAVLIB LIB(QMQM1) DEV(*SAVF) SAVF(QGPL/QMQM1)`
3. Create a save file for the queue manager data directory on GAMMA.
Use the queue manager name to name the save file.
`CRTSAVF FILE(QGPL/QMDQM1)`
4. Save the copy of the queue manager data from the local directory on GAMMA.
`SAV DEV('/QSYS.LIB/QGPL.LIB/QMDQM1.FILE') OBJ('/QIBM/Userdata/mqm/qmgrs/QM1')`

Commands to set up multi-instance queue managers:

WebSphere MQ has commands to simplify configuring journal replication, adding new queue manager instances, and configuring queue managers to use independent ASP.

The journal commands to create and manage local and remote journals are,

ADDMQMJRN

With this command you can create named local and remote journals for a queue manager instance, and configure whether replication is synchronous or asynchronous, what the synchronous timeout is, and if the remote journal is to be activated immediately.

CHGMQMJRN

The command modifies the timeout, status and delivery parameters affecting replica journals.

RMVMQMJRN

Removes named *remote* journals from a queue manager instance.

WRKMQMJRN

Lists the status of local and remote journals for a local queue manager instance.

Add and manage additional queue manager instances using the following commands, which modify the `mqs.ini` file.

ADDMQMINF

The command uses information you extracted from the `mqs.ini` file with `DSPMQMINF` command to add a new queue manager instance on a different IBM i server.

RMVMQMINF

Remove a queue manager instance. Use this command either to remove an instance of an existing queue manager, or to remove the configuration information for a queue manager that has been deleted from a different server.

The **CRTMQM** command has three parameters to assist configuring a multi-instance queue manager,

MQMDIRP(*DFT | *directory-prefix*)

Use this parameter to select a mount point that is mapped to queue manager data on networked storage.

ASP(*SYSTEM | *ASPDEV | *auxiliary-storage-pool-number*)

Specify `*SYSTEM`, or an *auxiliary-storage-pool-number* to place the queue manager journal on the system or a basic user ASP. Select the `*ASPDEV` option, and also set a device name using the `ASPDEV` parameter, to place the queue manager journal on an independent ASP.

ASPDEV(*ASP | *device-name*)

Specify a *device-name* of a primary or secondary independent ASP device. Selecting `*ASP` has the same result as specifying `ASP(*SYSTEM)`.

Performance and disk failover considerations

Use different auxiliary storage pools to improve performance and reliability.

If you use a large number of persistent messages or large messages in your applications, the time spent writing these message to disk becomes a significant factor in the performance of the system.

Ensure that you have sufficient disk activation to cope with this possibility, or consider a separate Auxiliary Storage Pool (ASP) in which to hold your queue manager journal receivers.

You can specify which ASP your queue manager library and journals are stored on when you create your queue manager using the `ASP` parameter of **CRTMQM**. By default, the queue manager library and journals and IFS data are stored in the system ASP.

ASPs allow isolation of objects on one or more specific disk units. This can also reduce the loss of data because of a disk media failure. In most cases, only the data that is stored on disk units in the affected ASP is lost.

You are recommended to store the queue manager library and journal data in separate user ASPs to that of the root IFS file system to provide failover and reduce disk contention.

For more information, see *i5/OS V4R4M0 Backup and Recovery*

Using SAVLIB to save WebSphere MQ libraries

You cannot use `SAVLIB LIB(*ALLUSR)` to save the WebSphere MQ libraries, because these libraries have names beginning with `Q`.

You can use `SAVLIB LIB(QM*)` to save all the queue manager libraries, but only if you are using a save device other than `*SAVF`. For `DEV(*SAVF)`, you must use a `SAVLIB` command for each and every queue manager library on your system.

Quiescing WebSphere MQ for IBM i

This section explains how to quiesce (end gracefully) WebSphere MQ for IBM i

To quiesce WebSphere MQ for IBM i:

1. Sign on to a new interactive WebSphere MQ for IBM i session, ensuring that you are not accessing any objects.
2. Ensure that you have:
 - *ALLOBJ authority , or object management authority for the QMQM library
 - Sufficient authority to use the ENDSBS command
3. Advise all users that you are going to stop WebSphere MQ for IBM i .
4. How you then proceed depends on whether you want to shut down (quiesce) a single queue manager (where others might exist) (see “Shutting down a single queue manager for WebSphere MQ for IBM i”) or all the queue managers (see “Shutting down all queue managers for WebSphere MQ for IBM i” on page 236).

ENDMQM parameter ENDCCTJOB(*YES)

The ENDMQM parameter ENDCCTJOB(*YES) works differently in WebSphere MQ for IBM i V6.0 and later compared to previous versions.

On previous versions, when you specify ENDCCTJOB(*YES), MQ forcibly terminates your applications for you.

On WebSphere MQ for IBM i V6.0 or later, when you specify ENDCCTJOB(*YES), your applications are not terminated but are instead disconnected from the queue manager.

If you specify ENDCCTJOB(*YES) and you have applications that are not written to detect that a queue manager is ending, the next time a new MQI call is issued, the call will return with a MQRC_CONNECTION_BROKEN (2009) error.

As an alternative to using ENDCCTJOB(*YES), use the parameter ENDCCTJOB(*NO) and use WRKMQM option 22 (Work with jobs) to manually end any application jobs that will prevent a queue manager restart.

Shutting down a single queue manager for WebSphere MQ for IBM i

Use this information to understand the three types of shutdown.

In the procedures that follow, we use a sample queue manager name of QMgr1 and a sample subsystem name of SUBX. Replace these names with your own values if necessary.

Planned shutdown

Planned shutdown of a queue manager on IBM i

1. Before shutdown, execute:
`RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(QMgr1) DSPJRNDTA(*YES)`
2. To shut down the queue manager, execute:
`ENDMQM MQMNAME(QMgr1) OPTION(*CNTRLD)`

If QMgr1 does not end, the channel or applications are probably busy.

3. If you must shut down QMgr1 immediately, execute the following:
`ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)`
`ENDCCTJOB(*YES) TIMEOUT(15)`

Unplanned shutdown

1. To shut down the queue manager, execute:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)
```

If QMgr1 does not end, the channel or applications are probably busy.

2. If you need to shut down QMgr1 immediately, execute the following:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(15)
```

Shutdown under abnormal conditions

1. To shut down the queue manager, execute:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)
```

If QMgr1 does not end, continue with step 3 providing that:

- QMgr1 is in its own subsystem, or
 - You can end all queue managers that share the same subsystem as QMgr1. Use the unplanned shutdown procedure for all such queue managers.
2. When you have taken all the steps in the procedure for all the queue managers sharing the subsystem (SUBX in our examples), execute:

```
ENDSBS SUBX *IMMED
```

If this command fails to complete, shut down all queue managers, using the unplanned shutdown procedure, and perform an IPL on your machine.

Warning: Do not use ENDJOBABN for WebSphere MQ jobs that fail to end as result of ENDJOB or ENDSBS, unless you are prepared to perform an IPL on your machine immediately after.

3. Start the subsystem by executing:

```
STRSBS SUBX
```

4. Shut down the queue manager immediately, by executing:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(10)
```

5. Restart the queue manager by executing:

```
STRMQM MQMNAME(QMgr1)
```

If this fails, and you:

- Have restarted your machine by performing an IPL, or
- Have only a single queue manager

Tidy up WebSphere MQ shared memory by executing:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(15)
```

before repeating step 5.

If the queue manager restart takes more than a few seconds, WebSphere MQ adds status messages intermittently to the job log detailing the startup progress.

If you still have problems restarting your queue manager, contact IBM support. Any further action you might take could damage the queue manager, leaving WebSphere MQ unable to recover.

Shutting down all queue managers for WebSphere MQ for IBM i

Use this information to understand the three types of shutdown.

The procedures are almost the same as for a single queue manager, but using *ALL instead of the queue manager name where possible, and otherwise using a command repeatedly using each queue manager name in turn. Throughout the procedures, we use a sample queue manager name of QMgr1 and a sample subsystem name of SUBX. Replace these with your own.

Planned shutdown

1. One hour before shutdown, execute:

```
RCDMQMIMG OBJ(*ALL) OBJTYPE(*ALL) MQMNAME(QMgr1) DSPJRNDTA(*YES)
```

Repeat this for every queue manager that you want to shut down.

2. To shut down the queue manager, execute:

```
ENDMQM MQMNAME(QMgr1) OPTION(*CNTRLD)
```

Repeat this for every queue manager that you want to shut down; separate commands can run in parallel.

If any queue manager does not end within a reasonable time (for example 10 minutes), proceed to step 3.

3. To shut down all queue managers immediately, execute the following:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(15)
```

Unplanned shutdown

1. To shut down a queue manager, execute:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)
```

Repeat this for every queue manager that you want to shut down; separate commands can run in parallel.

If queue managers do not end, the channel or applications are probably busy.

2. If you need to shut down the queue managers immediately, execute the following:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(15)
```

Shutdown under abnormal conditions

1. To shut down the queue managers, execute:

```
ENDMQM MQMNAME(QMgr1) OPTION(*IMMED)
```

Repeat this for every queue manager that you want to shut down; separate commands can run in parallel.

2. End the subsystems (SUBX in our examples), by executing:

```
ENDSBS SUBX *IMMED
```

Repeat this for every subsystem that you want to shut down; separate commands can run in parallel.

If this command fails to complete, perform an IPL on your system.

Warning: Do not use ENDJOBABN for MQSeries or WebSphere MQ jobs that fail to end as result of ENDJOB or ENDSBS, unless you are prepared to perform an IPL on your system immediately after.

3. Start the subsystems by executing:

```
STRSBS SUBX
```

Repeat this for every subsystem that you want to start.

4. Shut the queue managers down immediately, by executing:

```
ENDMQM MQMNAME(*ALL) OPTION(*IMMED)  
ENDCCTJOB(*YES) TIMEOUT(15)
```

5. Restart the queue managers by executing:

```
STRMQM MQMNAME(QMgr1)
```

Repeat this for every queue manager that you want to start.

If any queue manager restart takes more than a few seconds WebSphere MQ will show status messages intermittently detailing the startup progress.

If you still have problems restarting any queue manager, contact IBM support. Any further action you might take could damage the queue managers, leaving MQSeries or WebSphere MQ unable to recover.

Administering IBM WebSphere MQ for z/OS

Administering queue managers and associated resources includes the tasks that you perform frequently to activate and manage those resources. Choose the method you prefer to administer your queue managers and associated resources.

IBM WebSphere MQ for z/OS can be controlled and managed by a set of utilities and programs provided with the product. You can use the IBM WebSphere MQ Script (MQSC) commands or Programmable Command Formats (PCFs) to administer IBM WebSphere MQ for z/OS. For information about using commands for IBM WebSphere MQ for z/OS, see “Issuing commands” on page 238.

IBM WebSphere MQ for z/OS also provides a set of utility programs to help you with system administration. For information about the different utility programs and how to use them, see “The WebSphere MQ for z/OS utilities” on page 246.

For details of how to administer IBM WebSphere MQ for z/OS and the different administrative tasks you might have to undertake, see the following links:

- “Operating WebSphere MQ for z/OS” on page 248
- “Writing programs to administer WebSphere MQ” on page 272
- “Managing WebSphere MQ resources on z/OS” on page 285
- “Recovery and restart” on page 321
-  WebSphere MQ and CICS (*WebSphere MQ V7.1 Product Overview Guide*)
- “WebSphere MQ and IMS” on page 353

Related concepts:

 WebSphere MQ for z/OS concepts (*WebSphere MQ V7.1 Product Overview Guide*)

“Administering local WebSphere MQ objects” on page 72

“Administering remote WebSphere MQ objects” on page 110

“Administering IBM WebSphere MQ” on page 1

“Administering IBM i” on page 168

 Planning (*WebSphere MQ V7.1 Installing Guide*)

 Planning your IBM WebSphere MQ environment on z/OS (*WebSphere MQ V7.1 Installing Guide*)

 Configuring (*WebSphere MQ V7.1 Installing Guide*)

 Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*)

 Programmable command formats reference (*WebSphere MQ V7.1 Reference*)

 MQSC reference (*WebSphere MQ V7.1 Reference*)

 Using the WebSphere MQ for z/OS utilities (*WebSphere MQ V7.1 Reference*)

Issuing commands

WebSphere MQ script commands (MQSC) are a set of commands that you can use in a batch or interactive mode to control a queue manager.

WebSphere MQ for z/OS supports MQSC commands, which can be issued from the following sources:

- The z/OS console or equivalent (such as SDSF/TSO).
- The initialization input data sets.
- The supplied batch utility, CSQUTIL, processing a list of commands in a sequential data set.
- A suitably authorized application, by sending a command as a message to the command input queue. The application can be any of the following:
 - A batch region program
 - A CICS application
 - An IMS application
 - A TSO application
 - An application program or utility on another WebSphere MQ system

Table 9 on page 242 summarizes the MQSC commands and the sources from which they can be issued.

Much of the functionality of these commands is available in a convenient way from the WebSphere MQ for z/OS operations and controls panels.

Changes made to the resource definitions of a queue manager using the commands (directly or indirectly) are preserved across restarts of the WebSphere MQ subsystem.

WebSphere MQ for z/OS also supports Programmable Command Format (PCF) commands. These simplify the creation of applications for the administration of WebSphere MQ. MQSC commands are in human-readable text form, whereas PCF enables applications to create requests and read the replies without having to parse text strings. Like MQSC commands, applications issue PCF commands by sending them as messages to the command input queue. For more information about using PCF

commands and for details of the commands, see the  WebSphere MQ Programmable Command Formats and Administration Interface (*WebSphere MQ V7.1 Reference*) documentation.

Private and global definitions

When you define an object on WebSphere MQ for z/OS, you can choose whether you want to share that definition with other queue managers (a *global* definition), or whether the object definition is to be used by one queue manager only (a *private* definition). This is called the object *disposition*.

Global definition

If your queue manager belongs to a queue-sharing group, you can choose to share any object definitions you make with the other members of the group. This means that you have to define an object once only, reducing the total number of definitions required for the whole system.

Global object definitions are held in a *shared repository* (a Db2 shared database), and are available to all the queue managers in the queue-sharing group. These objects have a disposition of GROUP.

Private definition

If you want to create an object definition that is required by one queue manager only, or if your queue manager is not a member of a queue-sharing group, you can create object definitions that are not shared with other members of a queue-sharing group.

Private object definitions are held on page set zero of the defining queue manager. These objects have a disposition of QMGR.

You can create private definitions for all types of WebSphere MQ objects except CF structures (that is, channels, namelists, process definitions, queues, queue managers, storage class definitions, and authentication information objects), and global definitions for all types of objects except queue managers.

WebSphere MQ automatically copies the definition of a group object to page set zero of each queue manager that uses it. You can alter the copy of the definition temporarily if you want, and WebSphere MQ allows you to refresh the page set copies from the repository copy if required.

WebSphere MQ always tries to refresh the page set copies from the repository copy on start up (for channel commands, this is done when the channel initiator restarts), or if the group object is changed.

Note: The copy of the definition is refreshed from the definition of the group, only if the definition of the group has changed after you created the copy of the definition.

This ensures that the page set copies reflect the version on the repository, including any changes that were made when the queue manager was inactive. The copies are refreshed by generating DEFINE REPLACE commands, therefore there are circumstances under which the refresh is not performed, for example:

- If a copy of the queue is open, a refresh that changes the usage of the queue fails.
- If a copy of a queue has messages on it, a refresh that deletes that queue fails.
- If a copy of a queue would require ALTER with FORCE to change it.

In these circumstances, the refresh is not performed on that copy, but is performed on the copies on all other queue managers.

If the queue manager is shut down and then restarted stand-alone, any local copies of objects are deleted, unless for example, the queue has associated messages.

There is a third object disposition that applies to local queues only. This allows you to create shared queues. The definition for a shared queue is held on the shared repository and is available to all the queue managers in the queue-sharing group. In addition, the messages on a shared queue are also available to all the queue managers in the queue sharing group. This is described in  Shared queues and queue-sharing groups (*WebSphere MQ V7.1 Product Overview Guide*). Shared queues have an object disposition of SHARED.

The following table summarizes the effect of the object disposition options for queue managers started stand-alone, and as a member of a queue-sharing group.

Disposition	Stand-alone queue manager	Member of a queue-sharing group
QMGR	Object definition held on page set zero.	Object definition held on page set zero.
GROUP	Not allowed.	Object definition held in the shared repository. Local copy held on page set zero of each queue manager in the group.
SHARED	Not allowed.	Queue definition held in the shared repository. Messages available to any queue manager in the group.

Manipulating global definitions

If you want to change the definition of an object that is held in the shared repository, you need to specify whether you want to change the version on the repository, or the local copy on page set zero. Use the object disposition as part of the command to do this.

Directing commands to different queue managers

You can use the *command scope* to control on which queue manager the command runs.

You can choose to execute a command on the queue manager where it is entered, or on a different queue manager in the queue-sharing group. You can also choose to issue a particular command in parallel on all the queue managers in a queue-sharing group. This is possible for both MQSC commands and PCF commands.

This is determined by the *command scope*. The command scope is used with the object disposition to determine which version of an object you want to work with.

For example, you might want to alter some of the attributes of an object, the definition of which is held in the shared repository.

- You might want to change the version on one queue manager only, and not make any changes to the version on the repository or those in use by other queue managers.
- You might want to change the version in the shared repository for future users, but leave existing copies unchanged.
- You might want to change the version in the shared repository, but also want your changes to be reflected immediately on all the queue managers in the queue-sharing group that hold a copy of the object on their page set zero.

Use the command scope to specify whether the command is executed on this queue manager, another queue manager, or all queue managers. Use the object disposition to specify whether the object you are manipulating is in the shared repository (a group object), or is a local copy on page set zero (a queue manager object).

You do not have to specify the command scope and object disposition to work with a shared queue because every queue manager in the queue-sharing group handles the shared queue as a single queue.

Command summary

Use this topic as a reference of the main MQSC and PCF commands.

Table 8 summarizes the MQSC and PCF commands that are available on WebSphere MQ for z/OS to alter, define, delete and display WebSphere MQ objects.

Table 8. Summary of the main MQSC and PCF commands by object type

MQSC command	ALTER	DEFINE	DISPLAY	DELETE
PCF command	Change	Create/Copy	Inquire	Delete
AUTHINFO	X	X	X	X
CFSTATUS			X	
CFSTRUCT	X	X	X	X
CHANNEL	X	X	X	X
CHSTATUS			X	
NAMELIST	X	X	X	X
PROCESS	X	X	X	X
QALIAS	M	M	M	M
QCLUSTER			M	
QLOCAL	M	M	M	M
QMGR	X		X	
QMODEL	M	M	M	M
QREMOTE	M	M	M	M
QUEUE	P	P	X	P
QSTATUS			X	
STGCLASS	X	X	X	X

Key to table symbols:

- M = MQSC only
- P = PCF only
- X = both

There are many other MQSC and PCF commands which allow you to manage other WebSphere MQ resources, and carry out other actions in addition to those summarized in Table 8.

Table 9 on page 242 shows every MQSC command, and where each command can be issued from:

- CSQINP1 initialization input data set
- CSQINP2 initialization input data set
- z/OS console (or equivalent)
- SYSTEM.COMMAND.INPUT queue and command server (from applications, CSQUTIL, or the CSQINPX initialization input data set)

Table 9. Sources from which to run MQSC commands

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
ALTER AUTHINFO		X	X	X
ALTER BUFFPOOL		X	X	X
ALTER CFSTRUCT		X	X	X
ALTER CHANNEL		X	X	X
ALTER NAMELIST		X	X	X
ALTER PSID			X	X
ALTER PROCESS		X	X	X
ALTER QALIAS		X	X	X
ALTER QLOCAL		X	X	X
ALTER QMGR		X	X	X
ALTER QMODEL		X	X	X
ALTER QREMOTE		X	X	X
ALTER SECURITY	X	X	X	X
ALTER STGCLASS		X	X	X
ALTER SUB		X	X	X
ALTER TOPIC		X	X	X
ALTER TRACE	X	X	X	X
ARCHIVE LOG	X	X	X	X
BACKUP CFSTRUCT			X	X
CLEAR QLOCAL		X	X	X
DEFINE AUTHINFO		X	X	X
DEFINE BUFFPOOL	X			
DEFINE CFSTRUCT		X	X	X
DEFINE CHANNEL		X	X	X
DEFINE LOG			X	X
DEFINE NAMELIST		X	X	X
DEFINE PROCESS		X	X	X
DEFINE PSID	X		X	X
DEFINE QALIAS		X	X	X
DEFINE QLOCAL		X	X	X
DEFINE QMODEL		X	X	X
DEFINE QREMOTE		X	X	X
DEFINE STGCLASS		X	X	X
DEFINE SUB			X	X
DEFINE TOPIC		X	X	X
DELETE AUTHINFO		X	X	X
DELETE BUFFPOOL		X	X	X
DELETE CFSTRUCT		X	X	X
DELETE CHANNEL			X	X

Table 9. Sources from which to run MQSC commands (continued)

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
DELETE NAMELIST		X	X	X
DELETE PROCESS		X	X	X
DELETE PSID			X	X
DELETE QALIAS		X	X	X
DELETE QLOCAL		X	X	X
DELETE QMODEL		X	X	X
DELETE QREMOTE		X	X	X
DELETE STGCLASS		X	X	X
DELETE SUB		X	X	X
DELETE TOPIC		X	X	X
DISPLAY ARCHIVE	X	X	X	X
DISPLAY AUTHINFO		X	X	X
DISPLAY CFSTATUS			X	X
DISPLAY CFSTRUCT		X	X	X
DISPLAY CHANNEL		X	X	X
DISPLAY CHSTATUS			X	X
DISPLAY CLUSQMGR			X	X
DISPLAY CMDSERV	X	X	X	X
DISPLAY CONN		X	X	X
DISPLAY CHINIT		X	X	X
DISPLAY GROUP		X	X	X
DISPLAY LOG	X	X	X	X
DISPLAY NAMELIST		X	X	X
DISPLAY PROCESS		X	X	X
DISPLAY QALIAS		X	X	X
DISPLAY QCLUSTER		X	X	X
DISPLAY QLOCAL		X	X	X
DISPLAY QMGR		X	X	X
DISPLAY QMODEL		X	X	X
DISPLAY QREMOTE		X	X	X
DISPLAY QSTATUS		X	X	X
DISPLAY QUEUE		X	X	X
DISPLAY SECURITY			X	X
DISPLAY STGCLASS		X	X	X
DISPLAY SUB		X	X	X
DISPLAY TOPIC		X	X	X
DISPLAY SYSTEM	X	X	X	X
DISPLAY THREAD		X	X	X
DISPLAY TRACE	X	X	X	X

Table 9. Sources from which to run MQSC commands (continued)

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
DISPLAY USAGE		X	X	X
MOVE QLOCAL		X	X	X
PING CHANNEL			X	X
RECOVER BSDS	X	X	X	X
RECOVER CFSTRUCT			X	X
REFRESH CLUSTER		X	X	X
REFRESH QMGR		X	X	X
REFRESH SECURITY		X	X	X
RESET CHANNEL			X	X
RESET CLUSTER		X	X	X
RESET QSTATS		X	X	X
RESET TPIPE			X	X
RESOLVE CHANNEL			X	X
RESOLVE INDOUBT		X	X	X
RESUME QMGR			X	X
RVERIFY SECURITY		X	X	X
SET ARCHIVE	X	X	X	X
SET LOG	X	X	X	X
SET SYSTEM	X	X	X	X
START CHANNEL			X	X
START CHINIT		X	X	X
START CMDSERV	X	X	X	
START LISTENER			X	X
START QMGR			X	
START TRACE	X	X	X	X
STOP CHANNEL			X	X
STOP CHINIT			X	X
STOP CMDSERV	X	X	X	
STOP LISTENER			X	X
STOP QMGR			X	X
STOP TRACE	X	X	X	X
SUSPEND QMGR			X	X

In the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*), each command description identifies the sources from which that command can be run.

Initialization commands

Initialization commands can be used to control the queue manager startup.

Commands in the initialization input data sets are processed when WebSphere MQ is initialized on queue manager startup. Three types of command can be issued from the initialization input data sets:

- Commands to define WebSphere MQ entities that cannot be recovered (DEFINE BUFFPOOL and DEFINE PSID for example).

These commands must reside in the data set identified by the DDname CSQINP1. They are processed before the restart phase of initialization. They cannot be issued through the console, operations and control panels, or an application program. The responses to these commands are written to the sequential data set that you refer to in the CSQOUT1 statement of the started task procedure.

- Commands to define WebSphere MQ objects that are recoverable after restart. These definitions must be specified in the data set identified by the DDname CSQINP2. They are stored in page set zero. CSQINP2 is processed after the restart phase of initialization. The responses to these commands are written to the sequential data set that you refer to in the CSQOUT2 statement of the started task procedure.
- Commands to manipulate WebSphere MQ objects. These commands must also be specified in the data set identified by the DDname CSQINP2. For example, the WebSphere MQ-supplied sample contains an ALTER QMGR command to specify a dead-letter queue for the subsystem. The response to these commands is written to the CSQOUT2 output data set.

Note: If WebSphere MQ objects are defined in CSQINP2, WebSphere MQ attempts to redefine them each time the queue manager is started. If the queues already exist, the attempt to define them fails. If you need to define your objects in CSQINP2, you can avoid this problem by using the REPLACE parameter of the DEFINE commands, however, this overrides any changes that were made during the previous run of the queue manager.

Sample initialization data set members are supplied with WebSphere MQ for z/OS. They are described in



Sample definitions supplied with WebSphere MQ (*WebSphere MQ V7.1 Product Overview Guide*).

Initialization commands for distributed queuing

You can also use the CSQINP2 initialization data set for the START CHINIT command. If you need a series of other commands to define your distributed queuing environment (for example, starting listeners), WebSphere MQ provides a third initialization input data set, called CSQINPX, that is processed as part of the channel initiator started task procedure.

The MQSC commands contained in the data set are executed at the end of channel initiator initialization, and output is written to the data set specified by the CSQOUTX DD statement. You might use the CSQINPX initialization data set to start listeners for example.

A sample channel initiator initialization data set member is supplied with WebSphere MQ for z/OS. It is described in  Sample definitions supplied with WebSphere MQ (*WebSphere MQ V7.1 Product Overview Guide*).

Initialization commands for publish/Subscribe

If you need a series of commands to define your publish/subscribe environment (for example, when defining subscriptions), WebSphere MQ provides a fourth initialization input data set, called CSQINPT.

The MQSC commands contained in the data set are executed at the end of publish/subscribe initialization, and output is written to the data set specified by the CSQOUTT DD statement. You might use the CSQINPT initialization data set to define subscriptions for example.

A sample publish/subscribe initialization data set member is supplied with WebSphere MQ for z/OS. It is described in  Sample definitions supplied with WebSphere MQ (*WebSphere MQ V7.1 Product Overview Guide*).

The WebSphere MQ for z/OS utilities

WebSphere MQ for z/OS provides a set of utility programs that you can use to help with system administration.

WebSphere MQ for z/OS supplies a set of utility programs to help you perform various administrative tasks, including the following:

- Perform backup, restoration, and reorganization tasks.
- Issue commands and process object definitions.
- Generate data-conversion exits.
- Modify the bootstrap data set.
- List information about the logs.
- Print the logs.
- Set up Db2 tables and other Db2 utilities.
- Process messages on the dead-letter queue.

The CSQUTIL utility

This is a utility program provided to help you with backup, restore and reorganize tasks. See The CSQUTIL utility for more information.

The data conversion exit utility

The WebSphere MQ for z/OS data conversion exit utility (**CSQUCVX**) runs as a stand-alone utility to create data conversion exit routines.

The change log inventory utility

The WebSphere MQ for z/OS change log inventory utility program (**CSQJU003**) runs as a stand-alone utility to change the bootstrap data set (BSDS). You can use the utility to perform the following functions:

- Add or delete active or archive log data sets.
- Supply passwords for archive logs.

The print log map utility

The WebSphere MQ for z/OS print log map utility program (**CSQJU004**) runs as a stand-alone utility to list the following information:

- Log data set name and log RBA association for both copies of all active and archive log data sets. If dual logging is not active, there is only one copy of the data sets.
- Active log data sets available for new log data.
- Contents of the queue of checkpoint records in the bootstrap data set (BSDS).
- Contents of the archive log command history record.
- System and utility time stamps.

The log print utility

The log print utility program (**CSQ1LOGP**) is run as a stand-alone utility. You can run the utility specifying:

- A bootstrap data set (BSDS)

- Active logs (with no BSDS)
- Archive logs (with no BSDS)

The queue-sharing group utility

The queue-sharing group utility program (**CSQ5PQSG**) runs as a stand-alone utility to set up Db2 tables and perform other Db2 tasks required for queue-sharing groups.

The active log preformat utility

The active log preformat utility (**CSQJUFMT**) formats active log data sets before they are used by a queue manager. If the active log data sets are preformatted by the utility, log write performance is improved on the queue manager's first pass through the active logs.

The dead-letter queue handler utility

The dead-letter queue handler utility program (**CSQUDLQH**) runs as a stand-alone utility. It checks messages that are on the dead-letter queue and processes them according to a set of rules that you supply to the utility.

The CSQUTIL utility

The CSQUTIL utility program is provided with WebSphere MQ for z/OS to help you perform backup, restoration, and reorganization tasks, and to issue commands and process object definitions.

For more information about the CSQUTIL utility program, see  [WebSphere MQ utility program \(CSQUTIL\)](#) (*WebSphere MQ V7.1 Reference*). By using this utility program, you can invoke the following functions:

COMMAND

To issue MQSC commands, to record object definitions, and to make client-channel definition files.

COPY To read the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set, and put them into a sequential file and retain the original queue.

COPYPAGE

To copy whole page sets to larger page sets.

EMPTY

To delete the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set, retaining the definitions of the queues.

FORMAT

To format WebSphere MQ for z/OS page sets.

LOAD

To restore the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set from a sequential file created by the COPY function.

PAGEINFO

To extract page set information from one or more page sets.

RESETPAGE

To copy whole page sets to other page set data sets and reset the log information in the copy.

SCOPY

To copy the contents of a queue to a data set while the queue manager is offline.

SDEFS

To produce a set of define commands for objects while the queue manager is offline.

SLOAD

To restore messages from the destination data set of an earlier COPY or SCOPY operation. SLOAD processes a single queue.

XPARM

To convert a channel initiator parameter load module into queue manager attributes (for migration purposes).

Operating WebSphere MQ for z/OS

Use these basic procedures to operate WebSphere MQ for z/OS.

You can also perform the operations described in this section using the WebSphere MQ Explorer, which is distributed with WebSphere MQ for Windows and WebSphere MQ for Linux (x86 and x86-64 platforms). For more information, see “Administration using the IBM WebSphere MQ Explorer” on page 59.

This section contains information about the following topics:

Issuing queue manager commands

You can issue WebSphere MQ control commands from a z/OS console or with the utility program CSQUTIL. Commands can use command prefix string (CPF) to indicate which WebSphere MQ subsystem processes the command.

You can control most of the operational environment of WebSphere MQ using the WebSphere MQ commands. WebSphere MQ for z/OS supports both the MQSC and PCF types of these commands. This topic describes how to specify attributes using MQSC commands, and so it refers to those commands and attributes using their MQSC command names, rather than their PCF names. For details of the syntax of

the MQSC commands, see  The MQSC commands (*WebSphere MQ V7.1 Reference*). For details of the syntax of the PCF commands, see “Using Programmable Command Formats” on page 7. If you are a suitably authorized user, you can issue WebSphere MQ commands from:

- The initialization input data sets (described in “Initialization commands” on page 245).
- A z/OS console, or equivalent, such as SDSF
- The z/OS master get command routine, MGCRE (SVC 34)
- The WebSphere MQ utility, CSQUTIL (described in  WebSphere MQ utility program (CSQUTIL) (*WebSphere MQ V7.1 Reference*).
- A user application, which can be:
 - A CICS program
 - A TSO program
 - A z/OS batch program
 - An IMS program

See “Writing programs to administer WebSphere MQ” on page 272 for information about this.

Much of the functionality of these commands is provided in a convenient way by the operations and control panels, accessible from TSO and ISPF, and described in “Introducing the operations and control panels” on page 260.

For further information, see

- “Issuing commands from a z/OS console or its equivalent” on page 249
 - Command prefix strings
 - Using the z/OS console to issue commands
 - Command responses
- Issuing commands from the utility program CSQUTIL

Issuing commands from a z/OS console or its equivalent

You can issue all WebSphere MQ commands from a z/OS console or its equivalent. You can also issue WebSphere MQ commands from anywhere where you can issue z/OS commands, such as SDSF or by a program using the MGCRC macro.

The maximum amount of data that can be displayed as a result of a command typed in at the console is 32 KB.

Note:

1. You cannot issue WebSphere MQ commands using the IMS/SSR command format from an IMS terminal. This function is not supported by the IMS adapter.
2. The input field provided by SDSF might not be long enough for some commands, particularly those commands for channels.

Command prefix strings

Each WebSphere MQ command must be prefixed with a command prefix string (CPF), as shown in Figure 51.

Because more than one WebSphere MQ subsystem can run under z/OS, the CPF is used to indicate which WebSphere MQ subsystem processes the command. For example, to start the queue manager for a subsystem called CSQ1, where CPF is '+CSQ1', you issue the command +CSQ1 START QMGR from the operator console. This CPF must be defined in the subsystem name table (for the subsystem CSQ1). This is described in  Defining command prefix strings (CPFs) (*WebSphere MQ V7.1 Installing Guide*). In the examples, the string '+CSQ1' is used as the command prefix.

Using the z/OS console to issue commands

You can type simple commands from the z/OS console, for example, the DISPLAY command in Figure 51. However, for complex commands or for sets of commands that you issue frequently, the other methods of issuing commands are better.

```
+CSQ1 DISPLAY QUEUE(TRANSMIT.QUEUE.PROD) TYPE(QLOCAL)
```

Figure 51. Issuing a DISPLAY command from the z/OS console

Command responses

Direct responses to commands are sent to the console that issued the command. WebSphere MQ supports the *Extended Console Support* (EMCS) function available in z/OS, and therefore consoles with 4 byte IDs can be used. Additionally, all commands except START QMGR and STOP QMGR support the use of Command and Response Tokens (CARTs) when the command is issued by a program using the MGCRC macro.

Issuing commands from the utility program CSQUTIL

You can issue commands from a sequential data set using the COMMAND function of the utility program CSQUTIL. This utility transfers the commands, as messages, to the *system-command input queue* and waits for the response, which is printed together with the original commands in SYSPRINT. For details of this, see  WebSphere MQ utility program (CSQUTIL) (*WebSphere MQ V7.1 Reference*).

Starting and stopping a queue manager:

Use this topic as an introduction to stopping and starting a queue manager.

This section describes how to start and stop a queue manager. It contains information about the following topics:

- “Before you start WebSphere MQ”
- “Starting a queue manager”
- “Stopping a queue manager” on page 252

Starting and stopping a queue manager is relatively straightforward. When a queue manager stops under normal conditions, its last action is to take a termination checkpoint. This checkpoint, and the logs, give the queue manager the information it needs to restart.

This section contains information about the START and STOP commands, and contains a brief overview of start-up after an abnormal termination has occurred.

Before you start WebSphere MQ

After you have installed WebSphere MQ, it is defined as a formal z/OS subsystem. This message appears during any initial program load (IPL) of z/OS:

```
CSQ3110I +CSQ1 CSQ3UR00 - SUBSYSTEM ssnm INITIALIZATION COMPLETE
```

where *ssnm* is the WebSphere MQ subsystem name.

From now on, you can start the queue manager for that subsystem *from any z/OS console that has been authorized to issue system control commands*; that is, a z/OS SYS command group. You must issue the START command from the authorized console, you cannot issue it through JES or TSO.

If you are using queue-sharing groups, you must start RRS first, and then Db2, before you start the queue manager.

Starting a queue manager

You start a queue manager by issuing a START QMGR command. However, you cannot successfully use the START command unless you have appropriate authority. See the “Setting up security on z/OS” on page 520 for information about WebSphere MQ security. Figure 52 shows examples of the START command. (Remember that you must prefix a WebSphere MQ command with a command prefix string (CPF).)

```
+CSQ1 START QMGR  
+CSQ1 START QMGR PARM(NEWLOG)
```

Figure 52. Starting the queue manager from a z/OS console. The second example specifies a system parameter module name.

See  START QMGR (*WebSphere MQ V7.1 Reference*) for information about the syntax of the START QMGR command.

You cannot run the queue manager as a batch job or start it using a z/OS command START. These methods are likely to start an address space for WebSphere MQ that then ends abnormally. Nor can you start a queue manager from the CSQUTIL utility program or a similar user application.

You can, however, start a queue manager from an APF-authorized program by passing a START QMGR command to the z/OS MGCRE (SVC 34) service.

If you are using queue-sharing groups, the associated Db2 systems and RRS must be active when you start the queue manager.

Start options

When you start a queue manager, a system parameter module is loaded. You can specify the name of the system parameter module in one of two ways:

- With the PARM parameter of the /cpf START QMGR command, for example
/cpf START QMGR PARM(CSQ1ZPRM)
- With a parameter in the startup procedure, for example, code the JCL EXEC statement as
//MQM EXEC PGM=CSQYASCP,PARM='ZPARM(CSQ1ZPRM)'

A system parameter module provides information specified when the queue manager was customized. In “User messages issued when the queue manager starts” on page 254, the user message CSQY001I indicates the name of the system parameter module that was used, in this case, CSQ1ZPRM. .

You can also use the ENVPARM option to substitute one or more parameters in the JCL procedure for the queue manager.

For example, you can update your queue manager startup procedure, so that the DDname CSQINP2 is a variable. This means that you can change the CSQINP2 DDname without changing the startup procedure. This is useful for implementing changes, providing backouts for operators, and queue manager operations.

Suppose your start-up procedure for queue manager CSQ1 looked like Figure 53.

```
//CSQ1MSTR PROC INP2=NORM
//MQMESA EXEC PGM=CSQYASCP
//STEPLIB DD DISP=SHR,DSN=th1qua1.SCSQANLE
// DD DISP=SHR,DSN=th1qua1.SCSQAUTH
// DD DISP=SHR,DSN=db2qua1.SDSNLOAD
//BSDS1 DD DISP=SHR,DSN=myqua1.BSDS01
//BSDS2 DD DISP=SHR,DSN=myqua1.BSDS02
//CSQP0000 DD DISP=SHR,DSN=myqua1.PSID00
//CSQP0001 DD DISP=SHR,DSN=myqua1.PSID01
//CSQP0002 DD DISP=SHR,DSN=myqua1.PSID02
//CSQP0003 DD DISP=SHR,DSN=myqua1.PSID03
//CSQINP1 DD DISP=SHR,DSN=myqua1.CSQINP(CSQ1INP1)
//CSQINP2 DD DISP=SHR,DSN=myqua1.CSQINP(CSQ1&INP2.)
//CSQOUT1 DD SYSOUT=*
//CSQOUT2 DD SYSOUT=*
```

Figure 53. Sample start-up procedure

If you then start your queue manager with the command:

```
+CSQ1 START QMGR
```

the CSQINP2 used is a member called CSQ1NORM.

However, suppose you are putting a new suite of programs into production so that the next time you start queue manager CSQ1, the CSQINP2 definitions are to be taken from member CSQ1NEW. To do this, you would start the queue manager with this command:

```
+CSQ1 START QMGR ENVPARM('INP2=NEW')
```

and CSQ1NEW would be used instead of CSQ1NORM. Note: z/OS limits the KEYWORD=value specifications for symbolic parameters (as in INP2=NEW) to 255 characters.

Starting after an abnormal termination

WebSphere MQ automatically detects whether restart follows a normal shutdown or an abnormal termination.

Starting a queue manager after it ends abnormally is different from starting it after the STOP QMGR command has been issued. After STOP QMGR, the system finishes its work in an orderly way and takes a termination checkpoint before stopping. When you restart the queue manager, it uses information from the system checkpoint and recovery log to determine the system status at shutdown.

However, if the queue manager ends abnormally, it terminates without being able to finish its work or take a termination checkpoint. When you restart a queue manager after an abend, it refreshes its knowledge of its status at termination using information in the log, and notifies you of the status of various tasks. Normally, the restart process resolves all inconsistent states. But, in some cases, you must take specific steps to resolve inconsistencies.

User messages on start-up

When you start a queue manager successfully, it produces a set of start-up messages similar to the ones in “User messages issued when the queue manager starts” on page 254.

Stopping a queue manager

Before stopping a queue manager, all WebSphere MQ-related write-to-operator-with-reply (WTOR) messages must receive replies, for example, getting log requests. Each command in Figure 54 terminates a running queue manager.

```
+CSQ1 STOP QMGR
+CSQ1 STOP QMGR MODE(QUIESCE)
+CSQ1 STOP QMGR MODE(FORCE)
+CSQ1 STOP QMGR MODE(RESTART)
```

Figure 54. Stopping a queue manager

The command STOP QMGR defaults to STOP QMGR MODE(QUIESCE).

In QUIESCE mode, WebSphere MQ does not allow any new connection threads to be created, but allows existing threads to continue; it terminates only when all threads have ended. Applications can request to

be notified in the event of the queue manager quiescing. Therefore, use the QUIESCE mode where possible so that applications that have requested notification have the opportunity to disconnect. See

 What happens during termination (*WebSphere MQ V7.1 Product Overview Guide*) for details.

If the queue manager does not terminate in a reasonable time in response to a STOP QMGR MODE(QUIESCE) command, use the DISPLAY CONN command to determine whether any connection threads exist, and take the necessary steps to terminate the associated applications. If there are no threads, issue a STOP QMGR MODE(FORCE) command.

The STOP QMGR MODE(QUIESCE) and STOP QMGR MODE(FORCE) commands deregister WebSphere MQ from the MVS™ Automatic Restart Manager (ARM), preventing ARM from restarting the queue manager automatically. The STOP QMGR MODE(RESTART) command works in the same way as the STOP QMGR MODE(FORCE) command, except that it does not deregister WebSphere MQ from ARM. This means that the queue manager is eligible for immediate automatic restart.

If the WebSphere MQ subsystem is not registered with ARM, the STOP QMGR MODE(RESTART) command is rejected and the following message is sent to the z/OS console:

```
CSQY205I  ARM element arm-element is not registered
```

If this message is not issued, the queue manager is restarted automatically. For more information about ARM, see “Using the z/OS Automatic Restart Manager (ARM)” on page 339.

Only cancel the queue manager address space if STOP QMGR MODE(FORCE) does not terminate the queue manager.

If a queue manager is stopped by either canceling the address space or by using the command STOP QMGR MODE(FORCE), consistency is maintained with connected CICS or IMS systems. Resynchronization of resources is started when a queue manager restarts and is completed when the connection to the CICS or IMS system is established.

Note: When you stop your queue manager, you might find message IEF352I is issued. z/OS issues this message if it detects that failing to mark the address space as unusable would lead to an integrity exposure. You can ignore this message.

Stop messages

After issuing a STOP QMGR command, you get the messages CSQY009I and CSQY002I, for example:

```
CSQY009I +CSQ1 ' STOP QMGR' COMMAND ACCEPTED FROM  
USER(userid), STOP MODE(FORCE)  
CSQY002I +CSQ1 QUEUE MANAGER STOPPING
```

Where *userid* is the user ID that issued the STOP QMGR command, and the MODE parameter depends on that specified in the command.

When the STOP command has completed successfully, the following messages are displayed on the z/OS console:

```
CSQ9022I +CSQ1 CSQYASCP ' STOP QMGR' NORMAL COMPLETION
CSQ3104I +CSQ1 CSQ3EC0X - TERMINATION COMPLETE
```

If you are using ARM, and you did not specify MODE(RESTART), the following message is also displayed:

```
CSQY204I +CSQ1 ARM Deregister for element arm-element type
arm-element-type successful
```

You cannot restart the queue manager until the following message has been displayed:

```
CSQ3100I +CSQ1 CSQ3EC0X - SUBSYSTEM ssnm READY FOR START COMMAND
```

User messages issued when the queue manager starts:

User messages are output when WebSphere MQ, or a channel initiator, are started successfully.

When you start WebSphere MQ successfully, it produces a set of messages similar to the ones for subsystem MQ86, shown here:

```
$HASP373 MQ86MSTR STARTED
IEF403I MQ86MSTR - STARTED
CSQY000I MQ86 IBM WebSphere MQ for z/OS V7
CSQY001I MQ86 QUEUE MANAGER STARTING, USING PARAMETER MODULE CSQZMQ86
CSQ3111I MQ86 CSQYSCMD - EARLY PROCESSING PROGRAM IS V7 LEVEL 004-000
CSQY100I MQ86 SYSTEM parameters ...
CSQY101I MQ86 LOGLOAD=500000
CSQY102I MQ86 CMDUSER=CSQOPR, QMCCSID=0, ROUTCDE=( 1)
CSQY103I MQ86 SMFACCT=NO (00000000), SMFSTAT=NO (00000000), STATIME=30
CSQY104I MQ86 OTMACON= 303
( , ,DFSYDRU0,2147483647,CSQ)
CSQY105I MQ86 TRACSTR=( 1), TRACTBL=999
CSQY106I MQ86 EXITTCB=8, EXITLIM=5, WLMTIME=60, WLMTIMU=MINS
CSQY107I MQ86 QSGDATA=(QH03,DSN910P7,DH48,4,8)
CSQY108I MQ86 RESAUDIT=YES, QINDEXBLD=WAIT, CLCACHE=STATIC
CSQY110I MQ86 LOG parameters ...
CSQY111I MQ86 INBUFF=60, OUTBUFF=4000, MAXRTU=2, MAXARCH=500
CSQY112I MQ86 TWOACTV=NO, TWOARCH=NO, TWOBSDS=NO
CSQY113I MQ86 OFFLOAD=YES, WRTHRSH=256, DEALLCT=0
CSQY120I MQ86 ARCHIVE parameters ...
CSQY121I MQ86 UNIT=3390, UNIT2=, ALCUNIT=CYL, 314
PRIQTY=56, SECQTY=5, BLKSIZE=24576
CSQY122I MQ86 ARCPFX1=MQTST.SUBSYS.MQ86.ARC1, 315
ARCPFX2=MQTST.SUBSYS.MQ86.ARC2, TSTAMP=NO
CSQY123I MQ86 ARCRETN=0, ARCWTOR=YES, ARCWRTC=( 1 ,3 ,4)
CSQY124I MQ86 CATALOG=YES, COMPACT=YES, PROTECT=NO, QUIESCE=5
CSQY201I MQ86 CSQYSTRT ARM REGISTER for element 324
SYSQMGRMQ86 type SYSQMGR successful
IEC161I 056-084,MQ86MSTR,MQ86MSTR,BSDS1,,MQTST.SUBSYS.MQ86.BSDS01, 326
IEC161I MQTST.SUBSYS.MQ86.BSDS01.DATA,ICFCAT.PLEX7.CATALOGA
IEC161I 056-084,MQ86MSTR,MQ86MSTR,BSDS1,,MQTST.SUBSYS.MQ86.BSDS01, 327
IEC161I MQTST.SUBSYS.MQ86.BSDS01.INDEX,ICFCAT.PLEX7.CATALOGA
IEC161I 062-086,MQ86MSTR,MQ86MSTR,BSDS1,,MQTST.SUBSYS.MQ86.BSDS01, 328
IEC161I MQTST.SUBSYS.MQ86.BSDS01.DATA,ICFCAT.PLEX7.CATALOGA
```

```

CSQJ127I MQ86 SYSTEM TIME STAMP FOR BSDS=2008-01-29 10:58:32.83
CSQJ001I MQ86 CURRENT COPY 1 ACTIVE LOG DATA SET IS 330
DSNAME=MQTST.SUBSYS.MQ86.LOGCOPY1.DS01, STARTRBA=000000000000
ENDRBA=00000275FFFF
CSQJ099I MQ86 LOG RECORDING TO COMMENCE WITH 331
STARTRBA=00000108B000
CSQW130I MQ86 'GLOBAL' TRACE STARTED, ASSIGNED TRACE NUMBER 01
CSQ5001I MQ86 CSQ5CONN Connected to DB2 DH48
CSQP007I MQ86 Page set 1 uses buffer pool 1
CSQP007I MQ86 Page set 2 uses buffer pool 2
CSQP007I MQ86 Page set 3 uses buffer pool 3
CSQP007I MQ86 Page set 4 uses buffer pool 0
CSQP007I MQ86 Page set 5 uses buffer pool 1
CSQP007I MQ86 Page set 6 uses buffer pool 2
CSQP007I MQ86 Page set 7 uses buffer pool 3
CSQP007I MQ86 Page set 8 uses buffer pool 0
CSQP007I MQ86 Page set 9 uses buffer pool 1
CSQP007I MQ86 Page set 99 uses buffer pool 2
CSQP007I MQ86 Page set 0 uses buffer pool 3
CSQY220I MQ86 Queue manager is using 26 MB of local 362
storage, 1413 MB are free
CSQV452I MQ86 CSQVXLDL Cluster workload exits not available
CSQR001I MQ86 RESTART INITIATED
CSQR003I MQ86 RESTART - PRIOR CHECKPOINT RBA=000000002158
CSQR004I MQ86 RESTART - UR COUNTS - 379
IN COMMIT=1, INDOUBT=0, INFLIGHT=0, IN BACKOUT=0
CSQR007I MQ86 UR STATUS 380
T CON-ID          THREAD-XREF          S    URID          TIME
- - - - -
C IYRMQ86 281A2790C1D7F0F20000738C C00000108120E 2008-01-29 10:58:44
CSQI049I MQ86 Page set 0 has media recovery 381
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 1 has media recovery 382
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 2 has media recovery 383
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 3 has media recovery 384
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 4 has media recovery 385
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 5 has media recovery 386
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 6 has media recovery 387
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 7 has media recovery 388
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 8 has media recovery 389
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 9 has media recovery 390
RBA=000000002158, checkpoint RBA=000000002158
CSQI049I MQ86 Page set 99 has media recovery 391
RBA=000000002158, checkpoint RBA=000000002158
CSQE005I MQ86 Structure CSQ_ADMIN connected as 393
IXL014I IXLCONN REQUEST FOR STRUCTURE QH03CSQ_ADMIN 392
CSQEQH03MQ8604, version=C1DEA7C6AB88EB45 00040002
WAS SUCCESSFUL. JOBNAME: MQ86MSTR ASID: 03A6
CSQE021I MQ86 Structure CSQ_ADMIN connection as 394
CSQEQH03MQ8604 warning, RC=00000004 reason=02010407 codes=00000000
CONNECTOR NAME: CSQEQH03MQ8604 CFNAME: P7CF05
00000000 00000000

```

```

IXL014I IXLCONN REQUEST FOR STRUCTURE QH03APPL1 396
CSQE005I MQ86 Structure APPL1 connected as 397
CSQEQH03MQ8604, version=C1DEA8DB0311932E 00020002
WAS SUCCESSFUL. JOBNAME: MQ86MSTR ASID: 03A6
CONNECTOR NAME: CSQEQH03MQ8604 CFNAME: P7CF05
CSQE007I MQ86 EEPLDISCFAILCONNECTION event received 398
for structure APPL1 connection name CSQEQH03MQ8604
CSQE011I MQ86 Recovery phase 1 started for structure 399
APPL1 connection name CSQEQH03MQ8604
CSQE013I MQ86 Recovery phase 1 completed for 400
structure APPL1 connection name CSQEQH03MQ8604
CSQE012I MQ86 Recovery phase 2 started for structure 401
APPL1 connection name CSQEQH03MQ8604
CSQE014I MQ86 Recovery phase 2 completed for 402
structure APPL1 connection name CSQEQH03MQ8604
CSQE006I MQ86 Structure APPL1 connection name 407
CSQEQH03MQ8604 disconnected
IXL014I IXLCONN REQUEST FOR STRUCTURE QH03APPL2 427
CSQE005I MQ86 Structure APPL2 connected as 428
CSQEQH03MQ8604, version=C1DEA8E6BB67A14F 00040002
WAS SUCCESSFUL. JOBNAME: MQ86MSTR ASID: 03A6
CONNECTOR NAME: CSQEQH03MQ8604 CFNAME: P7CF05
CSQE007I MQ86 EEPLDISCFAILCONNECTION event received 429
for structure APPL2 connection name CSQEQH03MQ8604
CSQE011I MQ86 Recovery phase 1 started for structure 430
APPL2 connection name CSQEQH03MQ8604
CSQE013I MQ86 Recovery phase 1 completed for 431
structure APPL2 connection name CSQEQH03MQ8604
CSQE012I MQ86 Recovery phase 2 started for structure 432
APPL2 connection name CSQEQH03MQ8604
CSQE014I MQ86 Recovery phase 2 completed for 433
structure APPL2 connection name CSQEQH03MQ8604
CSQE006I MQ86 Structure APPL2 connection name 436
CSQEQH03MQ8604 disconnected
CSQE005I MQ86 Structure APPLSYS connected as 440
IXL014I IXLCONN REQUEST FOR STRUCTURE QH03APPLSYS 439
CSQEQH03MQ8604, version=C1DEA7CC633AC140 00040003
WAS SUCCESSFUL. JOBNAME: MQ86MSTR ASID: 03A6
CONNECTOR NAME: CSQEQH03MQ8604 CFNAME: P7CF05
CSQE007I MQ86 EEPLDISCFAILCONNECTION event received 441
for structure APPLSYS connection name CSQEQH03MQ8604
CSQE011I MQ86 Recovery phase 1 started for structure 442
APPLSYS connection name CSQEQH03MQ8604
CSQE013I MQ86 Recovery phase 1 completed for 443
structure APPLSYS connection name CSQEQH03MQ8604
CSQE012I MQ86 Recovery phase 2 started for structure 444
APPLSYS connection name CSQEQH03MQ8604
CSQE014I MQ86 Recovery phase 2 completed for 445
structure APPLSYS connection name CSQEQH03MQ8604
CSQE006I MQ86 Structure APPLSYS connection name 448
CSQEQH03MQ8604 disconnected
CSQR030I MQ86 Forward recovery log range 449
from RBA=000000002158 to RBA=00000108A1A2
CSQR005I MQ86 RESTART - FORWARD RECOVERY COMPLETE - 450
IN COMMIT=0, INDOUBT=0
CSQR032I MQ86 Backward recovery log range 451
from RBA=00000108A1A2 to RBA=00000108A1A2
CSQR006I MQ86 RESTART - BACKWARD RECOVERY COMPLETE - 452
INFLIGHT=0, IN BACKOUT=0
CSQH001I MQ86 CSQHINSQ Security using uppercase classes

```

```

CSQH021I MQ86 CSQHINSQ SUBSYSTEM security switch set 456
OFF, profile 'MQ86.NO.SUBSYS.SECURITY' found
CSQR002I MQ86 RESTART COMPLETED
CSQP018I MQ86 CSQPBCKW CHECKPOINT STARTED FOR ALL BUFFER POOLS
CSQP019I MQ86 CSQP1DWP CHECKPOINT COMPLETED FOR 461
MQ86 DISPLAY CONN(*) TYPE(CONN) ALL WHERE(UOWSTATE EQ UNRESOLVED)
BUFFER POOL 3, 10 PAGES WRITTEN
CSQP019I MQ86 CSQP1DWP CHECKPOINT COMPLETED FOR 463
BUFFER POOL 1, 17 PAGES WRITTEN
CSQP019I MQ86 CSQP1DWP CHECKPOINT COMPLETED FOR 464
BUFFER POOL 2, 14 PAGES WRITTEN
CSQP019I MQ86 CSQP1DWP CHECKPOINT COMPLETED FOR 465
BUFFER POOL 0, 47 PAGES WRITTEN
CSQP021I MQ86 Page set 0 new media recovery 466
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 1 new media recovery 467
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 2 new media recovery 468
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 3 new media recovery 469
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 4 new media recovery 470
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 5 new media recovery 471
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 6 new media recovery 472
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 7 new media recovery 473
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 8 new media recovery 474
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 9 new media recovery 475
RBA=00000108D142, checkpoint RBA=00000108D142
CSQP021I MQ86 Page set 99 new media recovery 476
RBA=00000108D142, checkpoint RBA=00000108D142
CSQI007I MQ86 CSQIRBLD BUILDING IN-STORAGE INDEX FOR 477
QUEUE SYSTEM.DURABLE.SUBSCRIBER.QUEUE
CSQI006I MQ86 CSQIRBLD COMPLETED IN-STORAGE INDEX FOR 478
QUEUE SYSTEM.DURABLE.SUBSCRIBER.QUEUE
CSQN011I MQ86 COMMAND SERVER STATUS IS ENABLED
CSQM297I MQ86 CSQMDRTC NO CONN FOUND MATCHING REQUEST CRITERIA
CSQ9022I MQ86 CSQMDRTC ' DISPLAY CONN' NORMAL COMPLETION
MQ86 DISPLAY SYSTEM
MQ86 DISPLAY LOG
CSQJ322I MQ86 DISPLAY SYSTEM report ... 484
Parameter Initial value SET value
-----
LOGLOAD 500000
CMDUSER CSQOPR
QMCCSID 0
ROUTCDE 1
SMFACCT NO
SMFSTAT NO
STATIME 30
OTMACON
GROUP
MEMBER
DRUEXIT DFSYDRU0
AGE 2147483647
TPIPEPFX CSQ

```

```

TRACSTR      1
TRACTBL     999
EXITTCB     8
EXITLIM     5
WLMTIME     60
MQ86 DISPLAY ARCHIVE
WLMTIMU     MINS
QSGDATA
  QSGNAME    QH03
  DSGNAME    DSN910P7
  DB2NAME    DH48
  DB2SERV    4
  DB2BLOB    8
RESAUDIT    YES
QINDXBLD    WAIT
CLCACHE     STATIC
End of SYSTEM report
CSQ9022I MQ86 CSQJC001 ' DISPLAY SYSTEM' NORMAL COMPLETION
CSQJ322I MQ86 DISPLAY LOG report ... 487
Parameter   Initial value      SET value

```

MQ86 DISPLAY USAGE

```

INBUFF      60
OUTBUFF     4000
MAXRTU      2
MAXARCH     500
TWOACTV     NO
TWOARCH     NO
TWOBSDS     NO
OFFLOAD     YES
WRTHRSH     256
DEALLCT     0

```

End of LOG report

CSQJ370I MQ86 LOG status report ... 489

Copy %Full DSName

```

  1      43  MQTST.SUBSYS.MQ86.LOGCOPY1.DS01
  2          Inactive

```

Restarted at 2008-01-29 11:00:23 using RBA=00000108B000

Latest RBA=0000010914B7

Offload task is AVAILABLE

Full logs to offload - 0 of 4

CSQ9022I MQ86 CSQJC001 ' DISPLAY LOG' NORMAL COMPLETION

CSQM050I MQ86 CSQMIGQA Intra-group queuing agent starting, TCB=0092AAD0

CSQJ322I MQ86 DISPLAY ARCHIVE report ... 492

Parameter Initial value SET value

CSQM073I MQ86 CSQMDURR Loading of durable subscribers started

UNIT 3390

UNIT2

CSQM074I MQ86 CSQMDURR Loading of durable subscribers finished

ALCUNIT CYL

PRIQTY 56

SECQTY 5

CSQM075I MQ86 CSQMDURR Consolidation of durable subscribers started

BLKSIZE 24576

ARCPFX1 MQTST.SUBSYS.MQ86.ARC1

ARCPFX2 MQTST.SUBSYS.MQ86.ARC2

TSTAMP NO

ARCRETN 0

ARCWTOR YES

```

ARCWRTC      1 ,3 ,4
CATALOG      YES
CSQM076I MQ86 CSQMDURR Consolidation of durable subscribers finished
COMPACT      YES
PROTECT      NO
QUIESCE      5
End of ARCHIVE report
CSQJ325I MQ86 ARCHIVE tape unit report ... 498
Addr St CorrelID VolSer DSName
-----

```

```

No tape archive reading activity
End of tape unit report

```

```

CSQ9022I MQ86 CSQJC001 ' DISPLAY ARCHIVE' NORMAL COMPLETION
CSQI010I MQ86 Page set usage ... 500

```

Page set	Buffer pool	Total pages	Unused pages	Persistent data pages	NonPersist data pages	Expansion	count
- 0	0	5038	5010	28	0	USER	0
- 1	1	50035	50033	2	0	USER	0
- 2	2	50035	50032	3	0	USER	0
- 3	3	50035	50035	0	0	USER	0
- 4	0	50035	50033	2	0	USER	0
CSQY022I MQ86 QUEUE MANAGER INITIALIZATION COMPLETE							
- 5	1	50035	50035	0	0	USER	0
- 6	2	50035	50035	0	0	USER	0
CSQ9022I MQ86 CSQYASCP 'START QMGR' NORMAL COMPLETION							
- 7	3	50035	50035	0	0	USER	0
- 8	0	50035	50035	0	0	USER	0
- 9	1	50035	50035	0	0	USER	0
- 99	0	50035	50035	0	0	USER	0

```

End of page set report

```

```

CSQP001I MQ86 Buffer pool 0 has 1000 buffers, 869 (86%) stealable
CSQP001I MQ86 Buffer pool 1 has 1000 buffers, 989 (98%) stealable
CSQP001I MQ86 Buffer pool 2 has 1000 buffers, 983 (98%) stealable
CSQP001I MQ86 Buffer pool 3 has 1000 buffers, 994 (99%) stealable
CSQI024I MQ86 CSQIDUSE Restart RBA for system as 508
configured=000000000000
CSQ9022I MQ86 CSQIDUSE ' DISPLAY USAGE' NORMAL COMPLETION
IXL014I IXLCONN REQUEST FOR STRUCTURE QH03APPLSYS 517
CSQE005I MQ86 Structure APPLSYS connected as 518
CSQEQH03MQ8604, version=C1DEA7CC633AC140 00040004
WAS SUCCESSFUL. JOBNAME: MQ86MSTR ASID: 03A6
CONNECTOR NAME: CSQEQH03MQ8604 CFNAME: P7CF05

```

1. The actual messages you get depend on the state of your system, the objects that you have defined, the parameters you specify, and the page sets and other data sets that you are using.
If you are starting the queue manager for the first time, or if the queue manager is not in a queue-sharing group, the messages are somewhat different.
2. If any of the values in message CSQR004I is not zero, message CSQR007I is issued to provide the restart status table.
3. Messages CSQP018I and CSQP019I are issued every time a checkpoint is taken. At checkpoint time, all pages that have not been changed for the two checkpoints are written out to DASD. Message CSQP019I is issued for each buffer pool, giving the number of pages written. You can use this information when balancing page sets in buffer pools.
4. There might be periods during startup when no messages are produced; for example, if you are using indexed queues, no messages are produced while the queue indexes are being rebuilt.
5. The messages provide a wide-ranging set of information about the state of the queue manager, including:

- system parameter values from the system parameter module, and as set by commands in the initialization data sets
- security switch settings
- page set and buffer pool status
- unresolved units of work
- log status

A number of commands are issued internally during startup (for example DISPLAY CONN and DISPLAY SYSTEM) to provide some of the information, in addition to information that is provided directly.

Messages when a channel initiator is started

When you successfully start a channel initiator, messages similar to those shown here are output:

```
CSQM138I MQ04 CSQMSCHI CHANNEL INITIATOR STARTING
CSQ9022I MQ04 CSQXCRPS ' START CHINIT' NORMAL COMPLETION
CSQM137I MQ04 CSQMDDQM DISPLAY DQM COMMAND ACCEPTED
CSQX830I MQ04 CSQXRDQM Channel initiator active
CSQX831I MQ04 CSQXRDQM 8 adapter subtasks started, 8 requested
CSQX832I MQ04 CSQXRDQM 5 dispatchers started, 5 requested
CSQX833I MQ04 CSQXRDQM 0 SSL server subtasks started, 0 requested
CSQX840I MQ04 CSQXRDQM 0 channels current, maximum 1520
CSQX841I MQ04 CSQXRDQM 0 channels active, maximum 1520, including 0 paused
CSQX842I MQ04 CSQXRDQM 0 channels starting,0 stopped, 0 retrying
CSQX836I MQ04 CSQXRDQM Maximum channels - TCP/IP 1520, LU 6.2 1520
CSQX845I MQ04 CSQXRDQM TCP/IP system name is TCPIP
CSQX848I MQ04 CSQXRDQM TCP/IP listener INDISP=QMGR not started
CSQX849I MQ04 CSQXRDQM LU 6.2 listener INDISP=QMGR not started
CSQ9022I MQ04 CSQXCRPS ' DISPLAY DQM' NORMAL COMPLETION
CSQM134I MQ04 CSQMSLIS START LISTENER TRPTYPE(TCP) COMMAND ACCEPTED
CSQ9022I MQ04 CSQXCRPS ' START LISTENER' NORMAL COMPLETION
CSQM134I MQ04 CSQMSLIS START LISTENER TRPTYPE(LU62) COMMAND ACCEPTED
CSQ9022I MQ04 CSQXCRPS ' START LISTENER' NORMAL COMPLETION
```

Introducing the operations and control panels

You can use the WebSphere MQ operations and control panels to perform administration tasks on WebSphere MQ objects. Use this topic as an introduction to the commands, and control panels.

You use these panels for defining, displaying, altering, or deleting WebSphere MQ objects. Use the panels for day-to-day administration and for making small changes to objects. If you are setting up or changing many objects, use the COMMAND function of the CSQUTIL utility program.

The operations and control panels support the controls for the channel initiator (for example, to start a channel or a TCP/IP listener), for clustering, and for security. They also enable you to display information about threads and page set usage.

The panels work by sending MQSC type WebSphere MQ commands to a queue manager, through the system command input queue.

Note:

1. There are no panels for the direct manipulation of topic objects or subscriptions. Using a command facility, you can administer publish/subscribe definitions and other system controls that are not directly available from other panels.
2. You cannot issue the WebSphere MQ commands directly from the command line in the panels.

3. To use the operations and control panels, you must have the correct security authorization; this is described in the “User IDs for command security and command resource security” on page 568.

Invocation and rules for the operations and control panels:

You can control WebSphere MQ and issue control commands through the ISPF panels.

If the ISPF/PDF primary options menu has been updated for WebSphere MQ, you can access the WebSphere MQ operations and control panels from that menu. For details about updating the menu, see the  Task 20: Set up the operations and control panels (*WebSphere MQ V7.1 Installing Guide*).

You can access the WebSphere MQ operations and control panels from the TSO command processor panel (typically option 6 on the ISPF/PDF primary options menu). The name of the exec that you run to do this is CSQOREXX. It has two parameters; `thlqual` is the high-level qualifier for the WebSphere MQ libraries to be used, and `langletter` is the letter identifying the national language libraries to be used (for example, E for U.S. English). The parameters can be omitted if the WebSphere MQ libraries are permanently installed in your ISPF setup. Alternatively, you can issue CSQOREXX from the TSO command line.

These panels are designed to be used by operators and administrators with a minimum of formal training. Read these instructions with the panels running and try out the different tasks suggested.

Note: While using the panels, temporary dynamic queues with names of the form SYSTEM.CSQOREXX.* are created.

Rules for the operations and control panels

See  Rules for naming IBM WebSphere MQ objects (*WebSphere MQ V7.1 Product Overview Guide*) about the general rules for WebSphere MQ character strings and names. However, there are some rules that apply only to the operations and control panels:

- Do not enclose strings, for example descriptions, in single or double quotation marks.
- If you include an apostrophe or quotation mark in a text field, you do not have to repeat it or add an escape character. The characters are saved exactly as you type them; for example:

This is Maria's queue

The panel processor doubles them for you to pass them to WebSphere MQ. However, if it has to truncate your data to do this, it does so.

- You can use uppercase or lowercase characters in most fields, and they are folded to uppercase characters when you press Enter. The exceptions are:
 - Storage class names and coupling facility structure names, which must start with uppercase A through Z and be followed by uppercase A through Z or numeric characters.
 - Certain fields that are not translated. These include:
 - Application ID
 - Description
 - Environment data
 - Object names (but if you use a lowercase object name, you might not be able to enter it at a z/OS console)
 - Remote system name
 - Trigger data

- User data
- In names, leading blanks and leading underscores are ignored. Therefore, you cannot have object names beginning with blanks or underscores.
- Underscores are used to show the extent of blank fields. When you press Enter, trailing underscores are replaced by blanks.
- Many description and text fields are presented in multiple parts, each part being handled by WebSphere MQ independently. This means that trailing blanks *are retained* and the text is not contiguous.

Blank fields

When you specify the **Define** action for a WebSphere MQ object, each field on the define panel contains a value. See the general help (extended help) for the display panels for information about where WebSphere MQ gets the values. If you type over a field with blanks, and blanks are not allowed, WebSphere MQ puts the installation default value in the field or prompts you to enter the required value.

When you specify the **Alter** action for a WebSphere MQ object, each field on the alter panel contains the current value for that field. If you type over a field with blanks, and blanks are not allowed, the value of that field is left unchanged.

Objects and actions:

The operations and control panels offer you many different types of object and a number of actions that you can perform on them.

The actions are listed on the initial panel and enable you to manipulate the objects and display information about them. These objects include all the WebSphere MQ objects, together with some extra ones. The objects fall into the following categories.

- Queues, processes, authentication information objects, namelists, storage classes and CF structures
- Channels
- Cluster objects
- Queue manager and security
- Connections
- System

Refer to Actions for a cross reference table of the actions which can be taken with the WebSphere MQ objects.

Queues, processes, authentication information objects, namelists, storage classes and CF structures

These are the basic WebSphere MQ objects. There can be many of each type. They can be listed, listed with filter, defined, and deleted, and have attributes that can be displayed and altered, using the LIST or DISPLAY, LIST with FILTER, DEFINE LIKE, MANAGE, and ALTER actions. (Objects are deleted using the MANAGE action.)

This category consists of the following objects:

QLOCAL	Local queue
QREMOTE	Remote queue
QALIAS	Alias queue for indirect reference to a queue
QMODEL	Model queue for defining queues dynamically
QUEUE	Any type of queue
QSTATUS	Status of a local queue
PROCESS	Information about an application to be started when a trigger event occurs
AUTHINFO	Authentication information: definitions required to perform Certificate Revocation List (CRL) checking using LDAP servers
NAMELIST	List of names, such as queues or clusters
STGCLASS	Storage class
CFSTRUCT	coupling facility (CF) structure
CFSTATUS	Status of a CF structure

Channels

Channels are used for distributed queuing. There can be many of each type, and they can be listed, listed with filter, defined, deleted, displayed, and altered. They also have other functions available using the START, STOP and PERFORM actions. PERFORM provides reset, ping, and resolve channel functions.

This category consists of the following objects:

CHANNEL	Any type of channel
SENDER	Sender channel
SERVER	Server channel
RECEIVER	Receiver channel
REQUESTER	Requester channel
CLUSRCVR	Cluster-receiver channel
CLUSSDR	Cluster-sender channel
SVRCONN	Server-connection channel
CLNTCONN	Client-connection channel
CHSTATUS	Status of a channel connection

Cluster objects

Cluster objects are created automatically for queues and channels that belong to a cluster. The base queue and channel definitions can be on another queue manager. There can be many of each type, and names can be duplicated. They can be listed, listed with filter, and displayed. PERFORM, START, and STOP are also available through the LIST actions.

This category consists of the following objects:

CLUSQ	Cluster queue, created for a queue that belongs to a cluster
CLUSCHL	Cluster channel, created for a channel that belongs to a cluster
CLUSQMGR	Cluster queue manager, the same as a cluster channel but identified by its queue manager name

Cluster channels and cluster queue managers do have the PERFORM, START and STOP actions, but only indirectly through the DISPLAY action.

Queue manager and security

Queue manager and security objects have a single instance. They can be listed, and have attributes that can be displayed and altered (using the LIST or DISPLAY, and ALTER actions), and have other functions available using the PERFORM action.

This category consists of the following objects:

MANAGER	Queue manager – the PERFORM action provides suspend and resume cluster functions
SECURITY	Security functions – the PERFORM action provides refresh and reverify functions

Connection

Connections can be listed, listed with filter and displayed.

This category consists only of the connection object, CONNECT.

System

A collection of other functions. This category consists of the following objects:

SYSTEM	System functions
CONTROL	Synonym for SYSTEM

The functions available are:

LIST or DISPLAY	Display queue-sharing group, distributed queuing, page set, or data set usage information.
PERFORM	Refresh or reset clustering
START	Start the channel initiator or listeners
STOP	Stop the channel initiator or listeners

Actions

The actions that you can perform for each type of object are shown in the following table:

Table 10. Valid operations and control panel actions for WebSphere MQ objects

Object	Alter	Define like	Manage (1)	List or Display	List with Filter	Perform	Start	Stop
AUTHINFO	X	X	X	X	X			
CFSTATUS				X				
CFSTRUCT	X	X	X	X	X			
CHANNEL	X	X	X	X	X	X	X	X
CHSTATUS				X	X			
CLNTCONN	X	X	X	X	X			
CLUSCHL				X	X	X(2)	X(2)	X(2)

Table 10. Valid operations and control panel actions for WebSphere MQ objects (continued)

Object	Alter	Define like	Manage (1)	List or Display	List with Filter	Perform	Start	Stop
CLUSQ				X	X			
CLUSQMGR				X	X	X(2)	X(2)	X(2)
CLUSRCVR	X	X	X	X	X	X	X	X
CLUSDR	X	X	X	X	X	X	X	X
CONNECT				X	X			
CONTROL				X		X	X	X
MANAGER	X			X		X		
NAMELIST	X	X	X	X	X			
PROCESS	X	X	X	X	X			
QALIAS	X	X	X	X	X			
QLOCAL	X	X	X	X	X			
QMODEL	X	X	X	X	X			
QREMOTE	X	X	X	X	X			
QSTATUS				X	X			
QUEUE	X	X	X	X	X			
RECEIVER	X	X	X	X	X	X	X	X
REQUESTER	X	X	X	X	X	X	X	X
SECURITY	X			X		X		
SENDER	X	X	X	X	X	X	X	X
SERVER	X	X	X	X	X	X	X	X
SVRCONN	X	X	X	X	X		X	X
STGCLASS	X	X	X	X	X			
SYSTEM				X		X	X	X

Note:
1. Provides **Delete** and other functions
2. Using the **List or Display** action

Object dispositions:

You can specify the *disposition* of the object with which you need to work. The disposition signifies where the object **definition** is kept, and how the object behaves.

The disposition is significant only if you are working with any of the following object types:

- queues
- channels
- processes
- namelists
- storage classes
- authentication information objects

If you are working with other object types, the disposition is disregarded.

Permitted values are:

- Q** QMGR. The object definitions are on the page set of the queue manager and are accessible only by the queue manager.
- C** COPY. The object definitions are on the page set of the queue manager and are accessible only by the queue manager. They are local copies of objects defined as having a disposition of GROUP.
- P** PRIVATE. The object definitions are on the page set of the queue manager and are accessible only by the queue manager. The objects have been defined as having a disposition of QMGR or COPY.
- G** GROUP. The object definitions are in the shared repository, and are accessible by all queue managers in the queue-sharing group.
- S** SHARED. This disposition applies only to local queues. The queue definitions are in the shared repository, and are accessible by all queue managers in the queue-sharing group.
- A** ALL. If the action queue manager is either the target queue manager, or *, objects of **all** dispositions are included; otherwise, objects of QMGR and COPY dispositions only are included. This is the default.

Selecting a queue manager, defaults, and levels using the ISPF control panel:

You can use the CSQOREXX exec in ISPF to control your queue managers.

While you are viewing the initial panel, you are not connected to any queue manager. However, as soon as you press Enter, you are connected to the queue manager, or a queue manager in the queue-sharing group named in the **Connect name** field. You can leave this field blank; this means that you are using the default queue manager for batch applications. This is defined in CSQBDEFV (see  Task 19: Set up Batch, TSO, and RRS adapters (*WebSphere MQ V7.1 Installing Guide*) for information about this).

Use the **Target queue manager** field to specify the queue manager where the actions you request are to be performed. If you leave this field blank, it defaults to the queue manager specified in the **Connect name** field. You can specify a target queue manager that is not the one you connect to. In this case, you would normally specify the name of a remote queue manager object that provides a queue manager alias definition (the name is used as the *ObjectQMgrName* when opening the command input queue). To do this, you must have suitable queues and channels set up to access the remote queue manager.

The **Action queue manager** allows you to specify a queue manager that is in the same queue-sharing group as the queue manager specified in the **Target queue manager** field to be the queue manager where the actions you request are to be performed. If you specify * in this field, the actions you request are performed on all queue managers in the queue-sharing group. If you leave this field blank, it defaults to the value specified in the **Target queue manager** field. The **Action queue manager** field corresponds to using the CMDSCOPE command modifier described in  The MQSC commands (*WebSphere MQ V7.1 Reference*).

Queue manager defaults

If you leave any queue manager fields blank, or choose to connect to a queue-sharing group, a secondary window opens when you press Enter. This window confirms the names of the queue managers you will be using. Press Enter to continue. When you return to the initial panel after having made some requests, you find fields completed with the actual names.

Queue manager levels

The Operations and Control panels work satisfactorily only with queue managers that are running on z/OS, and with command levels that match that of the panels, currently 600 or 710.

If these conditions are not met, it is likely that actions work only partially, incorrectly, or not at all, and that the replies from the queue manager are not recognized.

If the action queue manager is not at command level 710, some fields are not displayed, and some values cannot be entered. A few objects and actions are disallowed. In such cases, a secondary window opens asking for you to confirm that you want to proceed.

Using the function keys and command line with the ISPF control panels:

To use the panels, you must use the function keys or enter the equivalent commands in the ISPF control panel command area.

- Function keys
 - Processing your actions
 - “Displaying WebSphere MQ user messages”
 - Canceling your actions
 - Getting help
- Using the command line

Function keys

The function keys have special settings for WebSphere MQ. (This means that you cannot use the ISPF default values for the function keys; if you have previously used the KEYLIST OFF ISPF command anywhere, you must type KEYLIST ON in the command area of any operations and control panel and then press Enter to enable the WebSphere MQ settings.)

These function key settings can be displayed on the panels, as shown in Figure 55 on page 268. If the settings are not shown, type PFSHOW in the command area of any operations and control panel and then press Enter. To remove the display of the settings, use the command PFSHOW OFF.

The function key settings in the operations and control panels conform to CUA standards. Although you can change the key setting through normal ISPF procedures (such as the KEYLIST utility), you are not recommended to do so.

Note: Using the PFSHOW and KEYLIST commands affects any other logical ISPF screens that you have, and their settings remain when you leave the operations and control panels.

Processing your actions

Press Enter to carry out the action requested on a panel. The information from the panel is sent to the queue manager for processing.

Each time you press Enter in the panels, WebSphere MQ generates one or more operator messages. If the operation was successful, you get confirmation message CSQ9022I, otherwise you get some error messages.

Displaying WebSphere MQ user messages

Press function key F10 in any panel to see the WebSphere MQ user messages.

Canceling your actions

On the initial panel, both F3 and F12 exit the operations and control panels and return you to ISPF. No information is sent to the queue manager.

On any other panel, press function keys F3 or F12 to leave the current panel **ignoring any data you have typed since last pressing Enter**. Again, no information is sent to the queue manager.

- F3 takes you straight back to the initial panel.
- F12 takes you back to the previous panel.

Getting help

Each panel has help panels associated with it. The help panels use the ISPF protocols:

- Press function key F1 on any panel to see general help (extended help) about the task.

- Press function key F1 with the cursor on any field to see specific help about that field.
- Press function key F5 from any field help panel to get the general help.
- Press function key F3 to return to the base panel, that is, the panel from which you pressed function key F1.
- Press function key F6 from any help panel to get help about the function keys.

If the help information carries on into a second or subsequent pages, a **More** indicator is displayed in the upper-right of the panel. Use these function keys to navigate through the help pages:

- F11 to get to the next help page (if there is one).
- F10 to get back to the previous help page (if there is one).

Using the command line

You never need to use the command line to issue the commands used by the operations and control panels because they are available from function keys. The command line is provided to allow you to enter normal ISPF commands (like PFSHOW).

The ISPF command PANELID ON displays the name of the current CSQOREXX panel.

The command line is initially displayed in the default position at the bottom of the panels, regardless of what ISPF settings you have. You can use the SETTINGS ISPF command from any of the operations and control panels to change the position of the command line. The settings are remembered for subsequent sessions with the operations and control panels.

Using the operations and control panels

Use this topic to investigate the initial control panel displayed from CSQOREXX

Figure 55 shows the panel that is displayed when you start a panel session.

```

      IBM WebSphere MQ for z/OS - Main Menu

Complete fields. Then press Enter.

Action . . . . . 1      0. List with filter  4. Manage
                        1. List or Display  5. Perform
                        2. Define like    6. Start
                        3. Alter          7. Stop
                        8. Command

Object type . . . . . CHANNEL +
Name . . . . . *
Disposition . . . . . A  Q=Qmgr, C=Copy, P=Private, G=Group,
                        S=Shared, A=All

Connect name . . . . . MQ1C - local queue manager or group
Target queue manager . . . MQ1C
                        - connected or remote queue manager for command input
Action queue manager . . . MQ1C - command scope in group
Response wait time . . . 30  5 - 999 seconds

(C) Copyright IBM Corporation 1993, 2019. All rights reserved.

Command ==>
F1=Help      F2=Split    F3=Exit      F4=Prompt    F9=SwapNext  F10=Messages
F12=Cancel

```

Figure 55. The WebSphere MQ operations and control initial panel

From this panel you can perform actions such as:

- Choose the local queue manager you want and whether you want the commands issued on that queue manager, on a remote queue manager, or on another queue manager in the same queue-sharing group as the local queue manager. Over type the queue manager name if you need to change it.
- Select the action you want to perform by typing in the appropriate number in the **Action** field.
- Specify the object type that you want to work with. Press function key F1 for help about the object types if you are not sure what they are.

- Specify the disposition of the object type that you want to work with.
- Display a list of objects of the type specified. Type in an asterisk (*) in the **Name** field and press Enter to display a list of objects (of the type specified) that have already been defined on the action queue manager. You can then select one or more objects to work with in sequence. All the actions are available from the list.

Note: You are recommended to make choices that result in a list of objects being displayed, and then work from that list. Use the **Display** action, because that is allowed for all object types.

Using the Command Facility

Use the editor to enter or amend MQSC commands to be passed to the queue manager.

From the primary panel, CSQOPRIA, select option **8 Command**, to start the Command Facility.

You are presented with an edit session of a sequential file, *prefix.CSQUTIL.COMMANDS*, used as input to the CSQUTIL COMMAND function; see  Issuing commands to IBM WebSphere MQ (COMMAND) (*WebSphere MQ V7.1 Reference*).

You do not need to prefix commands with the command prefix string (CPF).

You can continue MQSC commands on subsequent lines by terminating the current line with the continuation characters + or -. Alternatively, use line edit mode to provide long MQSC commands or the values of long attribute values within the command.

line edit

To use line edit, move the cursor to the appropriate line in the edit panel and use **F4** to display a single line in a scrollable panel. A single line can be up to 32 760 bytes of data.

To leave line edit:

- **F3 exit** saves changes made to the line and exits
- **F12 cancel** returns to the edit panel discarding changes made to the line.

To discard changes made in the edit session, use **F12 cancel** to terminate the edit session leaving the contents of the file unchanged. Commands are not executed.

Executing commands

When you have finished entering MQSC commands, terminate the edit session with **F3 exit** to save the contents of the file and invoke CSQUTIL to pass the commands to the queue manager. The output from command processing is held in file *prefix.CSQUTIL.OUTPUT*. An edit session opens automatically on this file so that you can view the responses. Press **F3 exit** to exit this session and return to the main menu.

Working with WebSphere MQ objects

Many of the tasks described in this documentation involve manipulating WebSphere MQ objects. The object types are queue managers, queues, process definitions, namelists, channels, client connection channels, listeners, services, and authentication information objects.

- Defining simple queue objects
- Defining other types of objects
- Working with object definitions
- Working with namelists

Defining simple queue objects

To define a new object, use an existing definition as the basis for it. You can do this in one of three ways:

- By selecting an object that is a member of a list displayed as a result of options selected on the initial panel. You then enter action type 2 (**Define like**) in the action field next to the selected object. Your new object has the attributes of the selected object, except the disposition. You can then change any attributes in your new object as you require.
- On the initial panel, select the **Define like** action type, enter the type of object that you are defining in the **Object type** field, and enter the name of a specific existing object in the **Name** field. Your new object has the same attributes as the object you named in the **Name** field, except the disposition. You can then change any attributes in your new object definition as you require.
- By selecting the **Define like** action type, specifying an object type and then leaving the **Name** field blank. You can then define your new object and it has the default attributes defined for your installation. You can then change any attributes in your new object definition as you require.

Note: You do not enter the name of the object you are defining on the initial panel, but on the **Define** panel you are presented with.

The following example demonstrates how to define a local queue using an existing queue as a template.

Defining a local queue

To define a local queue object from the operations and control panels, use an existing queue definition as the basis for your new definition. There are several panels to complete. When you have completed all the panels and you are satisfied that the attributes are correct, press Enter to send your definition to the queue manager, which then creates the actual queue.

Use the **Define like** action either on the initial panel or against an object entry in a list displayed as a result of options selected on the initial panel.

For example, starting from the initial panel, complete these fields:

Action	2 (Define like)
Object type	QLOCAL
Name	QUEUE.YOU.LIKE. This is the name of the queue that provides the attributes for your new queue.

Press Enter to display the **Define a Local Queue** panel. The queue name field is blank so that you can supply the name for the new queue. The description is that of the queue upon which you are basing this new definition. Over type this field with your own description for the new queue.

The values in the other fields are those of the queue upon which you are basing this new queue, except the disposition. You can over type these fields as you require. For example, type Y in the **Put enabled** field (if it is not already Y) if suitably authorized applications can put messages on this queue.

You get field help by moving the cursor into a field and pressing function key F1. Field help provides information about the values that can be used for each attribute.

When you have completed the first panel, press function key F8 to display the second panel.

Hints:

1. Do *not* press Enter at this stage, otherwise the queue will be created before you have a chance to complete the remaining fields. (If you do press Enter prematurely, do not worry; you can always alter your definition later on.)
2. Do not press function keys F3 or F12, or the data you typed will be lost.

Press function key F8 repeatedly to see and complete the remaining panels, including the trigger definition, event control, and backout reporting panels.

When your local queue definition is complete

When your definition is complete, press Enter to send the information to the queue manager for processing. The queue manager creates the queue according to the definition you have supplied. If you do not want the queue to be created, press function key F3 to exit and cancel the definition.

Defining other types of objects

To define other types of object, use an existing definition as the base for your new definition as explained in Defining a local queue.

Use the **Define like** action either on the initial panel or against an object entry in a list displayed as a result of options selected on the initial panel.

For example, starting from the initial panel, complete these fields:

Action	2 (Define like)
Object type	QALIAS, NAMELIST, PROCESS, CHANNEL, and other resource objects.
Name	Leave blank or enter the name of an existing object of the same type.

Press Enter to display the corresponding DEFINE panels. Complete the fields as required and then press Enter again to send the information to the queue manager.

Like defining a local queue, defining another type of object generally requires several panels to be completed. Defining a namelist requires some additional work, as described in “Working with namelists” on page 272.

Working with object definitions

When an object has been defined, you can specify an action in the **Action** field, to alter, display, or manage it.

In each case, you can either:

- Select the object you want to work with from a list displayed as a result of options selected on the initial panel. For example, having entered 1 in the **Action** field to display objects, Queue in the **Object type** field, and * in the **Name** field, you are presented with a list of all queues defined in the system. You can then select from this list the queue with which you need to work.
- Start from the initial panel, where you specify the object you are working with by completing the **Object type** and **Name** fields.

Altering an object definition

To alter an object definition, specify action 3 and press Enter to see the ALTER panels. These panels are very similar to the DEFINE panels. You can alter the values you want. When your changes are complete, press Enter to send the information to the queue manager.

Displaying an object definition

If you want to see the details of an object without being able to change them, specify action 1 and press Enter to see the DISPLAY panels. Again, these panels are similar to the DEFINE panels except that you cannot change any of the fields. Change the object name to display details of another object.

Deleting an object

To delete an object, specify action 4 (Manage) and the **Delete** action is one of the actions presented on the resulting menu. Select the **Delete** action.

You are asked to confirm your request. If you press function key F3 or F12, the request is canceled. If you press Enter, the request is confirmed and passed to the queue manager. The object you specified is then deleted.

Note: You cannot delete most types of channel object unless the channel initiator is started.

Working with namelists

When working with namelists, proceed as you would for other objects.

For the actions DEFINE LIKE or ALTER, press function key F11 to add names to the list or to change the names in the list. This involves working with the ISPF editor and all the normal ISPF edit commands are available. Enter each name in the namelist on a separate line.

When you use the ISPF editor in this way, the function key settings are the normal ISPF settings, and **not** those used by the other operations and control panels.

If you need to specify lowercase names in the list, specify CAPS(OFF) on the editor panel command line. When you do this, all the namelists that you edit in the future are in lowercase until you specify CAPS(ON).

When you have finished editing the namelist, press function key F3 to end the ISPF edit session. Then press Enter to send the changes to the queue manager.

Attention: If you do not press Enter at this stage but press function key F3 instead, you lose any updates that you have typed in.

Writing programs to administer WebSphere MQ

You can write your own application programs to administer a queue manager. Use this topic to understand the requirements for writing your own administration programs.

Start of General-use programming interface information

This topic contains hints and guidance to enable you to issue WebSphere MQ commands from a WebSphere MQ application program.

It contains these sections:

- “Preparing queues for administration programs” on page 273
- “Using the command server” on page 274
- “Retrieving replies to your commands” on page 276
- “Interpreting the reply messages from the command server” on page 278
- Using DISPLAY commands
- “Examples of commands and their replies” on page 279
- “If you do not receive a reply” on page 279
- “Passing commands using MGCRC” on page 279

Note: In this topic, the MQI calls are described using C-language notation. For typical invocations of the calls in the COBOL, PL/I, and assembler languages, see  MQI function calls (*WebSphere MQ V7.1 Reference*).

Understanding how it all works

In outline, the procedure for issuing commands from an application program is as follows:

1. Build a WebSphere MQ command into a type of WebSphere MQ message called a *request message*. The command can be in MQSC or PCF format.

2. Send (use **MQPUT**) this message to a special queue called the system-command input queue. The WebSphere MQ command processor runs the command.
3. Retrieve (use **MQGET**) the results of the command as *reply messages* on the reply-to queue. These messages contain the user messages that you need to determine whether your command was successful and, if it was, what the results were.

Then it is up to your application program to process the results.

Preparing queues for administration programs

Administration programs require a number of predefined queues for system command input and receiving responses.

This section applies to commands in the MQSC format. For the equivalent in PCF, see “Using Programmable Command Formats” on page 7.

Before you can issue any **MQPUT** or **MQGET** calls, you must first define, and then open, the queues you are going to use.

Defining the system-command input queue

The system-command input queue is a local queue called `SYSTEM.COMMAND.INPUT`. The supplied `CSQINP2` initialization data set, `thlqual.SCSQPROC(CSQ4INSG)`, contains a default definition for the system-command input queue. For compatibility with WebSphere MQ on other platforms, an alias of this queue, called `SYSTEM.ADMIN.COMMAND.QUEUE` is also supplied.

See  Sample definitions supplied with WebSphere MQ (*WebSphere MQ V7.1 Product Overview Guide*) for more information.

Defining a reply-to queue

You must define a reply-to queue to receive reply messages from the WebSphere MQ command processor. It can be any queue with attributes that allow reply messages to be put on it. However, for normal operation, specify these attributes:

- `USAGE(NORMAL)`
- `NOTRIGGER` (unless your application uses triggering)

Avoid using persistent messages for commands, but if you choose to do so, the reply-to queue must not be a temporary dynamic queue.

The supplied `CSQINP2` initialization data set, `thlqual.SCSQPROC(CSQ4INSG)`, contains a definition for a model queue called `SYSTEM.COMMAND.REPLY.MODEL`. You can use this model to create a dynamic reply-to queue.

Note: Replies generated by the command processor can be up to 15 000 bytes in length.

If you use a permanent dynamic queue as a reply-to queue, your application should allow time for all `PUT` and `GET` operations to complete before attempting to delete the queue, otherwise `MQRC2055 (MQRC_Q_NOT_EMPTY)` can be returned. If this occurs, try the queue deletion again after a few seconds.

Opening the system-command input queue

Before you can open the system-command input queue, your application program must be connected to your queue manager. Use the MQI call **MQCONN** or **MQCONNX** to do this.

Then use the MQI call **MQOPEN** to open the system-command input queue. To use this call:

1. Set the *Options* parameter to `MQOO_OUTPUT`
2. Set the `MQOD` object descriptor fields as follows:

```
ObjectType
    MQOT_Q (the object is a queue)
```

ObjectName

SYSTEM.COMMAND.INPUT

ObjectQMgrName

If you want to send your request messages to your local queue manager, leave this field blank. This means that your commands are processed locally.

If you want your WebSphere MQ commands to be processed on a remote queue manager, put its name here. You must also have the correct queues and links set up,

as described in  Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*).

Opening a reply-to queue

To retrieve the replies from a WebSphere MQ command, you must open a reply-to queue. One way of doing this is to specify the model queue, SYSTEM.COMMAND.REPLY.MODEL in an **MQOPEN** call, to create a permanent dynamic queue as the reply-to queue. To use this call:

1. Set the *Options* parameter to MQOO_INPUT_SHARED
2. Set the MQOD object descriptor fields as follows:

ObjectType

MQOT_Q (the object is a queue)

ObjectName

The name of the reply-to queue. If the queue name you specify is the name of a model queue object, the queue manager creates a dynamic queue.

ObjectQMgrName

To receive replies on your local queue manager, leave this field blank.

DynamicQName

Specify the name of the dynamic queue to be created.

Using the command server

The command server is a WebSphere MQ component that works with the command processor component. You can send formatted messages to the command server which interprets the messages, runs the administration requests, and sends responses back to your administration application.

The command server reads request messages from the system-command input queue, verifies them, and passes the valid ones as commands to the command processor. The command processor processes the commands and puts any replies as reply messages on to the reply-to queue that you specify. The first reply message contains the user message CSQN205I. See “Interpreting the reply messages from the command server” on page 278 for more information. The command server also processes channel initiator and queue-sharing group commands, wherever they are issued from.

Identifying the queue manager that processes your commands

The queue manager that processes the commands you issue from an administration program is the queue manager that owns the system-command input queue that the message is put onto.

Starting the command server

Normally, the command server is started automatically when the queue manager is started. It becomes available as soon as the message CSQ9022I 'START QMGR' NORMAL COMPLETION is returned from the START QMGR command. The command server is stopped when all the connected tasks have been disconnected during the system termination phase.

You can control the command server yourself using the START CMDSERV and STOP CMDSERV commands. To prevent the command server starting automatically when WebSphere MQ is

restarted, you can add a STOP CMDSERV command to your CSQINP1 or CSQINP2 initialization data sets. However, this is not recommended as it prevents any channel initiator or queue-sharing group commands being processed.

The STOP CMDSERV command stops the command server as soon as it has finished processing the current message, or immediately if no messages are being processed.

If the command server has been stopped by a STOP CMDSERV command in the program, no other commands from the program can be processed. To restart the command server, you must issue a START CMDSERV command from the z/OS console.

If you stop and restart the command server while the queue manager is running, all the messages that are on the system-command input queue when the command server stops are processed when the command server is restarted. However, if you stop and restart the queue manager after the command server is stopped, only the persistent messages on the system-command input queue are processed when the command server is restarted. All nonpersistent messages on the system-command input queue are lost.

Sending commands to the command server

For each command, you build a message containing the command, then put it onto the system-command input queue.

Building a message that includes WebSphere MQ commands

You can incorporate WebSphere MQ commands in an application program by building request messages that include the required commands. For each such command you:

1. Create a buffer containing a character string representing the command.
2. Issue an MQPUT call specifying the buffer name in the *buffer* parameter of the call.

The simplest way to do this in C is to define a buffer using 'char'. For example:

```
char message_buffer[ ] = "ALTER QLOCAL(SALES) PUT(ENABLED)";
```

When you build a command, use a null-terminated character string. Do not specify a command prefix string (CPF) at the start of a command defined in this way. This means that you do not have to alter your command scripts if you want to run them on another queue manager. However, you must take into account that a CPF is included in any response messages that are put onto the reply-to queue.

The command server folds all lowercase characters to uppercase unless they are inside quotation marks.

Commands can be any length up to a maximum 32 762 characters.

Putting messages on the system-command input queue

Use the MQPUT call to put request messages containing commands on the system-command input queue. In this call you specify the name of the reply-to queue that you have already opened.

To use the MQPUT call:

1. Set these MQPUT parameters:

Hconn The connection handle returned by the MQCONN or MQCONNX call.

Hobj The object handle returned by the MQOPEN call for the system-command input queue.

BufferLength
The length of the formatted command.

Buffer The name of the buffer containing the command.

2. Set these MQMD fields:

MsgType

MQMT_REQUEST

Format MQFMT_STRING or MQFMT_NONE

If you are not using the same code page as the queue manager, set *CodedCharSetId* as appropriate and set MQFMT_STRING, so that the command server can convert the message. Do not set MQFMT_ADMIN, as that causes your command to be interpreted as PCF.

ReplyToQ

Name of your reply-to queue.

ReplyToQMGr

If you want replies sent to your local queue manager, leave this field blank. If you want your WebSphere MQ commands to be sent to a remote queue manager, put its name here. You must also have the correct queues and links set up, as described in



Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*).

3. Set any other MQMD fields, as required. You should normally use nonpersistent messages for commands.
4. Set any *PutMsgOpts* options, as required.
If you specify MQPMO_SYNCPOINT (the default), you must follow the **MQPUT** call with a syncpoint call.

Using MQPUT1 and the system-command input queue

If you want to put just one message on the system-command input queue, you can use the **MQPUT1** call. This call combines the functions of an **MQOPEN**, followed by an **MQPUT** of one message, followed by an **MQCLOSE**, all in one call. If you use this call, modify the parameters

accordingly. See  Putting one message on a queue using the MQPUT1 call (*WebSphere MQ V7.1 Programming Guide*) for details.

Retrieving replies to your commands

The command server sends a response to a reply queue for each request message it receives. Any administration application must receive, and handle the reply messages.

When the command processor processes your commands, any reply messages are put onto the reply-to queue specified in the **MQPUT** call. The command server sends the reply messages with the same persistence as the command message it received.

Waiting for a reply

Use the **MQGET** call to retrieve a reply from your request message. One request message can produce several reply messages. For details, see “Interpreting the reply messages from the command server” on page 278.

You can specify a time interval that an **MQGET** call waits for a reply message to be generated. If you do not get a reply, use the checklist beginning in topic “If you do not receive a reply” on page 279.

To use the **MQGET** call:

1. Set these parameters:

Hconn The connection handle returned by the **MQCONN** or **MQCONNX** call.

Hobj The object handle returned by the **MQOPEN** call for the reply-to queue.

Buffer The name of the area to receive the reply.

BufferLength

The length of the buffer to receive the reply. This must be a minimum of 80 bytes.

2. To ensure that you only get the responses from the command that you issued, you must specify the appropriate *MsgId* and *CorrelId* fields. These depend on the report options, MQMD_REPORT, you specified in the MQPUT call:

MQRO_NONE

Binary zero, '00...00' (24 nulls).

MQRO_NEW_MSG_ID

Binary zero, '00...00' (24 nulls).

This is the default if none of these options has been specified.

MQRO_PASS_MSG_ID

The *MsgId* from the MQPUT.

MQRO_NONE

The *MsgId* from the MQPUT call.

MQRO_COPY_MSG_ID_TO_CORREL_ID

The *MsgId* from the MQPUT call.

This is the default if none of these options has been specified.

MQRO_PASS_CORREL_ID

The *CorrelId* from the MQPUT call.

For more details on report options, see  Report options and message flags (*WebSphere MQ V7.1 Reference*).

3. Set the following *GetMsgOpts* fields:

Options

MQGMO_WAIT

If you are not using the same code page as the queue manager, set MQGMO_CONVERT, and set *CodedCharSetId* as appropriate in the MQMD.

WaitInterval

For replies from the local queue manager, try 5 seconds. Coded in milliseconds, this becomes 5 000. For replies from a remote queue manager, and channel control and status commands, try 30 seconds. Coded in milliseconds, this becomes 30 000.

Discarded messages

If the command server finds that a request message is not valid, it discards this message and writes the message CSQN205I to the named reply-to queue. If there is no reply-to queue, the CSQN205I message is put onto the dead-letter queue. The return code in this message shows why the original request message was not valid:

00D5020F	It is not of type MQMT_REQUEST.
00D50210	It has zero length.
00D50212	It is longer than 32 762 bytes.
00D50211	It contains all blanks.
00D5483E	It needed converting, but <i>Format</i> was not MQFMT_STRING.

Other

See  Command server codes (X'D5') (*WebSphere MQ V7.1 Reference*) manual.

The command server reply message descriptor

For any reply message, the following MQMD message descriptor fields are set:

<i>MsgType</i>	MQMT_REPLY
<i>Feedback</i>	MQFB_NONE
<i>Encoding</i>	MQENC_NATIVE
<i>Priority</i>	As for the MQMD in the message you issued.
<i>Persistence</i>	As for the MQMD in the message you issued.
<i>CorrelId</i>	Depends on the MQPUT report options.
<i>ReplyToQ</i>	None.

The command server sets the *Options* field of the MQPMO structure to MQPMO_NO_SYNCPOINT. This means that you can retrieve the replies as they are created, rather than as a group at the next syncpoint.

Interpreting the reply messages from the command server

Each request message correctly processed by WebSphere MQ produces at least two reply messages. Each reply message contains a single WebSphere MQ user message.

The length of a reply depends on the command that was issued. The longest reply you can get is from a DISPLAY NAMELIST, and that can be up to 13 000 bytes long.

The first user message, CSQN205I, always contains:

- A count of the replies (in decimal), which you can use as a counter in a loop to get the rest of the replies. The count includes this first message.
- The return code from the command preprocessor.
- A reason code, which is the return code from the command processor.

This message does not contain a CPF.

For example:

```
CSQN205I  COUNT=      4, RETURN=0000000C, REASON=00000008
```

The COUNT field is 8 bytes long and is right-justified. It always starts at position 18, that is, immediately after 'COUNT='. The RETURN field is 8 bytes long in character hexadecimal and is immediately after 'RETURN=' at position 35. The REASON field is 8 bytes long in character hexadecimal and is immediately after 'REASON=' at position 52.

If the RETURN= value is 00000000 and the REASON= value is 00000004, the set of reply messages is incomplete. After retrieving the replies indicated by the CSQN205I message, issue a further **MQGET** call to wait for a further set of replies. The first message in the next set of replies is again CSQN205I, indicating how many replies there are, and whether there are still more to come.

See the  *z/OS Messages and Codes (WebSphere MQ V7.1 Reference)* documentation for more details about the individual messages.

If you are using a non-English language feature, the text and layout of the replies are different from those shown here. However, the size and position of the count and return codes in message CSQN205I are the same.

If you do not receive a reply

There are a series of steps you can take if you do not receive a response to request to the command server.

If you do not receive a reply to your request message, work through this checklist:

- Is the command server running?
- Is the *WaitInterval* long enough?
- Are the system-command input and reply-to queues correctly defined?
- Were the **MQOPEN** calls to these queues successful?
- Are both the system-command input and reply-to queues enabled for **MQPUT** and **MQGET** calls?
- Have you considered increasing the **MAXDEPTH** and **MAXMSGL** attributes of your queues?
- Are you are using the *CorrelId* and *MsgId* fields correctly?
- Is the queue manager still running?
- Was the command built correctly?
- Are all your remote links defined and operating correctly?
- Were the **MQPUT** calls correctly defined?
- Has the reply-to queue been defined as a temporary dynamic queue instead of a permanent dynamic queue? (If the request message is persistent, you must use a permanent dynamic queue for the reply.)

When the command server generates replies but cannot write them to the reply-to queue that you specify, it writes them to the dead-letter queue.

Passing commands using MGCRE

With appropriate authorization, an application program can make requests to multiple queue managers using a z/OS service routine.

If you have the correct authorization, you can pass WebSphere MQ commands from your program to multiple queue managers by the MGCRE (SVC 34) z/OS service. The value of the CPF identifies the particular queue manager to which the command is directed. For information about CPFs, see “User IDs for command security and command resource security” on page 568 and “Issuing queue manager commands” on page 248.

If you use MGCRE, you can use a Command and Response Token (CART) to get the direct responses to the command.

Examples of commands and their replies

Use this topic as a series of examples of commands to the command server and the responses from the command server.

Here are some examples of commands that could be built into WebSphere MQ messages, and the user messages that are the replies. Unless otherwise stated, each line of the reply is a separate message.

- Messages from a **DEFINE** command
- Messages from a **DELETE** command
- Messages from **DISPLAY** commands
- Messages from commands with **CMDSCOPE**
- Messages from commands that generate commands with **CMDSCOPE**

Messages from a **DEFINE** command

The following command:

```
DEFINE QLOCAL(Q1)
```

produces these messages:

```
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000000  
CSQ9022I +CSQ1 CSQMMSGP ' DEFINE QLOCAL' NORMAL COMPLETION
```

These reply messages are produced on normal completion.

Messages from a DELETE command

The following command:

```
DELETE QLOCAL(Q2)
```

produces these messages:

```
CSQN205I  COUNT=      4, RETURN=00000000, REASON=00000008  
CSQM125I +CSQ1 CSQMUQLC QLOCAL (Q2) QSGDISP(QMGR) WAS NOT FOUND  
CSQM090E +CSQ1 CSQMUQLC FAILURE REASON CODE X'00D44002'  
CSQ9023E +CSQ1 CSQMUQLC ' DELETE QLOCAL' ABNORMAL COMPLETION
```

These messages indicate that a local queue called Q2 does not exist.

Messages from DISPLAY commands

The following examples show the replies from some DISPLAY commands.

Finding out the name of the dead-letter queue

If you want to find out the name of the dead-letter queue for a queue manager, issue this command from an application program:

```
DISPLAY QMGR DEADQ
```

The following three user messages are returned, from which you can extract the required name:

```
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000000  
CSQM409I +CSQ1 QMNAME(CSQ1) DEADQ(SYSTEM.DEAD.QUEUE      )  
CSQ9022I +CSQ1 CSQMDRTS ' DISPLAY QMGR' NORMAL COMPLETION
```

Messages from the DISPLAY QUEUE command

The following examples show how the results from a command depend on the attributes specified in that command.

Example 1

You define a local queue using the command:

```
DEFINE QLOCAL(Q1) DESCR('A sample queue') GET(ENABLED) SHARE
```

If you issue the following command from an application program:

```
DISPLAY QUEUE(Q1) SHARE GET DESCR
```

these three user messages are returned:

```
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000000
CSQM401I +CSQ1 QUEUE(Q1                                ) TYPE(
QLOCAL  ) QSGDISP(QMGR                                )
          DESCR(A sample queue
          ) SHARE  GET(ENABLED  )
CSQ9022I +CSQ1 CSQMMSG ' DISPLAY QUEUE' NORMAL COMPLETION
```

Note: The second message, CSQM401I, is shown here occupying four lines.

Example 2

Two queues have names beginning with the letter A:

- A1 is a local queue with its PUT attribute set to DISABLED.
- A2 is a remote queue with its PUT attribute set to ENABLED.

If you issue the following command from an application program:

```
DISPLAY QUEUE(A*) PUT
```

these four user messages are returned:

```
CSQN205I  COUNT=      4, RETURN=00000000, REASON=00000000
CSQM401I +CSQ1 QUEUE(A1                                ) TYPE(
QLOCAL  ) QSGDISP(QMGR                                )
          PUT(DISABLED )
CSQM406I +CSQ1 QUEUE(A2                                ) TYPE(
QREMOTE ) PUT(ENABLED  )
CSQ9022I +CSQ1 CSQMMSG ' DISPLAY QUEUE' NORMAL COMPLETION
```

Note: The second and third messages, CSQM401I and CSQM406I, are shown here occupying three and two lines.

Messages from the DISPLAY NAMELIST command

You define a namelist using the command:

```
DEFINE NAMELIST(N1) NAMES(Q1,SAMPLE_QUEUE)
```

If you issue the following command from an application program:

```
DISPLAY NAMELIST(N1) NAMES NAMCOUNT
```

the following three user messages are returned:

```
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000000
CSQM407I +CSQ1 NAMELIST(N1                               ) QS
GDISP(QMGR      ) NAMCOUNT(      2) NAMES(Q1
,SAMPLE_QUEUE
)
CSQ9022I +CSQ1 CSQMDMSG ' DISPLAY NAMELIST' NORMAL COMPLETION
```

Note: The second message, CSQM407I, is shown here occupying three lines.

Messages from commands with CMDSCOPE

The following examples show the replies from commands that have been entered with the CMDSCOPE attribute.

Messages from the ALTER PROCESS command

The following command:

```
ALT PRO(V4) CMDSCOPE(*)
```

produces the following messages:

```
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000004
CSQN137I !MQ25 'ALT PRO' command accepted for CMDSCOPE(*), sent to 2
CSQN205I  COUNT=      5, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'ALT PRO' command responses from MQ26
CSQM125I !MQ26 CSQMMSGP PROCESS(V4) QSGDISP(QMGR) WAS NOT FOUND
CSQM090E !MQ26 CSQMMSGP FAILURE REASON CODE X'00D44002'
CSQ9023E !MQ26 CSQMMSGP ' ALT PRO' ABNORMAL COMPLETION
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'ALT PRO' command responses from MQ25
CSQ9022I !MQ25 CSQMMSGP ' ALT PRO' NORMAL COMPLETION
CSQN205I  COUNT=      2, RETURN=0000000C, REASON=00000008
CSQN123E !MQ25 'ALT PRO' command for CMDSCOPE(*) abnormal completion
```

These messages tell you that the command was entered on queue manager MQ25 and sent to two queue managers (MQ25 and MQ26). The command was successful on MQ25 but the process definition did not exist on MQ26, so the command failed on that queue manager.

Messages from the DISPLAY PROCESS command

The following command:

```
DIS PRO(V*) CMDSCOPE(*)
```

produces the following messages:

```
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000004
CSQN137I !MQ25 'DIS PRO' command accepted for CMDSCOPE(*), sent to 2
CSQN205I  COUNT=      5, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DIS PRO' command responses from MQ26
CSQM408I !MQ26 PROCESS(V2) QSGDISP(COPY)
CSQM408I !MQ26 PROCESS(V3) QSGDISP(QMGR)
CSQ9022I !MQ26 CSQMDRTS ' DIS PROCESS' NORMAL COMPLETION
CSQN205I  COUNT=      7, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DIS PRO' command responses from MQ25
CSQM408I !MQ25 PROCESS(V2) QSGDISP(COPY)
CSQM408I !MQ25 PROCESS(V2) QSGDISP(GROUP)
CSQM408I !MQ25 PROCESS(V3) QSGDISP(QMGR)
CSQM408I !MQ25 PROCESS(V4) QSGDISP(QMGR)
CSQ9022I !MQ25 CSQMDRTS ' DIS PROCESS' NORMAL COMPLETION
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000000
CSQN122I !MQ25 'DIS PRO' command for CMDSCOPE(*) normal completion
```

These messages tell you that the command was entered on queue manager MQ25 and sent to two queue managers (MQ25 and MQ26). Information is displayed about all the processes on each queue manager with names starting with the letter V.

Messages from the DISPLAY CHSTATUS command

The following command:

```
DIS CHS(VT) CMDSCOPE(*)
```

produces the following messages:

```
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000004
CSQN137I !MQ25 'DIS CHS' command accepted for CMDSCOPE(*), sent to 2
CSQN205I  COUNT=      4, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DIS CHS' command responses from MQ25
CSQM422I !MQ25 CHSTATUS(VT) CHLDISP(PRIVATE) CONNAME( ) CURRENT STATUS(STOPPED)
CSQ9022I !MQ25 CSQXDRTS ' DIS CHS' NORMAL COMPLETION
CSQN205I  COUNT=      4, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DIS CHS' command responses from MQ26
CSQM422I !MQ26 CHSTATUS(VT) CHLDISP(PRIVATE) CONNAME( ) CURRENT STATUS(STOPPED)
CSQ9022I !MQ26 CSQXDRTS ' DIS CHS' NORMAL COMPLETION
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000000
CSQN122I !MQ25 'DIS CHS' command for CMDSCOPE(*) normal completion
```

These messages tell you that the command was entered on queue manager MQ25 and sent to two queue managers (MQ25 and MQ26). Information is displayed about channel status on each queue manager.

Messages from the STOP CHANNEL command

The following command:

```
STOP CHL(VT) CMDSCOPE(*)
```

produces these messages:

```
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000004
CSQN137I !MQ25 'STOP CHL' command accepted for CMDSCOPE(*), sent to 2
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'STOP CHL' command responses from MQ25
CSQM134I !MQ25 CSQMTCHL STOP CHL(VT) COMMAND ACCEPTED
SQN205I   COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'STOP CHL' command responses from MQ26
CSQM134I !MQ26 CSQMTCHL STOP CHL(VT) COMMAND ACCEPTED
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'STOP CHL' command responses from MQ26
CSQ9022I !MQ26 CSQXCRPS ' STOP CHL' NORMAL COMPLETION
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'STOP CHL' command responses from MQ25
CSQ9022I !MQ25 CSQXCRPS ' STOP CHL' NORMAL COMPLETION
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000000
CSQN122I !MQ25 'STOP CHL' command for CMDSCOPE(*) normal completion
```

These messages tell you that the command was entered on queue manager MQ25 and sent to two queue managers (MQ25 and MQ26). Channel VT was stopped on each queue manager.

Messages from commands that generate commands with CMDSCOPE

The following command:

```
DEF PRO(V2) QSGDISP(GROUP)
```

produces these messages:

```
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQM122I !MQ25 CSQMMSGP ' DEF PRO' COMPLETED FOR QSGDISP(GROUP)
CSQN138I !MQ25 'DEFINE PRO' command generated for CMDSCOPE(*), sent to 2
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DEFINE PRO' command responses from MQ25
CSQ9022I !MQ25 CSQMMSGP ' DEFINE PROCESS' NORMAL COMPLETION
CSQN205I  COUNT=      3, RETURN=00000000, REASON=00000004
CSQN121I !MQ25 'DEFINE PRO' command responses from MQ26
CSQ9022I !MQ26 CSQMMSGP ' DEFINE PROCESS' NORMAL COMPLETION
CSQN205I  COUNT=      2, RETURN=00000000, REASON=00000000
CSQN122I !MQ25 'DEFINE PRO' command for CMDSCOPE(*) normal completion
```

These messages tell you that the command was entered on queue manager MQ25. When the object was created on the shared repository, another command was generated and sent to all the active queue managers in the queue-sharing group (MQ25 and MQ26).

Managing WebSphere MQ resources on z/OS

Use the links in this topic to find out how to manage the resources used by IBM WebSphere MQ for z/OS, for example, managing log files, data sets, page sets, buffer pools, and coupling facility structures.

Use the following links for details of the different administrative tasks you might have to complete whilst using IBM WebSphere MQ for z/OS:

- “Managing the logs”
- “Managing the bootstrap data set (BSDS)” on page 292
- “Managing page sets” on page 300
- “How to back up and recover page sets” on page 307
- “How to back up and restore queues using CSQUTIL” on page 311
- “Managing buffer pools” on page 311
- “Managing queue-sharing groups and shared queues” on page 313

Related concepts:

 [WebSphere MQ for z/OS concepts \(WebSphere MQ V7.1 Product Overview Guide\)](#)

“Administering IBM WebSphere MQ for z/OS” on page 237

“Issuing commands” on page 238

“Recovery and restart” on page 321

 [Planning your IBM WebSphere MQ environment on z/OS \(WebSphere MQ V7.1 Installing Guide\)](#)

 [Configuring z/OS \(WebSphere MQ V7.1 Installing Guide\)](#)

 [Programmable command formats reference \(WebSphere MQ V7.1 Reference\)](#)

 [MQSC reference \(WebSphere MQ V7.1 Reference\)](#)

 [Using the WebSphere MQ for z/OS utilities \(WebSphere MQ V7.1 Reference\)](#)

Related reference:

“The WebSphere MQ for z/OS utilities” on page 246

Managing the logs

Use this topic to understand how to manage your WebSphere MQ log files, including the log archiving process, using log record compression, log record recovery, and printing log records.

This topic describes the tasks involved in managing the WebSphere MQ logs. It contains these sections:

- Archiving logs with the ARCHIVE LOG command
This topic also describes restarting the log archive process after a failure.
- Controlling archiving and logging
This topic describes how to control archiving and logging, how to control log compression, print log records, and recover logs.
- Discarding archive log data sets
This includes automatic and manual deleting of the archive log data sets.
- The effect of log shunting
This topic describes log shunting, which moves the log records to optimize the quantity of log data retained, and queue manager restart time.
- What to do when the used log range becomes critical
This topic describes how to reset the queue manager's logs and pagesets to prevent the log RBA wrapping from FFFFFFFFFF back to 0.

Archiving logs with the ARCHIVE LOG command:

An authorized operator can archive the current WebSphere MQ active log data sets whenever required using the ARCHIVE LOG command.

When you issue the ARCHIVE LOG command, WebSphere MQ truncates the current active log data sets, then runs an asynchronous offload process, and updates the BSDS with a record of the offload process.

The ARCHIVE LOG command has a MODE(QUIESCE) option. With this option, WebSphere MQ jobs and users are quiesced after a commit point, and the resulting point of consistency is captured in the current active log before it is offloaded.

Consider using the MODE(QUIESCE) option when planning a backup strategy for off site recovery. It creates a system-wide point of consistency, which minimizes the number of data inconsistencies when the archive log is used with the most current backup page set copy during recovery. For example:

```
ARCHIVE LOG MODE(QUIESCE)
```

If you issue the ARCHIVE LOG command without specifying a TIME parameter, the quiesce time period defaults to the value of the QUIESCE parameter of the CSQ6ARVP macro. If the time required for the ARCHIVE LOG MODE(QUIESCE) to complete is less than the time specified, the command completes successfully; otherwise, the command fails when the time period expires. You can specify the time period explicitly by using the TIME option, for example:

```
ARCHIVE LOG MODE(QUIESCE) TIME(60)
```

This command specifies a quiesce period of up to 60 seconds before ARCHIVE LOG processing occurs.

Attention: Using the TIME option when time is critical can significantly disrupt WebSphere MQ availability for all jobs and users that use WebSphere MQ resources.

By default, the command is processed asynchronously from the time you submit the command. (To process the command synchronously with other WebSphere MQ commands use the WAIT(YES) option with QUIESCE, but be aware that the z/OS console is locked from WebSphere MQ command input for the entire QUIESCE period.)

During the quiesce period:

- Jobs and users on the queue manager are allowed to go through commit processing, but are suspended if they try to update any WebSphere MQ resource after the commit.
- Jobs and users that only read data can be affected, since they might be waiting for locks held by jobs or users that were suspended.
- New tasks can start, but they cannot update data.

The output from the DISPLAY LOG command uses the message CSQV400I to indicate that a quiesce is in effect. For example:

```

CSQJ322I +CSQ1 DISPLAY LOG report ...
Parameter   Initial value      SET value
-----
INBUFF      60
OUTBUFF     4000
MAXRTU      2
MAXARCH     2
TWOACTV     YES
TWOARCH     YES
TWOBSDS     YES
OFFLOAD     YES
WRTHRS      20
DEALLCT     0
End of LOG report
CSQJ370I +CSQ1 LOG status report ...
Copy %Full DSName
  1      68 VICY.CSQ1.LOGCOPY1.DS01
  2      68 VICY.CSQ1.LOGCOPY2.DS01
Restarted at 2005-02-24 09:49:30 using RBA=00000891B000
Latest RBA=00000891CCF8
Offload task is AVAILABLE
Full logs to offload - 0 of 4
CSQV400I +CSQ1 ARCHIVE LOG QUIESCE CURRENTLY ACTIVE
CSQ9022I +CSQ1 CSQJ311E ' DISPLAY LOG' NORMAL COMPLETION

```

When all updates are quiesced, the quiesce history record in the BSDS is updated with the date and time that the active log data sets were truncated, and with the last-written RBA in the current active log data sets. WebSphere MQ truncates the current active log data sets, switches to the next available active log data sets, and issues message CSQJ311E stating that the offload process started.

If updates cannot be quiesced before the quiesce period expires, WebSphere MQ issues message CSQJ317I, and ARCHIVE LOG processing terminates. The current active log data sets are not truncated, nor switched to the next available log data sets, and the offload process is not started.

Whether the quiesce was successful or not, all suspended users and jobs are then resumed, and WebSphere MQ issues message CSQJ312I, stating that the quiesce is ended and update activity is resumed.

If ARCHIVE LOG is issued when the current active log is the last available active log data set, the command is not processed, and WebSphere MQ issues the following message:

```

CSQJ319I - csect-name CURRENT ACTIVE LOG DATA SET IS THE LAST
          AVAILABLE ACTIVE LOG DATA SET.  ARCHIVE LOG PROCESSING
          WILL BE TERMINATED

```

If ARCHIVE LOG is issued when another ARCHIVE LOG command is already in progress, the new command is not processed, and WebSphere MQ issues the following message:

CSQJ318I - ARCHIVE LOG COMMAND ALREADY IN PROGRESS

For information about the syntax of the ARCHIVE LOG command, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) documentation. For information about the messages issued during archiving, see the  z/OS Messages and Codes (*WebSphere MQ V7.1 Reference*) documentation.

Restarting the log archive process after a failure

If there is a problem during the log archive process (for example, a problem with allocation or tape mounts), the archiving of the active log might be suspended. You can cancel the archive process and restart it by using the ARCHIVE LOG CANCEL OFFLOAD command. This command cancels any offload processing currently in progress, and restarts the archive process. It starts with the oldest log data set that has not been archived, and proceeds through all active log data sets that need offloading. Any log archive operations that have been suspended are restarted.

Use this command only if you are sure that the current log archive task is no longer functioning, or if you want to restart a previous attempt that failed. This is because the command might cause an abnormal termination of the offload task, which might result in a dump.

Controlling archiving and logging:

You can control compression, printing, archiving, recovery and logging with using the CSQ6LOGP, CSQ6ARVP, and CSQ6SYSP macros.

Many aspects of archiving and logging are controlled by parameters set using the CSQ6LOGP, CSQ6ARVP and CSQ6SYSP macros of the system parameter module when the queue manager is customized. See  Task 17: Tailor your system parameter module (*WebSphere MQ V7.1 Installing Guide*) for details of these macros.

Some of these parameters can be changed while a queue manager is running using the WebSphere MQ MQSC SET LOG, SET SYSTEM and SET ARCHIVE commands. They are shown in Table 11:

Table 11. Archiving and logging parameters that can be changed while a queue manager is running

SET command	Parameters
LOG	WRTHRSR, MAXARCH, DEALLCT, MAXRTU, COMPLOG
ARCHIVE	All
SYSTEM	LOGLOAD

You can display the settings of all the parameters using the MQSC  DISPLAY LOG (*WebSphere MQ V7.1 Reference*),  DISPLAY ARCHIVE (*WebSphere MQ V7.1 Reference*) and  DISPLAY SYSTEM (*WebSphere MQ V7.1 Reference*) commands. These commands also show status information about archiving and logging.

Controlling log compression

You can enable and disable the compression of log records using either

- The SET and DISPLAY LOG commands in MQSC; see  The MQSC commands (*WebSphere MQ V7.1 Reference*)

- Invoking PCF interface. See “Introduction to Programmable Command Formats” on page 6
- Using the CSQ6LOGP macro in the system parameter module; see  Using CSQ6LOGP (*WebSphere MQ V7.1 Installing Guide*)

Printing log records

You can extract and print log records using the CSQ1LOGP utility. For instructions, see  The log print utility (CSQ1LOGP) (*WebSphere MQ V7.1 Reference*).

Recovering logs

Normally, you do not need to back up and restore the WebSphere MQ logs, especially if you are using dual logging. However, in rare circumstances, such as an I/O error on a log, you might need to recover the logs. Use Access Method Services to delete and redefine the data set, and then copy the corresponding dual log into it.

Discarding archive log data sets:

You can discard your archive log data sets and choose to discard the logs automatically or manually.

You must keep enough log data to be able to perform unit of work recovery, page set media recovery if a page set is lost, or CF structure media recovery if a CF structure is lost. Do not discard archive log data sets that might be required for recovery; if you discard these archive log data sets you might not be able to perform required recovery operations.

If you have confirmed that your archive log data sets can be discarded, you can do this in either of the following ways:

- Automatic archive log data set deletion
- Manually deleting archive log data sets

Automatic archive log data set deletion

You can use a DASD or tape management system to delete archive log data sets automatically. The retention period for WebSphere MQ archive log data sets is specified by the retention period field ARCRETN in the CSQ6ARVP installation macro (see the  Using CSQ6ARVP (*WebSphere MQ V7.1 Installing Guide*) for more information). This value is passed to the management system in the JCL parameter RETPD.

The default for the retention period specifies that archive logs are to be kept for 9999 days (the maximum). **You can change the retention period but you must ensure that you can accommodate the number of backup cycles that you have planned for.**

WebSphere MQ uses the retention period value as the value for the JCL parameter RETPD when archive log data sets are created.

The retention period set by MVS/DFP's storage management subsystem (SMS) can be overridden by this WebSphere MQ parameter. Typically, the retention period is set to the smaller value specified by either WebSphere MQ or SMS. The storage administrator and WebSphere MQ administrator must agree on a retention period value that is appropriate for WebSphere MQ.

Note: WebSphere MQ does not have an automated method to delete information about archive log data sets from the BSDS, because some tape management systems provide external manual overrides of retention periods. Therefore, information about an archive log data set can still be in the BSDS long after the data set retention period has expired and the data set has been scratched by the tape management

system. Conversely, the maximum number of archive log data sets might have been exceeded and the data from the BSDS might have been dropped before the data set has reached its expiration date.

If archive log data sets are deleted automatically, remember that the operation does not update the list of archive logs in the BSDS. You can update the BSDS with the change log inventory utility, as described in “Changing the BSDS” on page 294. The update is not essential. Recording old archive logs wastes space in the BSDS, but does no other harm.

Manually deleting archive log data sets

You must keep all the log records as far back as the lowest RBA identified in messages CSQI024I and CSQI025I. This RBA is obtained using the DISPLAY USAGE command that you issued when creating a point of recovery using Method 1: Full backup.

Read Creating a point of recovery for non-shared resources before discarding any logs.

Locate and discard archive log data sets

Having established the minimum log RBA required for recovery, you can find archive log data sets that contain only earlier log records by performing the following procedure:

1. Use the print log map utility to print the contents of the BSDS. For an example of the output, see  The print log map utility (CSQJU004) (*WebSphere MQ V7.1 Reference*).
2. Find the sections of the output titled “ARCHIVE LOG COPY n DATA SETS”. If you use dual logging, there are two sections. The columns labeled STARTRBA and ENDRBA show the range of RBAs contained in each volume. Find the volumes with ranges that include the minimum RBA you found with messages CSQI024I and CSQI025I. These are the earliest volumes you need to keep. If you are using dual-logging, there are two such volumes.

If no volumes have an appropriate range, one of the following cases applies:

- The minimum RBA has not yet been archived, and you can discard all archive log volumes.
- The list of archive log volumes in the BSDS wrapped around when the number of volumes exceeded the number allowed by the MAXARCH parameter of the CSQ6LOGP macro. If the BSDS does not register an archive log volume, that volume cannot be used for recovery. Therefore, consider adding information about existing volumes to the BSDS. For instructions, see “Changes for archive logs” on page 296.

Also consider increasing the value of MAXARCH. For information, see the  Using CSQ6LOGP (*WebSphere MQ V7.1 Installing Guide*).

3. Delete any archive log data set or volume with an ENDRBA value that is less than the STARTRBA value of the earliest volume you want to keep. If you are using dual logging, delete both such copies.

Because BSDS entries wrap around, the first few entries in the BSDS archive log section might be more recent than the entries at the bottom. Look at the combination of date and time and compare their ages. Do not assume that you can discard all entries *above* the entry for the archive log containing the minimum LOGRBA.

Delete the data sets. If the archives are on tape, erase the tapes. If they are on DASD, run a z/OS utility to delete each data set. Then, if you want the BSDS to list only existing archive volumes, use the change log inventory utility (CSQJU003) to delete entries for the discarded volumes. See “Changes for archive logs” on page 296 for an example.

The effect of log shunting:

Long running transactions can cause unit of work log records which span log data sets. WebSphere MQ handles this scenario by using log shunting, a technique which moves the log records to optimize the quantity of log data retained, and queue manager restart time.

When a unit of work is considered to be long, a representation of each log record is written further down the log. This is known as *log shunting*. It is described more fully in  Log files.

The queue manager uses these shunted log records instead of the originals after a failure, to ensure unit of work integrity. There are two benefits to this:

- the quantity of log data which must be retained for unit of work coordination is reduced
- less log data must be traversed at queue manager restart time, so the queue manager is restarted more quickly

Shunted log records do not contain sufficient information for media recovery operations.

Data held in the log is used for two distinct purposes; media recovery and unit of work coordination. If a media failure occurs which affects either a CF structure or page set, the queue manager can recover the media to the point of failure by restoring a prior copy and updating this using data contained in the log. Persistent activity performed in a unit of work is recorded on the log so that in the event of a failure, it can either be backed out or locks can be recovered on changed resources. The quantity of log data you need to retain to enable queue manager recovery is affected by these two elements.

For media recovery, you must retain sufficient log data to be able to perform media recovery from at least the most recent media copy and to be able to back out. (Your site may stipulate the ability to recover from older backups.) For unit of work integrity, you must retain the log data for your oldest in flight or indoubt units of work.

To assist you with managing the system, the queue manager detects old units of work at each log archive and reports them in messages CSQJ160 and CSQJ161. An internal task reads unit of work log information for these old units of work and rewrites it in a more succinct form to the current position in the log. Message CSQR026 indicates when this has happened. The MQSC command DISPLAY USAGE TYPE(DATASET) can also help you to manage the retention of log data. The command reports 3 pieces of recovery information:

1. how much of the log must be retained for unit of work recovery
2. how much of the log must be retained for media recovery of page sets
3. for a queue manager in a queue-sharing group, how much of the log must be retained for media recovery of CF structures

For each of these, an attempt is made to map the oldest log data required into a data set. As new units of work start and stop, we would expect (1) above to move to a more recent position in the log. If it is not moving, the long running UOW messages warn you that there is an issue. (2) relates to page set media recovery if the queue manager were to be shut down now and restarted. It does not know about when you last backed up your page sets, or which backup you might have to use if there was a page set failure. It normally moves to a more recent position in the log during checkpoint processing as changes held in the buffer pools are written to the page sets. In (3), the queue manager does know about CF structure backups taken either on this queue manager or on other queue managers in the queue sharing group. However, CF structure recovery requires a merge of log data from all queue managers in the queue-sharing group which have interacted with the CF structure since the last backup. This means that the log data is identified by a log record sequence number, (or LRSN), which is timestamp based and so applicable across the entire queue-sharing group rather than an RBA which would be different on different queue managers in the queue-sharing group. It normally moves to a more recent position in the log as BACKUP CFSTRUCT commands are performed on either this or other queue managers in the queue-sharing group.

Resetting the queue manager's log:

Use this topic to understand how to reset the queue manager's log.

You must not allow the queue manager log RBA to wrap around from a value of FFFFFFFFFF to 0 as this will lead to a queue manager outage and all persistent data will become unrecoverable.

The queue manager issues messages CSQI045I, CSQI046E, and CSQI047E to indicate that the used log range is significant and that you should plan to take action to reset logs to avoid an unplanned outage and data loss

The procedure to follow to reset the queue manager's log is as follows:

1. Resolve any unresolved units of work. The number of unresolved units of work is displayed at queue manager startup in message CSQR005I as the INDOUBT count. At each checkpoint, and at queue manager shutdown, the queue manager automatically issues the command **DISPLAY CONN(*) TYPE(CONN) ALL WHERE(UOWSTATE EQ UNRESOLVED)** to provide information about unresolved units of work.

See  How in-doubt units of recovery are resolved (*WebSphere MQ V7.1 Product Overview Guide*) for information on resolving units of recovery. The ultimate recourse is to use the **RESOLVE INDOUBT MQSC** command to manually resolve indoubt units of recovery.

2. Shutdown the queue manager cleanly.
You can use either **STOP QMGR** or **STOP QMGR MODE(FORCE)** as both these commands flush any changed pages from bufferpools to the pagesets.
3. If a queue manager is part of a queue sharing group, take CFSTRUCT backups on other queue managers for all structures in the queue sharing group. This ensures that the most recent backups are not in this queue manager's log, and that this queue manager's log is not required for CFSTRUCT recovery.
4. Define new logs and BSDS using CSQJU003 (see  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*) for more information on using the change log inventory utility).
5. Run **CSQUTIL RESETPAGE** against all the pagesets for this queue manager (see  Copying a page set and resetting the log (RESETPAGE) (*WebSphere MQ V7.1 Reference*) for more information on using this function). Note that pageset RBAs may be reset independently so multiple concurrent jobs (eg 1 per pageset) may be submitted to reduce the elapsed time for this step.
6. Restart the queue manager

Managing the bootstrap data set (BSDS)

The bootstrap data set (BSDS) is used to reference log data sets, and log records. Use this topic to understand how you can examine, change, and recover the BSDS.

For more information, see  The bootstrap data set (*WebSphere MQ V7.1 Product Overview Guide*).

This topic describes the tasks involved in managing the bootstrap data set. It contains these sections:

- "Finding out what the BSDS contains" on page 293
- "Changing the BSDS" on page 294
- "Recovering the BSDS" on page 298

Finding out what the BSDS contains:

You can use the print log map utility (CSQJU004) to examine the contents of the BSDS.

The print log map utility (CSQJU004) is a batch utility that lists the information stored in the BSDS. For instructions on running it, see  The print log map utility (CSQJU004) (*WebSphere MQ V7.1 Reference*).

The BSDS contains:

- Time stamps
- Active log data set status

Time stamps in the BSDS

The output of the print log map utility shows the time stamps, which are used to record the date and time of various system events, that are stored in the BSDS.

The following time stamps are included in the header section of the report:

SYSTEM TIMESTAMP

Reflects the date and time the BSDS was last updated. The BSDS time stamp can be updated when:

- The queue manager starts.
- The write threshold is reached during log write activities. Depending on the number of output buffers you have specified and the system activity rate, the BSDS might be updated several times a second, or might not be updated for several seconds, minutes, or even hours. For details of the write threshold, see the WRTHRSH parameter of the CSQ6LOGP macro in  Using CSQ6LOGP (*WebSphere MQ V7.1 Installing Guide*).
- WebSphere MQ drops into a single BSDS mode from its normal dual BSDS mode due to an error. This can occur when a request to get, insert, point to, update, or delete a BSDS record is unsuccessful. When this error occurs, WebSphere MQ updates the time stamp in the remaining BSDS to force a time stamp mismatch with the disabled BSDS.

UTILITY TIMESTAMP

The date and time the contents of the BSDS were altered by the change log inventory utility (CSQJU003).

The following time stamps are included in the active and archive log data sets portion of the report:

Active log date

The date the active log entry was created in the BSDS, that is, when the CSQJU003 NEWLOG was done.

Active log time

The time the active log entry was created in the BSDS, that is, when the CSQJU003 NEWLOG was done.

Archive log date

The date the archive log entry was created in the BSDS, that is, when the CSQJU003 NEWLOG was done or the archive itself was done.

Archive log time

The time the archive log entry was created in the BSDS, that is, when the CSQJU003 NEWLOG was done or the archive itself was done.

Active log data set status

The BSDS records the status of an active log data set as one of the following:

NEW The data set has been defined but never used by WebSphere MQ, or the log was truncated to a point before the data set was first used. In either case, the data set starting and ending RBA values are reset to zero.

REUSABLE

Either the data set has been defined but never used by WebSphere MQ, or the data set has been offloaded. In the print log map output, the start RBA value for the last REUSABLE data set is equal to the start RBA value of the last archive log data set.

NOT REUSABLE

The data set contains records that have not been offloaded.

STOPPED

The offload processor encountered an error while reading a record, and that record could not be obtained from the other copy of the active log.

TRUNCATED

Either:

- An I/O error occurred, and WebSphere MQ has stopped writing to this data set. The active log data set is offloaded, beginning with the starting RBA and continuing up to the last valid record segment in the truncated active log data set. The RBA of the last valid record segment is lower than the ending RBA of the active log data set. Logging is switched to the next available active log data set, and continues uninterrupted.

or

- An ARCHIVE LOG function has been called, which has truncated the active log.

The status appears in the output from the print log map utility.

Changing the BSDS:

You do not have to take special steps to keep the BSDS updated with records of logging events because WebSphere MQ does that automatically.

However, you might want to change the BSDS if you do any of the following:

- Add more active log data sets.
- Copy active log data sets to newly allocated data sets, for example, when providing larger active log allocations.
- Move log data sets to other devices.
- Recover a damaged BSDS.
- Discard outdated archive log data sets.

You can change the BSDS by running the change log inventory utility (CSQJU003). Only run this utility when the queue manager is inactive, or you might get inconsistent results. The action of the utility is controlled by statements in the SYSIN data set. This section shows several examples. For complete

instructions, see  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*).

You can copy an active log data set only when the queue manager is inactive because WebSphere MQ allocates the active log data sets as exclusive (DISP=OLD) at queue manager startup.

Changes for active logs:

Use this topic to understand how you can change the active logs using the BSDS.

You can add to, delete from, and record entries in the BSDS for active logs using the change log utility. Examples only are shown here; replace the data set names shown with the ones you want to use. For more details of the utility, see  [The change log inventory utility \(CSQJU003\) \(WebSphere MQ V7.1 Reference\)](#).

See these topics for more information:

- Adding record entries to the BSDS
- Deleting information about the active log data set from the BSDS
- Recording information about the log data set in the BSDS
- Increasing the size of the active log
- The use of CSQJUFMT

Adding record entries to the BSDS

If an active log has been flagged as “stopped”, it is not reused for logging; however, it continues to be used for reading. Use the access method services to define new active log data sets, then use the change log inventory utility to register the new data sets in the BSDS. For example, use:

```
NEWLOG DSNAMES=MQM111.LOGCOPY1.DS10,COPY1
NEWLOG DSNAMES=MQM111.LOGCOPY2.DS10,COPY2
```

If you are copying the contents of an old active log data set to the new one, you can also give the RBA range and the starting and ending time stamps on the NEWLOG function.

Deleting information about the active log data set from the BSDS

To delete information about an active log data set from the BSDS, you could use:

```
DELETE DSNAMES=MQM111.LOGCOPY1.DS99
DELETE DSNAMES=MQM111.LOGCOPY2.DS99
```

Recording information about the log data set in the BSDS

To record information about an existing active log data set in the BSDS, use:

```
NEWLOG DSNAMES=MQM111.LOGCOPY1.DS10,COPY2,STARTIME=19930212205198,
        ENDTIME=19930412205200,STARTRBA=6400,ENDRBA=94FF
```

You might need to insert a record containing this type of information in the BSDS because:

- The entry for the data set has been deleted, but is needed again.
- You are copying the contents of one active log data set to another data set.
- You are recovering the BSDS from a backup copy.

Increasing the size of the active log

This procedure must only be used when the queue manager is inactive:

1. Stop the queue manager. This step is required because WebSphere MQ allocates all active log data sets for its exclusive use when it is active.

2. Use Access Method Services ALTER with the NEWNAME option to rename your active log data sets.
3. Use Access Method Services DEFINE to define larger active log data sets.
By reusing the old data set names, you do not have to run the change log inventory utility to establish new names in the BSDSs. The old data set names and the correct RBA ranges are already in the BSDSs.
4. Use Access Method Services REPRO to copy the old (renamed) data sets into their appropriate new data sets.
5. Start the queue manager.

If all your log data sets are the same size, your system will be operationally more consistent and efficient. If the log data sets are not the same size, it is more difficult to track your system's logs, and so space can be wasted.

The use of CSQJUFMT

Do not run a CSQJUFMT format when increasing the size of an active log.

If you run CSQJUFMT (in order to provide a performance advantage the first time the queue manager writes to the new active log) you receive messages:

```
IEC070I 203-204,XS95GTLX,REPRO02,OUTPUT,B857,SPMG02, 358
IEC070I MG.W.MG4E.LOGCOPY1.DS02,MG.W.MG4E.LOGCOPY1.DS02.DATA,
IDC3302I ACTION ERROR ON MG.W.MG4E.LOGCOPY1.DS02
IDC3351I ** VSAM I/O RETURN CODE IS 28 - RPLFDBWD = X'2908001C'
IDC31467I MAXIMUM ERROR LIMIT REACHED.
```

```
IDC0005I NUMBER OF RECORDS PROCESSED WAS 0
```

In addition, if you use the Access Method Services REPRO, ensure that you define a new empty log.

If you use REPRO to copy the old (renamed) data set into its respective new data set, the default is NOREPLACE.

This means that REPRO does not replace a record that is already on the designated data set. When formatting is done on the data set, the RBA value is reset. The net result is a data set that is not empty after formatting.

Changes for archive logs:

Use this topic to understand how to change the archive logs.

You can add to, delete from, and change the password of, entries in the BSDS for archive logs. Examples only are shown here; replace the data set names shown with the ones you want to use. For more details of the utility, see  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*).

- Adding an archive log
- Deleting an archive log
- Changing the password of an archive log

Adding an archive log

When the recovery of an object depends on reading an existing archive log data set, the BSDS must contain information about that data set so that WebSphere MQ can find it. To register information about an existing archive log data set in the BSDS, use:

```
NEWLOG DSNAME=CSQARC1.ARCHLOG1.E00021.T2205197.A0000015,COPY1VOL=CSQV04,  
UNIT=TAPE,STARTRBA=3A190000,ENDRBA=3A1F0FFF,CATALOG=NO
```

Deleting an archive log

To delete an entire archive log data set on one or more volumes, use:

```
DELETE DSNAME=CSQARC1.ARCHLOG1.E00021.T2205197.A0000015,COPY1VOL=CSQV04
```

Changing the password of an archive log

If you change the password of an existing archive log data set, you must also change the information in the BSDS.

1. List the BSDS, using the print log map utility.
2. Delete the entry for the archive log data set with the changed password, using the DELETE function of the CSQJU003 utility (see topic  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*)).
3. Name the data set as for a new archive log data set. Use the NEWLOG function of the CSQJU003 utility (see topic  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*)), and give the new password, the starting and ending RBAs, and the volume serial numbers (which can be found in the print log map utility output, see  The print log map utility (CSQJU004) (*WebSphere MQ V7.1 Reference*)).

To change the password for new archive log data sets, use:

```
ARCHIVE PASSWORD=password
```

To stop placing passwords on new archive log data sets, use:

```
ARCHIVE NOPASSWD
```

Note: Only use the ARCHIVE utility function if you do not have an external security manager.

Changing the high-level qualifier (HLQ) for the logs and BSDS:

Use this topic to understand the procedure required to change the high-level qualifier (HLQ).

Before you begin

You must end the queue manager normally before copying any of the logs or data sets to the new data sets. This is to ensure that the data is consistent and no recovery is needed during restart.

About this task

This task provides information about how to change the HLQ for the logs and BSDS. To do this, follow these steps:

Procedure

1. Run the log print utility CSQJU004 to record the log data set information. This data set is needed later.
2. You can either:
 - a. run DSS backup or restore with rename on the log and BSDS data sets to be renamed, or
 - b. use AMS DEFINE and REPRO to create the HLQ data sets and copy the data from the old data sets.
3. Modify the MSTR and CHIN proc to point to the new data sets.
4. Ensure that the new libraries have been APF authorized.
5. Delete the old logs. Add the new logs to the new BSDS, but keep the other information about each log the same, including the RBA. The HLQ should be the only thing that looks different.
6. Delete the old log information in the new BSDS using CSQJU003.
7. Define the new log data sets to the new BSDS using CSQJU003, then use the NEWLOG option and specify the RBA values.
8. The new BSDS should reflect the same information that was recorded for the old logs in the old BSDS.

What to do next

Compare the CSQJU004 output for the old and new BSDS to ensure that they look EXACTLY the same (except for the HLQs) before starting the queue manager.

Note: Care must be taken when performing these operations. Incorrect actions might lead to unrecoverable situations. Check the PRINT LOG MAP UTILITY output and make sure that all the information needed for recovery or restart has been included.

Recovering the BSDS:

If WebSphere MQ is operating in dual BSDS mode and one BSDS becomes damaged, forcing WebSphere MQ into single BSDS mode, WebSphere MQ continues to operate without a problem (until the next restart).

To return the environment to dual BSDS mode:

1. Use Access Method Services to rename or delete the damaged BSDS and to define a new BSDS with the same name as the damaged BSDS. Example control statements can be found in job CSQ4BREC in thlqual.SCSQPROC.
2. Issue the WebSphere MQ command RECOVER BSDS to make a copy of the valid BSDS in the newly allocated data set and to reinstate dual BSDS mode.

If WebSphere MQ is operating in single BSDS mode and the BSDS is damaged, or if WebSphere MQ is operating in dual BSDS mode and both BSDSs are damaged, the queue manager stops and does not restart until the BSDS data sets are repaired. In this case:

1. Locate the BSDS associated with the most recent archive log data set. The data set name of the most recent archive log appears on the job log in the last occurrence of message CSQJ003I, which indicates that offload processing has been completed successfully. In preparation for the rest of this procedure, it is a good practice to keep a log of all successful archives noted by that message:
 - If archive logs are on DASD, the BSDS is allocated on any available DASD. The BSDS name is like the corresponding archive log data set name; change only the first letter of the last qualifier, from A to B, as in this example:

Archive log name

CSQ.ARCHLOG1.A0000001

BSDS copy name

CSQ.ARCHLOG1.B0000001

- If archive logs are on tape, the BSDS is the first data set of the first archive log volume. The BSDS is not repeated on later volumes.
- 2. If the most recent archive log data set has no copy of the BSDS (for example, because an error occurred when offloading it), locate an earlier copy of the BSDS from earlier offload processing.
- 3. Rename *damaged* BSDSs using the Access Method Services ALTER command with the NEWNAME option. If you want to delete a damaged BSDS, use the Access Method Services DELETE command. For each damaged BSDS, use Access Method Services to define a new BSDS as a replacement data set. Job CSQ4BREC in thlqual.SCSQPROC contains Access Method Services control statements to define a new BSDS.
- 4. Use the Access Method Services REPRO command to copy the BSDS from the archive log to one of the replacement BSDSs you defined in step 3. Do not copy any data to the second replacement BSDS, you do that in step 5 on page 300.
 - a. Print the contents of the replacement BSDS.

Use the print log map utility (CSQJU004) to print the contents of the replacement BSDS. This enables you to review the contents of the replacement BSDS before continuing your recovery work.
 - b. Update the archive log data set inventory in the replacement BSDS.

Examine the output from the print log map utility and check that the replacement BSDS does not contain a record of the archive log from which the BSDS was copied. If the replacement BSDS is an old copy, its inventory might not contain all archive log data sets that were created more recently. The BSDS inventory of the archive log data sets must be updated to reflect the current subsystem inventory.

Use the change log inventory utility (CSQJU003) NEWLOG statement to update the replacement BSDS, adding a record of the archive log from which the BSDS was copied. If the archive log data set is password-protected, use the PASSWORD option of the NEWLOG function. Also, if the archive log data set is cataloged, ensure that the CATALOG option of the NEWLOG function is properly set to CATALOG=YES. Use the NEWLOG statement to add any additional archive log data sets that were created later than the BSDS copy.
 - c. Update passwords in the replacement BSDS.

The BSDS contains passwords for the archive log data sets and for the active log data sets. To ensure that the passwords in the replacement BSDS reflect the current passwords used by your installation, use the change log inventory ARCHIVE utility function with the PASSWORD option.
 - d. Update the active log data set inventory in the replacement BSDS.

In unusual circumstances, your installation might have added, deleted, or renamed active log data sets since the BSDS was copied. In this case, the replacement BSDS does not reflect the actual number or names of the active log data sets your installation currently has in use.

If you need to delete an active log data set from the replacement BSDS log inventory, use the change log inventory utility DELETE function.

If you need to add an active log data set to the replacement BSDS log inventory, use the change log inventory utility NEWLOG function. Ensure that the RBA range is specified correctly on the NEWLOG function. If the active log data set is password-protected, use the PASSWORD option.

If you need to rename an active log data set in the replacement BSDS log inventory, use the change log inventory utility DELETE function, followed by the NEWLOG function. Ensure that the RBA range is specified correctly on the NEWLOG function. If the active log data set is password-protected, use the PASSWORD option.
 - e. Update the active log RBA ranges in the replacement BSDS.

Later, when the queue manager restarts, it compares the RBAs of the active log data sets listed in the BSDS with the RBAs found in the actual active log data sets. If the RBAs do not agree, the queue manager does not restart. The problem is magnified when an old copy of the BSDS is used.

To solve this problem, use the change log inventory utility (CSQJU003) to adjust the RBAs found in the BSDS using the RBAs in the actual active log data sets. You do this by:

- Using the print log records utility (CSQ1LOGP) to print a summary report of the active log data set. This shows the starting and ending RBAs.
- Comparing the actual RBA ranges with the RBA ranges you have just printed, when the RBAs of all active log data sets are known.

If the RBA ranges are equal for all active log data sets, you can proceed to the next recovery step without any additional work.

If the RBA ranges are not equal, adjust the values in the BSDS to reflect the actual values. For each active log data set that needs to have the RBA range adjusted, use the change log inventory utility DELETE function to delete the active log data set from the inventory in the replacement BSDS. Then use the NEWLOG function to redefine the active log data set to the BSDS. If the active log data sets are password-protected, use the PASSWORD option of the NEWLOG function.

- f. If only two active log data sets are specified for each copy of the active log, WebSphere MQ can have difficulty during queue manager restart. The problem can arise when one of the active log data sets is full and has not been offloaded, while the second active log data set is close to filling. In this case, add a new active log data set for each copy of the active log and define each new active log data set in the replacement BSDS log inventory.

Use the Access Method Services DEFINE command to define a new active log data set for each copy of the active log and use the change log inventory utility NEWLOG function to define the new active log data sets in the replacement BSDS. You do not need to specify the RBA ranges on the NEWLOG statement. However, if the active log data sets are password-protected, use the PASSWORD option of the NEWLOG function. Example control statements to accomplish this task can be found in job CSQ4LREC in thlqual.SCSQPROC.

5. Copy the updated BSDS to the second new BSDS data set. The BSDSs are now identical. Use the print log map utility (CSQJU004) to print the contents of the second replacement BSDS at this point.
6. See “Active log problems” on page 1359 for information about what to do if you have lost your current active log data set.
7. Restart the queue manager using the newly constructed BSDS. WebSphere MQ determines the current RBA and what active logs need to be archived.

Managing page sets

Use this topic to understand how to manage the page sets associated with a queue manager.

This topic describes how to add, copy, and generally manage the page sets associated with a queue manager. It contains these sections:

- How to change the high-level qualifier (HLQ) for the page sets
- “How to add a page set to a queue manager” on page 301
- “What to do when one of your page sets becomes full” on page 301
- “How to balance loads on page sets” on page 302
- How to increase the size of a page set
- “How to reduce a page set” on page 305
- “How to reintroduce a page set” on page 306
- “How to back up and recover page sets” on page 307
- “How to delete page sets” on page 311
- How to back up and restore queues using CSQUTIL

See the  Page sets (*WebSphere MQ V7.1 Product Overview Guide*) for a description of page sets, storage classes, buffers, and buffer pools, and some of the performance considerations that apply.

How to change the high-level qualifier (HLQ) for the page sets

This task gives information on how to change the HLQ for the page sets. To perform this task, do the following:

1. Define the new HLQ page sets.
2. If the size allocation is the same as the old page sets, copy the existing page set using REPRO to the empty new HLQ page sets. If you are increasing the size of the page sets, use the FORMAT function of CSQUTIL to format the destination page set. For more information, see  Formatting page sets (FORMAT) (*WebSphere MQ V7.1 Reference*).
3. Use the COPYPAGE function of CSQUTIL to copy all the messages from the source page set to the destination page set. For more information, see  Expanding a page set (COPYPAGE) (*WebSphere MQ V7.1 Reference*).
4. Change the CSQP00xx DD statement in the queue manager procedure to point to the new HLQ page sets.

Restart the queue manager and verify the changes to the page sets.

How to add a page set to a queue manager

This description assumes that you have a queue manager that is already running. You might need to add a page set if, for example, your queue manager has to cope with new applications using new queues.

To add a new page set, use the following procedure:

1. Define and format the new page set. You can use the sample JCL in thlqual.SCSQPROC(CSQ4PAGE) as a basis. For more information, see  Formatting page sets (FORMAT) (*WebSphere MQ V7.1 Reference*).
Take care not to format any page sets that are in use, unless this is what you intend. If so, use the FORCE option of the FORMAT utility function.
2. Use the DEFINE PSID command with the DSN option to associate the page set with a buffer pool.
3. Add the appropriate storage class definitions for your page set by issuing DEFINE STGCLASS commands.
4. To ensure the page set remains available when you restart your queue manager:
 - a. Add the new page set to the started task procedure for your queue manager.
 - b. Add a definition for the new page set to your CSQINP1 initialization data set.

For details of the DEFINE PSID and DEFINE STGCLASS commands, see  DEFINE PSID (*WebSphere MQ V7.1 Reference*) and  DEFINE STGCLASS (*WebSphere MQ V7.1 Reference*).

What to do when one of your page sets becomes full

You can find out about the utilization of page sets by using the WebSphere MQ command DISPLAY USAGE. For example, the command:

```
DISPLAY USAGE PSID(03)
```

displays the current state of the page set 03. This tells you how many free pages this page set has.

If you have defined secondary extents for your page sets, they are dynamically expanded each time they fill up. Eventually, all secondary extents are used, or no further disk space is available. If this happens, an application receives the return code MQRC_STORAGE_MEDIUM_FULL.

If an application receives a return code of MQRC_STORAGE_MEDIUM_FULL from an MQI call, this is a clear indication that there is not enough space left on the page set. If the problem persists or is likely to recur, you must do something to solve it.

You can approach this problem in a number of ways:

- Balance the load between page sets by moving queues from one page set to another.
- Expand the page set. See “How to increase the size of a page set” on page 304 for instructions.
- Redefine the page set so that it can expand beyond 4 GB to a maximum size of 64 GB. See “Defining a page set to be larger than 4 GB” on page 304 for instructions.

How to balance loads on page sets

Load balancing on page sets means moving the messages associated with one or more queues from one page set to another, less used, page set. Use this technique if it is not practical to expand the page set.

To identify which queues are using a page set, use the appropriate WebSphere MQ commands. For example, to find out which queues are mapped to page set 02, first, find out which storage classes map to page set 02, by using the command:

```
DISPLAY STGCLASS(*) PSID(02)
```

Then use the following command to find out which queues use which storage class:

```
DISPLAY QUEUE(*) TYPE(QLOCAL) STGCLASS
```

Moving a non-shared queue

To move queues and their messages from one page set to another, use the MQSC MOVE QLOCAL command (described in  MOVE QLOCAL (*WebSphere MQ V7.1 Reference*)). When you have identified the queue or queues that you want to move to a new page set, follow this procedure for each of these queues:

1. Ensure that the queue you want to move is not in use by any applications (that is, IPPROCS and OPPROCS values from the DISPLAY QSTATUS command are zero) and that it has no uncommitted messages (the UNCOM value from the DISPLAY QSTATUS command is NO).

Note: The only way to ensure that this state continues is to change your security settings temporarily. If you cannot do this, later stages in this procedure might fail if applications start to use the queue despite precautionary steps such as setting PUT(DISABLED). However, messages can never be lost by this procedure.

2. Prevent applications from putting messages on the queue being moved by altering the queue definition to disable MQPUTs. Change the queue definition to PUT(DISABLED).

3. Define a temporary queue with the same attributes as the queue that is being moved, using the command:

```
DEFINE QL(TEMP_QUEUE) LIKE(Queue_TO_MOVE) PUT(ENABLED) GET(ENABLED)
```

Note: If this temporary queue already exists from a previous run, delete it before doing the define.

4. Move the messages to the temporary queue using the following command:

```
MOVE QLOCAL(Queue_TO_MOVE) TOQLOCAL(TEMP_QUEUE)
```

5. Delete the queue you are moving, using the command:

```
DELETE QLOCAL(Queue_TO_MOVE)
```

6. Define a new storage class that maps to the required page set, for example:

```
DEFINE STGCLASS(NEW) PSID(nn)
```

Add the new storage class definition to the CSQINP2 data sets ready for the next queue manager restart.

7. Redefine the queue that you are moving, by changing the storage class attribute:

```
DEFINE QL(Queue_TO_MOVE) LIKE(TEMP_QUEUE) STGCLASS(NEW)
```

When the queue is redefined, it is based on the temporary queue created in step 3.

8. Move the messages back to the new queue, using the command:

```
MOVE QLOCAL(TEMP) TOQLOCAL(Queue_TO_MOVE)
```

9. The queue created in step 3 is no longer required. Use the following command to delete it:

```
DELETE QL(TEMP_QUEUE)
```

10. If the queue being moved was defined in the CSQINP2 data sets, change the STGCLASS attribute of the appropriate DEFINE QLOCAL command in the CSQINP2 data sets. Add the REPLACE keyword so that the existing queue definition is replaced.

Figure 56 on page 304 shows an extract from a load balancing job.

```

//UTILITY EXEC PGM=CSQUTIL,PARM=('CSQ1')
//STEPLIB DD DSN=thlqual.SCSQANLE,DISP=SHR
// DD DSN=thlqual.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(MOVEQ)
/*
//MOVEQ DD *
ALTER QL(Queue_To_Move) PUT(DISABLED)
DELETE QL(TEMP_QUEUE) PURGE
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED)
MOVE QLOCAL(Queue_To_Move) TOQLOCAL(TEMP_QUEUE)
DELETE QL(Queue_To_Move)
DEFINE STGCLASS(NEW) PSID(2)
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) STGCLASS(NEW)
MOVE QLOCAL(TEMP_QUEUE) TOQLOCAL(Queue_To_Move)
DELETE QL(TEMP_QUEUE)
/*

```

Figure 56. Extract from a load balancing job for a page set

How to increase the size of a page set

You can initially allocate a page set larger than 4 GB, See Defining a page set to be larger than 4 GB

A page set can be defined to be automatically expanded as it becomes full by specifying EXPAND(SYSTEM) or EXPAND(USER). If your page set was defined with EXPAND(NONE), you can expand it in either of two ways:

- Alter its definition to allow automatic expansion. See Altering a page set to allow automatic expansion
- Create a new, larger page set and copy the messages from the old page set to the new one. See Moving messages to a new, larger page set

Defining a page set to be larger than 4 GB

WebSphere MQ can use a page set up to 64 GB in size, provided the data set is defined with 'extended addressability' to VSAM. Extended addressability is an attribute which is conferred by an SMS data class. In the example shown in the following sample JCL, the management class 'EXTENDED' is defined to SMS with 'Extended addressability'. If your existing page set is not currently defined as having extended addressability, use the following method to migrate to an extended addressability format data set.

1. Stop the queue manager.
2. Use Access Method Services to rename the existing page set.
3. Define a destination page set, the same size as the existing page set, but with DATACLASS(EXTENDED).
4. Use the COPYPAGE function of CSQUTIL to copy all the messages from the source page set to the destination page set. See  Expanding a page set (COPYPAGE) (WebSphere MQ V7.1 Reference) for more details.
5. Restart the queue manager.
6. Alter the page set to use system expansion, to allow it to continue growing beyond its current allocation.

The following JCL shows example Access Method Services commands:

```

//S1      EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN   DD   *
ALTER 'VICY.CSQ1.PAGE01' -
      NEWNAME('VICY.CSQ1.PAGE01.OLD')
ALTER 'VICY.CSQ1.PAGE01.DATA' -
      NEWNAME('VICY.CSQ1.PAGE01.DATA.OLD')
DEFINE CLUSTER (NAME('VICY.CSQ1.PAGE01') -
              MODEL('VICY.CSQ1.PAGE01.OLD') -
              DATACLASS(EXTENDED))
/*

```

Altering a page set to allow automatic expansion

Use the ALTER PSID command with the EXPAND(USER) or EXPAND(SYSTEM) options. See

 ALTER PSID (*WebSphere MQ V7.1 Reference*) and  Expanding a page set (COPYPAGE) (*WebSphere MQ V7.1 Reference*) for general information on expanding page sets.

Moving messages to a new, larger page set

This technique involves stopping and restarting the queue manager. This deletes any nonpersistent messages that are not on shared queues at restart time. If you have nonpersistent messages that you do not want to be deleted, use load balancing instead. For more details, see “How to balance loads on page sets” on page 302. In this description, the page set that you want to expand is referred to as the *source* page set; the new, larger page set is referred to as the *destination* page set.

Follow these steps:

1. Stop the queue manager.
2. Define the destination page set, ensuring that it is larger than the source page set, with a larger secondary extent value.
3. Use the FORMAT function of CSQUTIL to format the destination page set. See  Formatting page sets (FORMAT) (*WebSphere MQ V7.1 Reference*) for more details.
4. Use the COPYPAGE function of CSQUTIL to copy all the messages from the source page set to the destination page set. See  Expanding a page set (COPYPAGE) (*WebSphere MQ V7.1 Reference*) for more details.
5. Restart the queue manager using the destination page set by doing one of the following:
 - Change the queue manager started task procedure to reference the destination page set.
 - Use Access Method Services to delete the source page set and then rename the destination page set, giving it the same name as that of the source page set.

Attention:

Before you delete any WebSphere MQ page set, be sure that you have made the required backup copies.

How to reduce a page set

Prevent all users, other than the WebSphere MQ administrator, from using the queue manager. For example; by changing the access security settings.

If you have a large page set that is mostly empty (as shown by the DISPLAY USAGE command), you might want to reduce its size. The procedure to do this involves using the COPY, FORMAT, and LOAD

functions of CSQUTIL (see  WebSphere MQ utility program (CSQUTIL) (*WebSphere MQ V7.1 Reference*)). This procedure does not work for page set zero (0), as it is not practical to reduce the size of this page set; the only way to do so is by reinitializing your queue manager (see “Reinitializing a queue manager” on page 325). The prerequisite of this procedure is to try and remove all users from the system so that all UOWs are complete and the page sets are consistent.

1. Use the **STOP QMGR** command with the QUIESCE or FORCE attribute to stop the queue manager.
2. Run the **SCOPY** function of **CSQUTIL** with the PSID option, to copy all message data from the large page set and save them in a sequential data set.
3. Define a new smaller page set data set to replace the large page set.
4. Run the FORMAT TYPE(NEW) function of CSQUTIL against the page set that you created in step 3.
5. Restart the queue manager using the page set created in step 3.
6. Run the LOAD function of CSQUTIL to load back all the messages saved during step 2.
7. Allow all users access to the queue manager.
8. Delete the old large page set.

How to reintroduce a page set

In certain scenarios it is useful to be able to bring an old page set online again to the queue manager. Unless specific action is taken, when the old page set is brought online the queue manager will recognize that the page set recovery RBA stored in the page set itself and in the checkpoint records is old, and will therefore automatically start media recovery of the page set to bring it up to date.

Such media recovery can only be performed at queue manager restart, and is likely to take a considerable length of time, especially if archive logs held on tape must be read. However, normally in this circumstance, the page set has been offline for the intervening period and so the log contains no information pertinent to the page set recovery.

The following three choices are available:

Allow full media recovery to be performed.

1. Stop the queue manager.
2. Ensure definitions are available for the page set in both the started task procedure for the queue manager and in the CSQINP1 initialization data set.
3. Restart the queue manager.

Allow any messages on the page set to be destroyed.

This choice is useful where a page set has been offline for a long time (some months, for example) and it has now been decided to reuse it for a different purpose.

1. Format the page set using the FORMAT function of CSQUTIL with the TYPE(NEW) option.
2. Add definitions for the page set to both the started task procedure for the queue manager and the CSQINP1 initialization data set.
3. Restart the queue manager.

Using the TYPE(NEW) option for formatting clears the current contents of the page set and tells the queue manager to ignore any historical information in the checkpoint about the page set.

Bring the page set online avoiding the media recovery process.

Use this technique only if you are sure that the page set has been offline since a clean shutdown of the queue manager. This choice is most appropriate where the page set has been offline for a short period, typically due to operational issues such as a backup running while the queue manager is being started.

1. Format the page set using the FORMAT function of CSQUTIL with the TYPE(REPLACE) option.

2. Either add the page set back into the queue manager dynamically using the DEFINE PSID command with the DSN option or allow it to be added at a queue manager restart.

Using the TYPE(REPLACE) option for formatting checks that the page set was cleanly closed by the queue manager, and marks it so that media recovery will not be performed. No other changes are made to the contents of the page set.

How to back up and recover page sets

There are different mechanisms available for back up and recovery. Use this topic to understand these mechanisms.

This section describes the following topics:

- “Creating a point of recovery for non-shared resources”
- “Backing up page sets” on page 308
- “Recovering page sets” on page 309
- How to delete page sets

For information about how to create a point of recovery for shared resources, see “Recovering shared queues” on page 315.

Creating a point of recovery for non-shared resources

WebSphere MQ can recover objects and non-shared persistent messages to their current state if both:

1. Copies of page sets from an earlier point exist.
2. All the WebSphere MQ logs are available to perform recovery from that point.

These represent a point of recovery for non-shared resources.

Both objects and messages are held on page sets. Multiple objects and messages from different queues can exist on the same page set. For recovery purposes, objects and messages cannot be backed up in isolation, so a page set must be backed up as a whole to ensure the correct recovery of the data.

The WebSphere MQ recovery log contains a record of all persistent messages and changes made to objects. If WebSphere MQ fails (for example, due to an I/O error on a page set), you can recover the page set by restoring the backup copy and restarting the queue manager. WebSphere MQ applies the log changes to the page set from the point of the backup copy.

There are two ways of creating a point of recovery:

Full backup

Stop the queue manager, which forces all updates on to the page sets.

This allows you to restart from the point of recovery, using only the backed up page set data sets and the logs from that point on.

Fuzzy backup

Take *fuzzy* backup copies of the page sets without stopping the queue manager.

If you use this method, and your associated logs later become damaged or lost, you cannot use the fuzzy page set backup copies to recover. This is because the fuzzy page set backup copies contain an inconsistent view of the state of the queue manager and are dependent on the logs being available. If the logs are not available, you need to return to the last set of backup page set copies taken while the subsystem was inactive (Method 1) and accept the loss of data from that time.

Method 1: Full backup

This method involves shutting the queue manager down. This forces all updates on to the page sets so that the page sets are in a consistent state.

1. Stop all the WebSphere MQ applications that are using the queue manager (allowing them to complete first). This can be done by changing the access security or queue settings, for example.
2. When all activity has completed, display and resolve any in-doubt units of recovery. (Use the commands DISPLAY CONN and RESOLVE INDOUBT, as described in  DISPLAY CONN (*WebSphere MQ V7.1 Reference*) and  RESOLVE INDOUBT (*WebSphere MQ V7.1 Reference*).) This brings the page sets to a consistent state; if you do not do this, your page sets might not be consistent, and you are effectively doing a fuzzy backup.
3. Issue the ARCHIVE LOG command to ensure that the latest log data is written out to the log data sets.
4. Issue the STOP QMGR MODE(QUIESCE) command. Record the lowest RBA value in the CSQI024I or CSQI025I messages (see  CSQI024I: csect-name Restart RBA for system as configured = restart-rba (*WebSphere MQ V7.1 Reference*) and  CSQI025I: csect-name Restart RBA including offline page sets = restart-rba (*WebSphere MQ V7.1 Reference*) for more information). You should keep the log data sets starting from the one indicated by the RBA value up to the current log data set.
5. Take backup copies of all the queue manager page sets (see “Backing up page sets”).

Method 2: Fuzzy backup

This method does not involve shutting the queue manager down. Therefore, updates might be in virtual storage buffers during the backup process. This means that the page sets are not in a consistent state, and can only be used for recovery with the logs.

1. Issue the DISPLAY USAGE TYPE(ALL) command, and record the RBA value in the CSQI024I or CSQI025I messages (see  CSQI024I: csect-name Restart RBA for system as configured = restart-rba (*WebSphere MQ V7.1 Reference*) and  CSQI025I: csect-name Restart RBA including offline page sets = restart-rba (*WebSphere MQ V7.1 Reference*) for more information).
2. Take backup copies of the page sets (see “Backing up page sets”).
3. Issue the ARCHIVE LOG command, to ensure that the latest log data is written out to the log data sets. To restart from the point of recovery, you must keep the log data sets starting from the log data set indicated by the RBA value up to the current log data set.

Backing up page sets

To recover a page set, WebSphere MQ needs to know how far back in the log to go. WebSphere MQ maintains a log RBA number in page zero of each page set, called the *recovery log sequence number (LSN)*. This number is the starting RBA in the log from which WebSphere MQ can recover the page set. When you back up a page set, this number is also copied.

If the copy is later used to recover the page set, WebSphere MQ must have access to all the log records from this RBA value to the current RBA. That means you must keep enough of the log records to enable WebSphere MQ to recover from the oldest backup copy of a page set you intend to keep.

Use ADRDSSU COPY function to copy the page sets.

For more information, see  COPY DATASET Command Syntax for Logical Data Set.

For example:

```
//STEP2 EXEC PGM=ADRDSSU,REGION=6M
//SYSPRINT DD SYSOUT=H
//SYSIN DD *
```

```

COPY -
DATASET(INCLUDE(SCENDATA.MQPA.PAGESET.*)) -
RENAMEU(SCENDATA.MQPA.PAGESET.** ,SCENDATA.MQPA.BACKUP1.** ) -
SPHERE -
REPUNC -
FASTREPLICATION(PREF )-
CANCELERROR -
TOL(ENQF)
/*
//

```

If you copy the page set while the queue manager is running you must use a copy utility that copies page zero of the page set first. If you do not do this you could corrupt the data in your page set.

If the process of dynamically expanding a page set is interrupted, for example by power to the system being lost, you can still use ADRDSSU to take a backup of a page set.

If you perform an Access Method Services IDCAMS LISTCAT ENT('page set data set name') ALLOC, you will see that the HI-ALLOC-RBA is higher than the HI-USED-RBA.

The next time this page set fills up it is extended again, if possible, and the pages between the high used RBA and the highest allocated RBA are used, along with another new extent.

Backing up your object definitions

You should also back up copies of your object definitions. To do this, use the MAKEDEF feature of the CSQUTIL COMMAND function (described in  Issuing commands to IBM WebSphere MQ (COMMAND) (*WebSphere MQ V7.1 Reference*)).

Back up your object definitions whenever you take a backup copy of your queue manager, and keep the most current version.

Recovering page sets

If the queue manager has terminated due to a failure, the queue manager can normally be restarted with all recovery being performed during restart. However, such recovery is not possible if any of your page sets or log data sets are not available. The extent to which you can now recover depends on the availability of backup copies of page sets and log data sets.

To restart from a point of recovery you must have:

- A backup copy of the page set that is to be recovered.
- If you used the “fuzzy” backup process described in “Method 2: Fuzzy backup” on page 308, the log data set that included the recorded RBA value, the log data set that was made by the ARCHIVE LOG command, and all the log data sets between these.
- If you used full backup, but you do not have the log data sets following that made by the ARCHIVE LOG command, you do **not** need to run the FORMAT TYPE(REPLACE) function of the CSQUTIL utility against all your page sets.

To recover a page set to its current state, you must also have all the log data sets and records since the ARCHIVE LOG command.

There are two methods for recovering a page set. To use either method, the queue manager must be stopped.

Simple recovery

This is the simpler method, and is appropriate for most recovery situations.

1. Delete and redefine the page set you want to restore from backup.
2. Use ADRDSSU to recover your page set. Define your new page set with a secondary extent value so that it can be expanded dynamically.

Alternatively, you can rename your backup copy to the original name, or change the CSQP00xx DD statement in your queue manager procedure to point to your backup page set. However, if you then lose or corrupt the page set, you will no longer have a backup copy to restore from.

3. Restart the queue manager.
4. When the queue manager has restarted successfully, you can restart your applications
5. Reinstall your normal backup procedures for the restored page.

Advanced recovery

This method provides performance advantages if you have a large page set to recover, or if there has been much activity on the page set since the last backup copy was taken. However, it requires more manual intervention than the simple method, which might increase the risk of error and the time taken to perform the recovery.

1. Delete and redefine the page set you want to restore from backup.
2. Use ADRDSSU to copy the backup copy of the page set into the new page set. Define your new page set with a secondary extent value so that it can be expanded dynamically.
Alternatively, you can rename your backup copy to the original name, or change the CSQP00xx DD statement in your queue manager procedure to point to your backup page set. However, if you then lose or corrupt the page set, you will no longer have a backup copy to restore from.
3. Change the CSQINP1 definitions for your queue manager to make the buffer pool associated with the page set being recovered as large as possible. By making the buffer pool large, you might be able to keep all the changed pages resident in the buffer pool and reduce the amount of I/O to the page set.
4. Restart the queue manager.
5. When the queue manager has restarted successfully, stop it (using quiesce) and then restart it using the normal buffer pool definition for that page set. After this second restart completes successfully, you can restart your applications
6. Reinstall your normal backup procedures for the restored page.

What happens when the queue manager is restarted

When the queue manager is restarted, it applies all changes made to the page set that are registered in the log, beginning at the restart point for the page set. WebSphere MQ can recover multiple page sets in this way. The page set is dynamically expanded, if required, during media recovery.

During restart, WebSphere MQ determines the log RBA to start from by taking the lowest value from the following:

- Recovery LSN from the checkpoint log record for each page set.
- Recovery LSN from page zero in each page set.
- The RBA of the oldest incomplete unit of recovery in the system at the time the backup was taken.

All object definitions are stored on page set zero. Messages can be stored on any available page set.

Note: The queue manager cannot restart if page set zero is not available.

How to delete page sets

You delete a page set by using the DELETE PSID command; see  [DELETE PSID \(WebSphere MQ V7.1 Reference\)](#) for details of this command.

You cannot delete a page set that is still referenced by any storage class. Use DISPLAY STGCLASS to find out which storage classes reference a page set.

The data set is deallocated from WebSphere MQ but is not deleted. It remains available for future use, or can be deleted using z/OS facilities.

Remove the page set from the started task procedure for your queue manager.

Remove the definition of the page set from your CSQINP1 initialization data set.

How to back up and restore queues using CSQUTIL

Use this topic as a reference for further information about back up and restore using CSQUTIL.

You can use the CSQUTIL utility functions for backing up and restoring queues. To back up a queue, use the COPY or SCOPY function to copy the messages from a queue onto a data set. To restore the queue, use the complementary function LOAD or SLOAD. For more information, see  [WebSphere MQ utility program \(CSQUTIL\) \(WebSphere MQ V7.1 Reference\)](#).

Managing buffer pools

Use this topic if you want to change or delete your buffer pools.

This topic describes how to alter and delete buffer pools. It contains these sections:

- “How to change the number of buffers in a buffer pool”
- “How to delete a buffer pool” on page 312

Buffer pools are defined during queue manager initialization, using DEFINE BUFFPOOL commands issued from the initialization input data set CSQINP1. Their sizes can be altered in response to business requirements while the queue manager is running, using the processes detailed in this topic. The queue manager records the current buffer pool sizes in checkpoint log records. These are automatically restored on subsequent queue manager restart. Use the DISPLAY USAGE command to display the current buffer pools, their sizes, and page set to buffer pool mapping.

You can also define buffer pools dynamically using the DEFINE PSID command with the DSN option, see  [WebSphere MQ Script \(MQSC\) Command Reference \(WebSphere MQ V7.1 Reference\)](#).

If you change buffer pools dynamically, you should also update their definitions in the initialization data set CSQINP1.

See the  [Planning on z/OS \(WebSphere MQ V7.1 Installing Guide\)](#) for a description of page sets, storage classes, buffers, and buffer pools, and some of the performance considerations that apply.

How to change the number of buffers in a buffer pool

If a buffer pool is too small, the condition can result in message CSQP020E on the console, you can allocate more buffers to it using the ALTER BUFFPOOL command as follows:

1. Determine how much space is available for new buffers by looking at the CSQY220I messages in the log. The available space is reported in MB. As a buffer has a size of 4 KB, each MB of available space allows you to allocate 256 buffers. Do not allocate all the free space to buffers, as some is required for other tasks.
2. If the reported free space is inadequate, release some buffers from another buffer pool using the command
`ALTER BUFFPOOL(buf-pool-id) BUFFERS(integer)`

where *buf-pool-id* is the buffer pool from which you want to reclaim space and *integer* is the new number of buffers to be allocated to this buffer pool, which must be smaller than the original number of buffers allocated to it.

3. Add buffers to the buffer pool you want to expand using the command
`ALTER BUFFPOOL(buf-pool-id) BUFFERS(integer)`

where *buf-pool-id* is the buffer pool to be expanded and *integer* is the new number of buffers to be allocated to this buffer pool, which must be larger than the original number of buffers allocated to it.

How to delete a buffer pool

When a buffer pool is no longer used by any page sets, delete it to release the virtual storage allocated to it.

You delete a buffer pool using the `DELETE BUFFPOOL` command. See  [WebSphere MQ Script \(MQSC\) Command Reference \(WebSphere MQ V7.1 Reference\)](#) for details. The command fails if any page sets are using this buffer pool.

See “How to delete page sets” on page 311 for information about how to delete page sets.

How to change the number of buffers in a buffer pool:

Using this topic to understand how to interpret log messages relating to storage usage, and how to alter your buffer pool settings.

If a buffer pool is too small, as shown by message CSQP020E on the console, you can allocate more buffers to it using the `ALTER BUFFPOOL` command as follows:

1. Determine how much space is available for new buffers by looking at the CSQY220I messages in the log. The available space is reported in MB or GB as appropriate. As a buffer has a size of 4 KB, each MB of available space allows you to allocate 256 buffers. Do not allocate all the free space to buffers, as some is required for other tasks.
2. If the reported free space is inadequate, release some buffers from another buffer pool using the command
`ALTER BUFFPOOL(buf-pool-id) BUFFERS(integer)`

where *buf-pool-id* is the buffer pool from which you want to reclaim space and *integer* is the new number of buffers to be allocated to this buffer pool, which must be smaller than the original number of buffers allocated to it.

3. Add buffers to the buffer pool you want to expand using the command
`ALTER BUFFPOOL(buf-pool-id) BUFFERS(integer)`

where *buf-pool-id* is the buffer pool to be expanded and *integer* is the new number of buffers to be allocated to this buffer pool, which must be larger than the original number of buffers allocated to it.

How to delete a buffer pool:

Use this topic to understand when, and how to delete your buffer pools.

When a buffer pool is no longer used by any page sets, delete it to release the virtual storage allocated to it.

You delete a buffer pool using the DELETE BUFFPOOL command. See  DEFINE NAMELIST (*WebSphere MQ V7.1 Reference*) for details. The command fails if any page sets are using this buffer pool.

See “How to back up and recover page sets” on page 307 for information on how to delete page sets.

Managing queue-sharing groups and shared queues

WebSphere MQ can use different types of shared resources, for example queue-sharing groups, shared queues, and the coupling facility. Use this topic to review the procedures needed to manage these shared resources.

This section contains information about the following topics:

- “Managing queue-sharing groups”
- “Managing shared queues” on page 315
- “Managing group objects” on page 320
- “Managing the coupling facility” on page 321

Managing queue-sharing groups:

Use this topic to understand how to add or remove a queue manager to a queue-sharing group, and manage the associated Db2 tables.

This topic has sections about the following tasks:

- “Adding a queue-sharing group to the Db2 tables”
- “Adding a queue manager to a queue-sharing group”
- “Removing a queue manager from a queue-sharing group” on page 314
- “Removing a queue-sharing group from the Db2 tables” on page 314
- “Validating the consistency of Db2 definitions” on page 315

Adding a queue-sharing group to the Db2 tables

To add a queue-sharing group to the Db2 tables, use the ADD QSG function of the queue-sharing group utility (CSQ5PQSG). This program is described in  The queue-sharing group utility (CSQ5PQSG) (*WebSphere MQ V7.1 Reference*). A sample is provided in thlqual.SCSQPROC(CSQ45AQS).

Adding a queue manager to a queue-sharing group

A queue manager can be added to a queue-sharing group and this topic describes some of the limitations.

To add a queue manager to a queue-sharing group, use the ADD QMGR function of the queue-sharing group utility (CSQ5PQSG). This program is described in  The queue-sharing group utility (CSQ5PQSG) (*WebSphere MQ V7.1 Reference*). A sample is provided in thlqual.SCSQPROC(CSQ45AQM).

The queue-sharing group must exist before you can add queue managers to it.

Queue-sharing groups have a name of up to four characters. The name must be unique in your network, and must be different from any queue manager names.

A queue manager can only be a member of one queue-sharing group.

Note: To add a Version 7.0 queue manager to an existing queue-sharing group containing Version 5.3 or 5.3.1 queue managers, you must first apply the WebSphere MQ for z/OS Version 5 Release 3 Coexistence PTF.

Removing a queue manager from a queue-sharing group

You can only remove a queue manager from a queue-sharing group if the queue manager's logs are not needed by another process. The logs are needed if they contain:

- the latest backup of one of the coupling facility (CF) application structures used by the queue-sharing group
- data needed by a future restore process, that is, the queue manager has used a recoverable structure since the time described by the last backup exclusion interval value.

If either or both of these points apply, the queue manager cannot be removed. To determine which queue managers' logs are needed for a future restore process, use the MQSC DISPLAY CFSTATUS command (see  DISPLAY CFSTATUS (*WebSphere MQ V7.1 Reference*) for details of this command).

If the queue manager logs are not needed, take the following steps to remove the queue manager from the queue-sharing group:

1. Resolve any indoubt units of work involving this queue manager.
2. Shut the queue manager down cleanly using STOP QMGR MODE(QUIESCE).
3. Wait for an interval at least equivalent to the value of the EXCLINT parameter you will specify in the BACKUP CFSTRUCT command in the next step.
4. On another queue manager, run a CF structure backup for each recoverable CF structure by using the MQSC BACKUP CFSTRUCT command and specifying an EXCLINT value as required in the previous step.
5. Use the REMOVE QMGR function of the CSQ5PQSG utility to remove the queue manager from the queue-sharing group. This program is described in  The queue-sharing group utility (CSQ5PQSG) (*WebSphere MQ V7.1 Reference*). A sample is provided in thlqual.SCSQPROC(CSQ45RQM).
6. Before restarting the queue manager, reset the QSGDATA system parameter to its default value. See  Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*) for information about how to tailor your system parameters.

Note, that when removing the last queue manager in a queue sharing group, you must use the FORCE option, rather than REMOVE. This removes the queue manager from the queue sharing group, whilst not performing the consistency checks of the queue manager logs being required for recovery. You should only perform this operation if you are deleting the queue sharing group.

Removing a queue-sharing group from the Db2 tables

To remove a queue-sharing group from the Db2 tables, use the REMOVE QSG function of the queue-sharing group utility (CSQ5PQSG). This program is described in  The queue-sharing group utility (CSQ5PQSG) (*WebSphere MQ V7.1 Reference*). A sample is provided in thlqual.SCSQPROC(CSQ45RQS).

You can only remove a queue-sharing group from the common Db2 data-sharing group tables after you have removed all the queue managers from the queue-sharing group (as described in “Removing a queue manager from a queue-sharing group” on page 314).

When the queue-sharing group record is deleted from the queue-sharing group administration table, all objects and administrative information relating to that queue-sharing group are deleted from other WebSphere MQ Db2 tables. This includes shared queue and group object information.

Validating the consistency of Db2 definitions

Problems for shared queues within a queue-sharing group can occur if the Db2 object definitions have, for any reason, become inconsistent.

To validate the consistency of the Db2 object definitions for queue managers, CF structures, and shared queues, use the VERIFY QSG function of the queue-sharing group utility (CSQ5PQSG). This program is described in  The queue-sharing group utility (CSQ5PQSG) (*WebSphere MQ V7.1 Reference*).

Managing shared queues:

Use this topic to understand how to recover, move, and migrate shared queues.

This section describes the following tasks:

- “Recovering shared queues”
- “Moving shared queues” on page 316
- “Migrating non-shared queues to shared queues” on page 319
- Suspending a Db2 connection

Recovering shared queues

WebSphere MQ can recover persistent messages on shared queues if all:

- Backups of the CF structures containing the messages have been performed.
- All the logs for all queue managers in the queue-sharing group are available, to perform recovery from the point the backups are taken.
- Db2 is available and the structure backup table is more recent than the most recent CF structure backup.

The messages on a shared queue are stored in a coupling facility (CF) structure. Persistent messages can be put onto shared queues, and like persistent messages on non-shared queues, they are copied to the queue manager log. The MQSC BACKUP CFSTRUCT and RECOVER CFSTRUCT commands are provided to allow the recovery of a CF structure in the unlikely event of a coupling facility failure. In such circumstances, any nonpersistent messages stored in the affected structure are lost, but persistent messages can be recovered. Any further application activity using the structure is prevented until the structure has been recovered.

To enable recovery, you must back up your coupling facility list structures frequently using the MQSC BACKUP CFSTRUCT command. The messages in the CF structure are written onto the active log data set of the queue manager making the backup. It writes a record of the backup to Db2: the name of the CF structure being backed up, the name of the queue manager doing the backup, the RBA range for this backup on that queue manager log, and the backup time. Back up CF list structures even if you are not actively using shared queues, for example, if you have set up a queue-sharing group intending to use it in the future.

You can recover a CF structure by issuing an MQSC RECOVER CFSTRUCT command to the queue manager that can perform the recovery; you can use any queue manager in the queue-sharing group. You can specify a single CF structure to be recovered, or you can recover several CF structures simultaneously.

As noted previously, it is important that you back up your CF list structures frequently, otherwise recovering a CF structure can take a long time. Moreover, the recovery process cannot be canceled.

The definition of a shared queue is kept in a Db2 database and can therefore be recovered if necessary using standard Db2 database procedures.

For more information about the BACKUP CFSTRUCT and RECOVER CFSTRUCT commands, see

 WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*).

For more information about planning your backup and recovery strategy for queue-sharing groups, see

 Planning on z/OS (*WebSphere MQ V7.1 Installing Guide*).

Moving shared queues

This section describes how to perform load balancing by moving a shared queue from one coupling facility structure to another. It also describes how to move a non-shared queue to a shared queue, and how to move a shared queue to a non-shared queue.

When you move a queue, you need to define a temporary queue as part of the procedure. This is because every queue must have a unique name, so you cannot have two queues of the same name, even if the queues have different queue dispositions. WebSphere MQ tolerates having two queues with the same name (as in step 2), but you cannot use the queues.

- Moving a queue from one coupling facility structure to another
- Moving a non-shared queue to a shared queue
- Moving a shared queue to a non-shared queue

Moving a queue from one coupling facility structure to another

To move queues and their messages from one CF structure to another, use the MQSC MOVE QLOCAL command (described in  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*)). When you have identified the queue or queues that you want to move to a new CF structure, use the following procedure to move each queue:

1. Ensure that the queue you want to move is not in use by any applications, that is, the queue attributes IPPROCS and OPPROCS are zero on all queue managers in the queue-sharing group.
2. Prevent applications from putting messages on the queue being moved by altering the queue definition to disable **MQPUTs**. Change the queue definition to PUT(DISABLED).
3. Define a temporary queue with the same attributes as the queue that is being moved using the following command:

```
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED) QSGDISP(QMGR)
```

Note: If this temporary queue exists from a previous run, delete it before doing the define.

4. Move the messages to the temporary queue using the following command:

```
MOVE QLOCAL(Queue_To_Move) TOQLOCAL(TEMP_QUEUE)
```

5. Delete the queue you are moving, using the command:

```
DELETE QLOCAL(Queue_To_Move)
```

6. Redefine the queue that is being moved, changing the CFSTRUCT attribute, using the following command:

```
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) CFSTRUCT(NEW) QSGDISP(SHARED)
```

When the queue is redefined, it is based on the temporary queue created in step 3 on page 316.

7. Move the messages back to the new queue using the command:

```
MOVE QLOCAL(TEMP) TOQLOCAL(Queue_To_Move)
```

8. The queue created in step 3 on page 316 is no longer required. Use the following command to delete it:

```
DELETE QL(TEMP_QUEUE)
```

9. If the queue being moved was defined in the CSQINP2 data sets, change the CFSTRUCT attribute of the appropriate DEFINE QLOCAL command in the CSQINP2 data sets. Add the REPLACE keyword so that the existing queue definition is replaced.

Figure 57 on page 318 shows a sample job for moving a queue from one CF structure to another.

```

//UTILITY EXEC PGM=CSQUTIL,PARM=('CSQ1')
//STEPLIB DD DSN=thlqual.SCSQANLE,DISP=SHR
// DD DSN=thlqual.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(MOVEQ)
/*
//MOVEQ DD *
ALTER QL(Queue_To_Move) PUT(DISABLED)
DELETE QL(TEMP_QUEUE) PURGE
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED) QSGDISP(QMGR)
MOVE QLOCAL(Queue_To_Move) TOQLOCAL(TEMP_QUEUE)
DELETE QL(Queue_To_Move)
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) CFSTRUCT(NEW) QSGDISP(SHARED)
MOVE QLOCAL(TEMP_QUEUE) TOQLOCAL(Queue_To_Move)
DELETE QL(TEMP_QUEUE)
/*

```

Figure 57. Sample job for moving a queue from one CF structure to another

Moving a non-shared queue to a shared queue

The procedure for moving a non-shared queue to a shared queue is like the procedure for moving a queue from one CF structure to another (see “Moving a queue from one coupling facility structure to another” on page 316). Figure 58 gives a sample job to do this.

Note: Remember that messages on shared queues are subject to certain restrictions on the maximum message size, message persistence, and queue index type, so you might not be able to move some non-shared queues to a shared queue.

```

//UTILITY EXEC PGM=CSQUTIL,PARM=('CSQ1')
//STEPLIB DD DSN=thlqual.SCSQANLE,DISP=SHR
// DD DSN=thlqual.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(MOVEQ)
/*
//MOVEQ DD *
ALTER QL(Queue_To_Move) PUT(DISABLED)
DELETE QL(TEMP_QUEUE) PURGE
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED)
MOVE QLOCAL(Queue_To_Move) TOQLOCAL(TEMP_QUEUE)
DELETE QL(Queue_To_Move)
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) CFSTRUCT(NEW) QSGDISP(SHARED)
MOVE QLOCAL(TEMP_QUEUE) TOQLOCAL(Queue_To_Move)
DELETE QL(TEMP_QUEUE)
/*

```

Figure 58. Sample job for moving a non-shared queue to a shared queue

Moving a shared queue to a non-shared queue

The procedure for moving a shared queue to a non-shared queue is like the procedure for moving a queue from one CF structure to another (see “Moving a queue from one coupling facility structure to another” on page 316).

Figure 59 on page 319 gives a sample job to do this.

```

//UTILITY EXEC PGM=CSQUTIL,PARM=('CSQ1')
//STEPLIB DD DSN=thlqua1.SCSQANLE,DISP=SHR
// DD DSN=thlqua1.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(MOVEQ)
/*
//MOVEQ DD *
ALTER QL(Queue_To_Move) PUT(DISABLED)
DELETE QL(TEMP_QUEUE) PURGE
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED) QSGDISP(QMGR)
MOVE QLOCAL(Queue_To_Move) TOQLOCAL(TEMP_QUEUE)
DELETE QL(Queue_To_Move)
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) STGCLASS(NEW) QSGDISP(QMGR)
MOVE QLOCAL(TEMP_QUEUE) TOQLOCAL(Queue_To_Move)
DELETE QL(TEMP_QUEUE)
/*

```

Figure 59. Sample job for moving a shared queue to a non-shared queue

Migrating non-shared queues to shared queues

There are two stages to migrating non-shared queues to shared queues:

- Migrating the first (or only) queue manager in the queue-sharing group
- Migrating any other queue managers in the queue-sharing group

Migrating the first (or only) queue manager in the queue-sharing group

Figure 58 on page 318 shows an example job for moving a non-shared queue to a shared queue. Do this for each queue that needs migrating.

Note:

1. Messages on shared queues are subject to certain restrictions on the maximum message size, message persistence, and queue index type, so you might not be able to move some non-shared queues to a shared queue.
2. You must use the correct index type for shared queues. If you migrate a transmission queue to be a shared queue, the index type must be MSGID.

If the queue is empty, or you do not need to keep the messages that are on it, migrating the queue is simpler. Figure 60 shows an example job to use in these circumstances.

```

//UTILITY EXEC PGM=CSQUTIL,PARM=('CSQ1')
//STEPLIB DD DSN=thlqua1.SCSQANLE,DISP=SHR
// DD DSN=thlqua1.SCSQAUTH,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
COMMAND DDNAME(MOVEQ)
/*
//MOVEQ DD *
DELETE QL(TEMP_QUEUE) PURGE
DEFINE QL(TEMP_QUEUE) LIKE(Queue_To_Move) PUT(ENABLED) GET(ENABLED)
DELETE QL(Queue_To_Move)
DEFINE QL(Queue_To_Move) LIKE(TEMP_QUEUE) CFSTRUCT(NEW) QSGDISP(SHARED)
DELETE QL(TEMP_QUEUE)
/*

```

Figure 60. Sample job for moving a non-shared queue without messages to a shared queue

Migrating any other queue managers in the queue-sharing group

1. For each queue that does not have the same name as an existing shared queue, move the queue as described in Figure 58 on page 318 or Figure 60 on page 319.
2. For queues that have the same name as an existing shared queue, move the messages to the shared queue using the commands shown in Figure 61.

```
MOVE QLOCAL(Queue_To_Move) QSGDISP(QMGR) TOQLOCAL(Queue_To_Move)
DELETE QLOCAL(Queue_To_Move) QSGDISP(QMGR)
```

Figure 61. Moving messages from a non-shared queue to an existing shared queue

Suspending a connection to Db2

If you want to apply maintenance or service to the Db2 tables or package related to shared queues without stopping your queue manager, you must temporarily disconnect queue managers in the data sharing group (DSG) from Db2.

To do this:

1. Use the MQSC command  `SUSPEND QMGR (WebSphere MQ V7.1 Reference) FACILITY(Db2)`.
2. Do the binds.
3. Reconnect to Db2 using the MQSC command  `RESUME QMGR (WebSphere MQ V7.1 Reference) FACILITY(Db2)`.

There are restrictions on the use of these commands. See  `WebSphere MQ Script (MQSC) Command Reference (WebSphere MQ V7.1 Reference)` for details.

Attention: While the Db2 connection is suspended, the following operations will not be available. Therefore, you need to do this work during a time when your enterprise is at its least busy.

- Access to Shared queue objects for administration (define, delete, alter)
- Starting shared channels
- Storing messages in Db2
- Backup or recover CFSTRUCT

Managing group objects:

Use this topic to understand how to work with group objects.

WebSphere MQ automatically copies the definition of a group object to page set zero of each queue manager that uses it. You can alter the copy of the definition temporarily, and WebSphere MQ allows you to refresh the page set copies from the repository copy. WebSphere MQ always tries to refresh the page set copies from the repository copy on start-up (for channel objects, this is done when the channel initiator restarts). This ensures that the page set copies reflect the version on the repository, including any changes that were made when the queue manager was inactive.

There are circumstances under which the refresh is not performed, for example:

- If a copy of the queue is open, a refresh that would change the usage of the queue fails.
- If a copy of a queue has messages on it, a refresh that would delete that queue fails.

In these circumstances, the refresh is not performed on that copy, but is performed on the copies on all other queue managers. Check for and correct any problems with copy objects after adding, changing, or deleting a group object, and at queue manager or channel initiator restart.

Managing the coupling facility:

Use this topic to understand how to add or remove coupling facility (CF) structures.

This section describes the following tasks:

- “Adding a coupling facility structure”
- “Removing a coupling facility structure”

Adding a coupling facility structure

There are no WebSphere MQ actions required when you add a coupling facility structure. The information about setting up the coupling facility in  Task 10: Set up the coupling facility (*WebSphere MQ V7.1 Installing Guide*) describes the rules for naming coupling facility structures, and how to define structures in the CFRM policy data set.

Removing a coupling facility structure

To remove a coupling facility structure, follow this procedure:

- Use the following command to get a list of all the queues using the coupling facility structure that you want to delete:

```
DISPLAY QUEUE(*) QSGDISP(SHARED) CFSTRUCT(structure-name)
```

- Delete all the queues that use the structure.
- Stop and restart each queue manager in the queue-sharing group in turn to disconnect WebSphere MQ and Db2 from the structure and delete information about it. (You do not need to stop all the queue managers at once; just one at a time.)
- Remove the structure definition from your CFRM policy data set and run the IXCMIAPU utility. (This is the reverse of customization task 10 (set up the coupling facility) described in  Task 10: Set up the coupling facility (*WebSphere MQ V7.1 Installing Guide*).)

Recovery and restart

Use this topic to understand the recovery and restart mechanisms used by WebSphere MQ.

Restarting WebSphere MQ

After a queue manager terminates there are different restart procedures needed depending on how the queue manager terminated. Use this topic to understand the different restart procedures that you can use.

This topic contains information about how to restart your queue manager in the following circumstances:

- “Restarting after a normal shutdown” on page 322
- “Restarting after an abnormal termination” on page 322
- “Restarting if you have lost your page sets” on page 322
- “Restarting if you have lost your log data sets” on page 322
- Restarting if you have lost your CF structures

Restarting after a normal shutdown

If the queue manager was stopped with the STOP QMGR command, the system finishes its work in an orderly way and takes a termination checkpoint before stopping. When you restart the queue manager, it uses information from the system checkpoint and recovery log to determine the system status at shutdown.

To restart the queue manager, issue the START QMGR command as described in “Starting and stopping a queue manager” on page 250.

Restarting after an abnormal termination

WebSphere MQ automatically detects whether restart follows a normal shutdown or an abnormal termination.

Starting the queue manager after it has terminated abnormally is different from starting it after the STOP QMGR command has been issued. If the queue manager terminates abnormally, it terminates without being able to finish its work or take a termination checkpoint.

To restart the queue manager, issue the START QMGR command as described in “Starting and stopping a queue manager” on page 250. When you restart a queue manager after an abnormal termination, it refreshes its knowledge of its status at termination using information in the log, and notifies you of the status of various tasks.

Normally, the restart process resolves all inconsistent states. But, in some cases, you must take specific steps to resolve inconsistencies. This is described in “Recovering units of work manually” on page 344.

Restarting if you have lost your page sets

If you have lost your page sets, you need to restore them from your backup copies before you can restart the queue manager. This is described in “How to back up and recover page sets” on page 307.

The queue manager might take a long time to restart under these circumstances because of the length of time needed for media recovery.

Restarting if you have lost your log data sets

If, after stopping a queue manager (using the STOP QMGR command), both copies of the log are lost or damaged, you can restart the queue manager providing you have a consistent set of page sets (produced using Method 1: Full backup).

Follow this procedure:

1. Define new page sets to correspond to each existing page set in your queue manager. See  Task 15: Define your page sets (*WebSphere MQ V7.1 Installing Guide*) for information about page set definition. Ensure that each new page set is larger than the corresponding source page set.
2. Use the FORMAT function of CSQUTIL to format the destination page set. See  Formatting page sets (FORMAT) (*WebSphere MQ V7.1 Reference*) for more details.
3. Use the RESETPAGE function of CSQUTIL to copy the existing page sets or reset them in place, and reset the log RBA in each page. See  Copying a page set and resetting the log (RESETPAGE) (*WebSphere MQ V7.1 Reference*) for more information about this function.
4. Redefine your queue manager log data sets and BSDS using CSQJU003 (see  The change log inventory utility (CSQJU003) (*WebSphere MQ V7.1 Reference*)).

5. Restart the queue manager using the new page sets. To do this, you do one of the following:
 - Change the queue manager started task procedure to reference the new page sets. See  Task 6: Create procedures for the WebSphere MQ queue manager (*WebSphere MQ V7.1 Installing Guide*) for more information.
 - Use Access Method Services to delete the old page sets and then rename the new page sets, giving them the same names as the old page sets.

Attention: Before you delete any WebSphere MQ page set, ensure that you have made the required backup copies.

If the queue manager is a member of a queue-sharing group, GROUP and SHARED object definitions are not normally affected by lost or damaged logs. However, if any shared-queue messages are involved in a unit of work that was covered by the lost or damaged logs, the effect on such uncommitted messages is unpredictable.

Note: If logs are damaged and the queue manager is a member of a queue-sharing group, the ability to recover shared persistent messages might be lost. Issue a BACKUP CFSTRUCT command immediately on another active queue manager in the queue-sharing group for all CF structures with the RECOVER(YES) attribute.

Restarting if you have lost your CF structures

You do not need to restart if you lose your CF structures, because the queue manager does not terminate.

Recovering a queue-sharing group at the alternative site:

If you need to recover a queue-sharing group, you must understand the steps needed to prepare the recovery, and the steps needed to recover the queue managers in the queue-sharing group.

Before you can recover the queue-sharing group, you need to prepare the environment:

1. If you have old information in your coupling facility from practice startups when you installed the queue-sharing group, you need to clean this out first (if you do not have old information in the coupling facility, you can omit this step:
 - a. Enter the following z/OS command to display the CF structures for this queue-sharing group:

```
D XCF,STRUCTURE,STRNAME=qsgname
```

- b. For all structures that start with the queue-sharing group name, use the z/OS command SETXCF FORCE CONNECTION to force the connection off those structures:

```
SETXCF FORCE,CONNECTION,STRNAME=strname,CONNAME=ALL
```

- c. Delete all the CF structures using the following command for each structure:

```
SETXCF FORCE,STRUCTURE,STRNAME=strname
```

2. Restore Db2 systems and data-sharing groups.
3. Recover the CSQ.ADMIN_B_STRBACKUP table so that it contains information about the most recent structure backups taken at the prime site.

Note: It is important that the STRBACKUP table contains the most recent structure backup information. Older structure backup information might require data sets that you have discarded as a result of the information given by a recent DISPLAY USAGE TYPE(DATASET) command, which would mean that your recovered CF structure would not contain accurate information.

4. Run the ADD QMGR command of the CSQ5PQSG utility for every queue manager in the queue-sharing group. This will restore the XCF group entry for each queue manager.

When you run the utility in this scenario, the following messages are normal:

```
CSQU566I Unable to get attributes for admin structure, CF not found
or not allocated
CSQU546E Unable to add QMGR <queue-manager-name> entry,
already exists in DB2 table CSQ.ADMIN_B_QMGR
CSQU148I CSQ5PQSG Utility completed, return code=4
```

To recover the queue managers in the queue-sharing group:

1. Define new page set data sets and load them with the data in the copies of the page sets from the primary site.
2. Define new active log data sets.
3. Define a new BSDS data set and use Access Method Services REPRO to copy the *most recent* archived BSDS into it.
4. Use the print log map utility CSQJU004 to print information from this most recent BSDS. At the time this BSDS was archived, the most recent archived log you have would have just been truncated as an active log, and does not appear as an archived log. Record the STARTRBA, STARTLRSN, ENDRBA, and ENDLRSN values of this log.
5. Use the change log inventory utility, CSQJU003, to register this latest archive log data set in the BSDS that you have just restored, using the values recorded in Step 4.
6. Use the DELETE option of CSQJU003 to remove all active log information from the BSDS.
7. Use the NEWLOG option of CSQJU003 to add active logs to the BSDS, do not specify STARTRBA or ENDRBA.
8. Calculate the *recoverylrsn* for the queue-sharing group. The *recoverylrsn* is the lowest of the ENDLRSNs across all queue managers in the queue-sharing group (as recorded in Step 4), minus 1. For example, if there are two queue managers in the queue-sharing group, and the ENDLRSN for one of them is B713 3C72 22C5, and for the other is B713 3D45 2123, the *recoverylrsn* is B713 3C72 22C4.
9. Use CSQJU003 to add a restart control record to the BSDS. Specify:

```
CRESTART CREATE,ENDLRSN=recoverylrsn
```

where *recoverylrsn* is the value you recorded in Step 8.

The BSDS now describes all active logs as being empty, all the archived logs you have available, and no checkpoints beyond the end of your logs.

You must add the CRESTART record to the BSDS for each queue manager within the queue-sharing group.

10. Restart each queue manager in the queue-sharing group with the usual START QMGR command. During initialization, an operator reply message such as the following is issued:

```
CSQJ245D +CSQ1 RESTART CONTROL INDICATES TRUNCATION AT RBA highrba.  
REPLY Y TO CONTINUE, N TO CANCEL
```

Reply Y to start the queue manager. The queue manager starts, and recovers data up to ENDRBA specified in the CRESTART statement.

At WebSphere MQ Version 7.0.1 and later, the first queue manager started can rebuild the admin structure partitions for other members of the queue sharing group as well as its own, and it is no longer necessary to restart each queue manager in the queue sharing group at this stage.

11. When the admin structure data for all queue managers has been rebuilt, issue a RECOVER CFSTRUCT command for each CF application structure.

If you issue the RECOVER CFSTRUCT command for all structures on a single queue manager, the log merge process is only performed once, so is quicker than issuing the command on a different queue manager for each CF structure, where each queue manager has to perform the log merge step.

When conditional restart processing is used in a queue sharing group, WebSphere MQ Version 7.0.1 and later queue managers, performing peer admin rebuild, check that peers BSDS contain the same CRESTART LRSN as their own. This is to ensure the integrity of the rebuilt admin structure. It is therefore important to restart other peers in the QSG, so they can process their own CRESTART information, before the next unconditional restart of any member of the group.

Reinitializing a queue manager:

If the queue manager has terminated abnormally you might not be able to restart it. This could be because your page sets or logs have been lost, truncated, or corrupted. If this has happened, you might have to reinitialize the queue manager (perform a cold start).

Attention

Only perform a cold start if you cannot restart the queue manager any other way. Performing a cold start enables you to recover your queue manager and your object definitions; you will **not** be able to recover your message data. Check that none of the other restart scenarios described in this topic work for you before you do this.

When you have restarted, all your WebSphere MQ objects are defined and available for use, but there is no message data.

Note: Do not reinitialize a queue manager while it is part of a cluster. You must first remove the queue manager from the cluster (using RESET CLUSTER commands on the other queue managers in the cluster), then reinitialize it, and finally reintroduce it to the cluster as a new queue manager.

This is because during reinitialization, the queue manager identifier (QMID) is changed, so any cluster object with the old queue manager identifier must be removed from the cluster.

For further information see the following sections:

- Reinitializing a queue manager that is not in a queue-sharing group
- Reinitializing queue managers in a queue-sharing group

Reinitializing a queue manager that is not in a queue-sharing group

To reinitialize a queue manager, follow this procedure:

1. Prepare the object definition statements that to be used when you restart the queue manager. To do this, either:

- If page set zero is available, use the CSQUTIL SDEFS function (see  Producing a list of WebSphere MQ define commands (SDEFS) (*WebSphere MQ V7.1 Reference*)). You must get definitions for all object types (authentication information objects, CF structures, channels, namelists, processes, queues, and storage classes).
 - If page set zero is not available, use the definitions from the last time you backed up your object definitions.
2. Redefine your queue manager data sets (do not do this until you have completed step 1 on page 325).
 3. Restart the queue manager using the newly defined and initialized log data sets, BSDS, and page sets. Use the object definition input statements that you created in step 1 on page 325 as input in the CSQINP2 initialization input data set.

Reinitializing queue managers in a queue-sharing group

In a queue-sharing group, reinitializing a queue manager is more complex. It might be necessary to reinitialize one or more queue managers because of page set or log problems, but there might also be problems with Db2 or the coupling facility to deal with. Because of this, there are a number of alternatives:

Cold start

Reinitializing the entire queue-sharing group involves forcing all the coupling facilities structures, clearing all object definitions for the queue-sharing group from Db2, deleting or redefining the logs and BSDS, and formatting page sets for all the queue managers in the queue-sharing group.

Shared definitions retained

Delete or redefine the logs and BSDS, format page sets for all queue managers in the queue-sharing group, and force all the coupling facilities structures. On restart, all messages will have been deleted. The queue managers re-create COPY objects that correspond to GROUP objects that still exist in the Db2 database. Any shared queues still exist and can be used.

Single queue manager reinitialized

Delete or redefine the logs and BSDS, and format page sets for the single queue manager (this deletes all its private objects and messages). On restart, the queue manager re-creates COPY objects that correspond to GROUP objects that still exist in the Db2 database. Any shared queues still exist, as do the messages on them, and can be used.

Point in time recovery of a queue-sharing group

This is the alternative site disaster recovery scenario.

Shared objects are recovered to the point in time achieved by Db2 recovery (described in “A Db2 system fails” on page 1374). Each queue manager can be recovered to a point in time achievable from the backup copies available at the alternative site.

Persistent messages can be used in queue-sharing groups, and can be recovered using the MQSC RECOVER CFSTRUCT command. Note that this command recovers to the time of failure. However, there is no recovery of nonpersistent shared queue messages; they are lost unless you have made backup copies independently using the COPY function of the CSQUTIL utility program.

It is not necessary to try to restore each queue manager to the same point in time because there are no interdependencies between the local objects on different queue managers (which are what is actually being recovered), and the queue manager resynchronization with Db2 on restart creates or deletes COPY objects as necessary on a queue manager by queue manager basis.

Starting and stopping a queue manager

Use this topic as an introduction to stopping and starting a queue manager.

This section describes how to start and stop a queue manager. It contains information about the following topics:

- “Before you start WebSphere MQ” on page 250
- “Starting a queue manager” on page 250
- “Stopping a queue manager” on page 252

Starting and stopping a queue manager is relatively straightforward. When a queue manager stops under normal conditions, its last action is to take a termination checkpoint. This checkpoint, and the logs, give the queue manager the information it needs to restart.

This section contains information about the START and STOP commands, and contains a brief overview of start-up after an abnormal termination has occurred.

Before you start WebSphere MQ

After you have installed WebSphere MQ, it is defined as a formal z/OS subsystem. This message appears during any initial program load (IPL) of z/OS:

```
CSQ3110I +CSQ1 CSQ3UR00 - SUBSYSTEM ssnm INITIALIZATION COMPLETE
```

where *ssnm* is the WebSphere MQ subsystem name.

From now on, you can start the queue manager for that subsystem *from any z/OS console that has been authorized to issue system control commands*; that is, a z/OS SYS command group. You must issue the START command from the authorized console, you cannot issue it through JES or TSO.

If you are using queue-sharing groups, you must start RRS first, and then Db2, before you start the queue manager.

Starting a queue manager

You start a queue manager by issuing a START QMGR command. However, you cannot successfully use the START command unless you have appropriate authority. See the “Setting up security on z/OS” on page 520 for information about WebSphere MQ security. Figure 52 on page 250 shows examples of the START command. (Remember that you must prefix a WebSphere MQ command with a command prefix string (CPF).)

```
+CSQ1 START QMGR  
+CSQ1 START QMGR PARM(NEWLOG)
```

Figure 62. Starting the queue manager from a z/OS console. The second example specifies a system parameter module name.

See  START QMGR (*WebSphere MQ V7.1 Reference*) for information about the syntax of the START QMGR command.

You cannot run the queue manager as a batch job or start it using a z/OS command START. These methods are likely to start an address space for WebSphere MQ that then ends abnormally. Nor can you start a queue manager from the CSQUTIL utility program or a similar user application.

You can, however, start a queue manager from an APF-authorized program by passing a START QMGR command to the z/OS MGCRC (SVC 34) service.

If you are using queue-sharing groups, the associated Db2 systems and RRS must be active when you start the queue manager.

Start options

When you start a queue manager, a system parameter module is loaded. You can specify the name of the system parameter module in one of two ways:

- With the PARM parameter of the /cpf START QMGR command, for example
/cpf START QMGR PARM(CSQ1ZPRM)
- With a parameter in the startup procedure, for example, code the JCL EXEC statement as
//MQM EXEC PGM=CSQYASCP,PARM='ZPARM(CSQ1ZPRM)'

A system parameter module provides information specified when the queue manager was customized. In “User messages issued when the queue manager starts” on page 254, the user message CSQY001I indicates the name of the system parameter module that was used, in this case, CSQ1ZPRM. .

You can also use the ENVPARM option to substitute one or more parameters in the JCL procedure for the queue manager.

For example, you can update your queue manager startup procedure, so that the DDname CSQINP2 is a variable. This means that you can change the CSQINP2 DDname without changing the startup procedure. This is useful for implementing changes, providing backouts for operators, and queue manager operations.

Suppose your start-up procedure for queue manager CSQ1 looked like Figure 53 on page 251.

```
//CSQ1MSTR PROC INP2=NORM
//MQMESA EXEC PGM=CSQYASCP
//STEPLIB DD DISP=SHR,DSN=th1qua1.SCSQANLE
// DD DISP=SHR,DSN=th1qua1.SCSQAUTH
// DD DISP=SHR,DSN=db2qua1.SDSNLOAD
//BDS1 DD DISP=SHR,DSN=myqua1.BSDS01
//BDS2 DD DISP=SHR,DSN=myqua1.BSDS02
//CSQP0000 DD DISP=SHR,DSN=myqua1.PSID00
//CSQP0001 DD DISP=SHR,DSN=myqua1.PSID01
//CSQP0002 DD DISP=SHR,DSN=myqua1.PSID02
//CSQP0003 DD DISP=SHR,DSN=myqua1.PSID03
//CSQINP1 DD DISP=SHR,DSN=myqua1.CSQINP(CSQ1INP1)
//CSQINP2 DD DISP=SHR,DSN=myqua1.CSQINP(CSQ1&INP2.)
//CSQOUT1 DD SYSOUT=*
//CSQOUT2 DD SYSOUT=*
```

Figure 63. Sample start-up procedure

If you then start your queue manager with the command:

```
+CSQ1 START QMGR
```

the CSQINP2 used is a member called CSQ1NORM.

However, suppose you are putting a new suite of programs into production so that the next time you start queue manager CSQ1, the CSQINP2 definitions are to be taken from member CSQ1NEW. To do this, you would start the queue manager with this command:

```
+CSQ1 START QMGR ENVPARM('INP2=NEW')
```

and CSQ1NEW would be used instead of CSQ1NORM. Note: z/OS limits the KEYWORD=value specifications for symbolic parameters (as in INP2=NEW) to 255 characters.

Starting after an abnormal termination

WebSphere MQ automatically detects whether restart follows a normal shutdown or an abnormal termination.

Starting a queue manager after it ends abnormally is different from starting it after the STOP QMGR command has been issued. After STOP QMGR, the system finishes its work in an orderly way and takes a termination checkpoint before stopping. When you restart the queue manager, it uses information from the system checkpoint and recovery log to determine the system status at shutdown.

However, if the queue manager ends abnormally, it terminates without being able to finish its work or take a termination checkpoint. When you restart a queue manager after an abend, it refreshes its knowledge of its status at termination using information in the log, and notifies you of the status of various tasks. Normally, the restart process resolves all inconsistent states. But, in some cases, you must take specific steps to resolve inconsistencies.

User messages on start-up

When you start a queue manager successfully, it produces a set of start-up messages similar to the ones in “User messages issued when the queue manager starts” on page 254.

Stopping a queue manager

Before stopping a queue manager, all WebSphere MQ-related write-to-operator-with-reply (WTOR) messages must receive replies, for example, getting log requests. Each command in Figure 54 on page 252 terminates a running queue manager.

```
+CSQ1 STOP QMGR
+CSQ1 STOP QMGR MODE(QUIESCE)
+CSQ1 STOP QMGR MODE(FORCE)
+CSQ1 STOP QMGR MODE(RESTART)
```

Figure 64. Stopping a queue manager

The command STOP QMGR defaults to STOP QMGR MODE(QUIESCE).

In QUIESCE mode, WebSphere MQ does not allow any new connection threads to be created, but allows existing threads to continue; it terminates only when all threads have ended. Applications can request to

be notified in the event of the queue manager quiescing. Therefore, use the QUIESCE mode where possible so that applications that have requested notification have the opportunity to disconnect. See

 What happens during termination (*WebSphere MQ V7.1 Product Overview Guide*) for details.

If the queue manager does not terminate in a reasonable time in response to a STOP QMGR MODE(QUIESCE) command, use the DISPLAY CONN command to determine whether any connection threads exist, and take the necessary steps to terminate the associated applications. If there are no threads, issue a STOP QMGR MODE(FORCE) command.

The STOP QMGR MODE(QUIESCE) and STOP QMGR MODE(FORCE) commands deregister WebSphere MQ from the MVS Automatic Restart Manager (ARM), preventing ARM from restarting the queue manager automatically. The STOP QMGR MODE(RESTART) command works in the same way as the STOP QMGR MODE(FORCE) command, except that it does not deregister WebSphere MQ from ARM. This means that the queue manager is eligible for immediate automatic restart.

If the WebSphere MQ subsystem is not registered with ARM, the STOP QMGR MODE(RESTART) command is rejected and the following message is sent to the z/OS console:

```
CSQY205I  ARM element arm-element is not registered
```

If this message is not issued, the queue manager is restarted automatically. For more information about ARM, see “Using the z/OS Automatic Restart Manager (ARM)” on page 339.

Only cancel the queue manager address space if STOP QMGR MODE(FORCE) does not terminate the queue manager.

If a queue manager is stopped by either canceling the address space or by using the command STOP QMGR MODE(FORCE), consistency is maintained with connected CICS or IMS systems. Resynchronization of resources is started when a queue manager restarts and is completed when the connection to the CICS or IMS system is established.

Note: When you stop your queue manager, you might find message IEF352I is issued. z/OS issues this message if it detects that failing to mark the address space as unusable would lead to an integrity exposure. You can ignore this message.

Stop messages

After issuing a STOP QMGR command, you get the messages CSQY009I and CSQY002I, for example:

```
CSQY009I +CSQ1 ' STOP QMGR' COMMAND ACCEPTED FROM  
USER(userid), STOP MODE(FORCE)  
CSQY002I +CSQ1 QUEUE MANAGER STOPPING
```

Where *userid* is the user ID that issued the STOP QMGR command, and the MODE parameter depends on that specified in the command.

When the STOP command has completed successfully, the following messages are displayed on the z/OS console:

```
CSQ9022I +CSQ1 CSQYASCP ' STOP QMGR' NORMAL COMPLETION
CSQ3104I +CSQ1 CSQ3EC0X - TERMINATION COMPLETE
```

If you are using ARM, and you did not specify MODE(RESTART), the following message is also displayed:

```
CSQY204I +CSQ1 ARM Deregister for element arm-element type
arm-element-type successful
```

You cannot restart the queue manager until the following message has been displayed:

```
CSQ3100I +CSQ1 CSQ3EC0X - SUBSYSTEM ssnm READY FOR START COMMAND
```

Alternative site recovery

You can recover a single queue manager or a queue-sharing group, or consider disk mirroring.

See the following sections for more details:

- Recovering a single queue manager at an alternative site
- Recovering a queue-sharing group.
 - CF structure media recovery
 - Backing up the queue-sharing group at the prime site
 - Recovering a queue-sharing group at the alternative site
- Using disk mirroring

Recovering a single queue manager at an alternative site

If a total loss of a WebSphere MQ computing center occurs, you can recover on another queue manager or queue-sharing group at a recovery site. (See “Recovering a queue-sharing group at the alternative site” on page 335 for the alternative site recovery procedure for a queue-sharing group.)

To recover on another queue manager at a recovery site, you must regularly back up the page sets and the logs. As with all data recovery operations, the objectives of disaster recovery are to lose as little data, workload processing (updates), and time, as possible.

At the recovery site:

- The recovery queue managers **must** have the same names as the lost queue managers.
- The system parameter module (for example, CSQZPARM) used on each recovery queue manager must contain the same parameters as the corresponding lost queue manager.

When you have done this, reestablish all your queue managers as described in the following procedure. This can be used to perform disaster recovery at the recovery site for a single queue manager. It assumes that all that is available are:

- Copies of the archive logs and BSDSs created by normal running at the primary site (the active logs will have been lost along with the queue manager at the primary site).
- Copies of the page sets from the queue manager at the primary site that are the same age or older than the most recent archive log copies available.

You can use dual logging for the active and archive logs, in which case you need to apply the BSDS updates to both copies:

1. Define new page set data sets and load them with the data in the copies of the page sets from the primary site.
2. Define new active log data sets.
3. Define a new BSDS data set and use Access Method Services REPRO to copy the *most recent* archived BSDS into it.
4. Use the print log map utility CSQJU004 to print information from this most recent BSDS. At the time this BSDS was archived, the most recent archived log you have would have just been truncated as an active log, and does not appear as an archived log. Record the STARTRBA and ENDRBA of this log.
5. Use the change log inventory utility, CSQJU003, to register this latest archive log data set in the BSDS that you have just restored, using the STARTRBA and ENDRBA recorded in Step 4.
6. Use the DELETE option of CSQJU003 to remove all active log information from the BSDS.
7. Use the NEWLOG option of CSQJU003 to add active logs to the BSDS, do not specify STARTRBA or ENDRBA.
8. Use CSQJU003 to add a restart control record to the BSDS. Specify CRESTART CREATE, ENDRBA=highrba, where highrba is the high RBA of the most recent archive log available (found in Step 4), plus 1.
The BSDS now describes all active logs as being empty, all the archived logs you have available, and no checkpoints beyond the end of your logs.
9. Restart the queue manager with the START QMGR command. During initialization, an operator reply message such as the following is issued:

```
CSQJ245D +CSQ1 RESTART CONTROL INDICATES TRUNCATION AT RBA highrba.  
REPLY Y TO CONTINUE, N TO CANCEL
```

Type Y to start the queue manager. The queue manager starts, and recovers data up to ENDRBA specified in the CRESTART statement.

See  Using the WebSphere MQ Utilities (*WebSphere MQ V7.1 Reference*) for information about using CSQJU003 and CSQJU004.

Figure 65 on page 333 shows sample input statements for CSQJU003 for steps 6, 7, and 8:

```

* Step 6
DELETE DSNAME=MQM2.LOGCOPY1.DS01
DELETE DSNAME=MQM2.LOGCOPY1.DS02
DELETE DSNAME=MQM2.LOGCOPY1.DS03
DELETE DSNAME=MQM2.LOGCOPY1.DS04
DELETE DSNAME=MQM2.LOGCOPY2.DS01
DELETE DSNAME=MQM2.LOGCOPY2.DS02
DELETE DSNAME=MQM2.LOGCOPY2.DS03
DELETE DSNAME=MQM2.LOGCOPY2.DS04

* Step 7
NEWLOG DSNAME=MQM2.LOGCOPY1.DS01,COPY1
NEWLOG DSNAME=MQM2.LOGCOPY1.DS02,COPY1
NEWLOG DSNAME=MQM2.LOGCOPY1.DS03,COPY1
NEWLOG DSNAME=MQM2.LOGCOPY1.DS04,COPY1
NEWLOG DSNAME=MQM2.LOGCOPY2.DS01,COPY2
NEWLOG DSNAME=MQM2.LOGCOPY2.DS02,COPY2
NEWLOG DSNAME=MQM2.LOGCOPY2.DS03,COPY2
NEWLOG DSNAME=MQM2.LOGCOPY2.DS04,COPY2

* Step 8
CRESTART CREATE,ENDRBA=063000

```

Figure 65. Sample input statements for CSQJU003

The things you need to consider for restarting the channel initiator at the recovery site are like those faced when using ARM to restart the channel initiator on a different z/OS image. See “Using ARM in a WebSphere MQ network” on page 342 for more information. Your recovery strategy should also cover recovery of the WebSphere MQ product libraries and the application programming environments that use WebSphere MQ (CICS, for example).

Other functions of the change log inventory utility (CSQJU003) can also be used in disaster recovery scenarios. The HIGHRBA function allows the update of the highest RBA written and highest RBA offloaded values within the bootstrap data set. The CHECKPT function allows the addition of new checkpoint queue records or the deletion of existing checkpoint queue records in the BSDS.

Attention: These functions might affect the integrity of your WebSphere MQ data. Only use them in disaster recovery scenarios under the guidance of IBM service personnel.

Fast copy techniques

If copies of all the page sets and logs are made while the queue manager is frozen, the copies will be a consistent set that can be used to restart the queue manager at an alternative site. They typically enable a much faster restart of the queue manager, as there is little media recovery to be performed.

Use the SUSPEND QMGR LOG command to freeze the queue manager. This command flushes buffer pools to the page sets, takes a checkpoint, and stops any further log write activity. Once log write activity has been suspended, the queue manager is effectively frozen until you issue a RESUME QMGR LOG command. While the queue manager is frozen, the page sets and logs can be copied.

By using copying tools such as FLASHCOPY or SNAPSHOT to rapidly copy the page sets and logs, the time during which the queue manager is frozen can be reduced to a minimum.

Within a queue-sharing group, however, the SUSPEND QMGR LOG command might not be such a good solution. To be effective, the copies of the logs must all contain the same point in time for recovery, which means that the SUSPEND QMGR LOG command must be issued on all queue managers within the queue-sharing group simultaneously, and therefore the entire queue-sharing group will be frozen for some time.

Recovering a queue-sharing group

In the event of a prime site disaster, you can restart a queue-sharing group at a remote site using backup data sets from the prime site. To recover a queue-sharing group you need to coordinate the recovery across all the queue managers in the queue-sharing group, and coordinate with other resources, primarily Db2. This section describes these tasks in detail.

- CF structure media recovery
- Backing up the queue-sharing group at the prime site
- Recovering a queue-sharing group at the alternative site

CF structure media recovery

Media recovery of a CF structure used to hold persistent messages on a shared queue, relies on having a backup of the media that can be forward recovered by the application of logged updates. Take backups of your CF structures periodically using the MQSC `BACKUP CFSTRUCT` command. All updates to shared queues (**MQGETs** and **MQPUTs**) are written on the log of the queue manager where the update is performed. To perform media recovery of a CF structure you must apply logged updates to that backup from the logs of **all** the queue managers that have used that CF structure. When you use the MQSC `RECOVER CFSTRUCT` command, WebSphere MQ automatically merges the logs from relevant queue managers, and applies the updates to the most recent backup.

The CF structure backup is written to the log of the queue manager that processed the `BACKUP CFSTRUCT` command, so there are no additional data sets to be collected and transported to the alternative site.

Backing up the queue-sharing group at the prime site

At the prime site you need to establish a consistent set of backups on a regular basis, which can be used in the event of a disaster to rebuild the queue-sharing group at an alternative site. For a single queue manager, recovery can be to an arbitrary point in time, typically the end of the logs available at the remote site. However, where persistent messages have been stored on a shared queue, the logs of all the queue managers in the queue-sharing group must be merged to recover shared queues, as any queue manager in the queue-sharing group might have performed updates (**MQPUTs** or **MQGETs**) on the queue.

For recovery of a queue-sharing group, you need to establish a point in time that is within the log range of the log data of all queue managers. However, as you can only **forward** recover media from the log, this point in time must be after the `BACKUP CFSTRUCT` command has been issued and after any page set backups have been performed. (Typically, the point in time for recovery might correspond to the end of a business day or week.)

The following diagram shows time lines for two queue managers in a queue-sharing group. For each queue manager, fuzzy backups of page sets are taken (see Method 2: Fuzzy backup). On queue manager A, a `BACKUP CFSTRUCT` command is issued. Subsequently, an `ARCHIVE LOG` command is issued on each queue manager to truncate the active log, and copy it to media offline from the queue manager, which can be transported to the alternative site. 'End of log' identifies the time at which the `ARCHIVE LOG` command was issued, and therefore marks the extent of log data typically available at the alternative site. The point in time for recovery must lie between the end of any page set or CF structure backups, and the earliest end of log available at the alternative site.

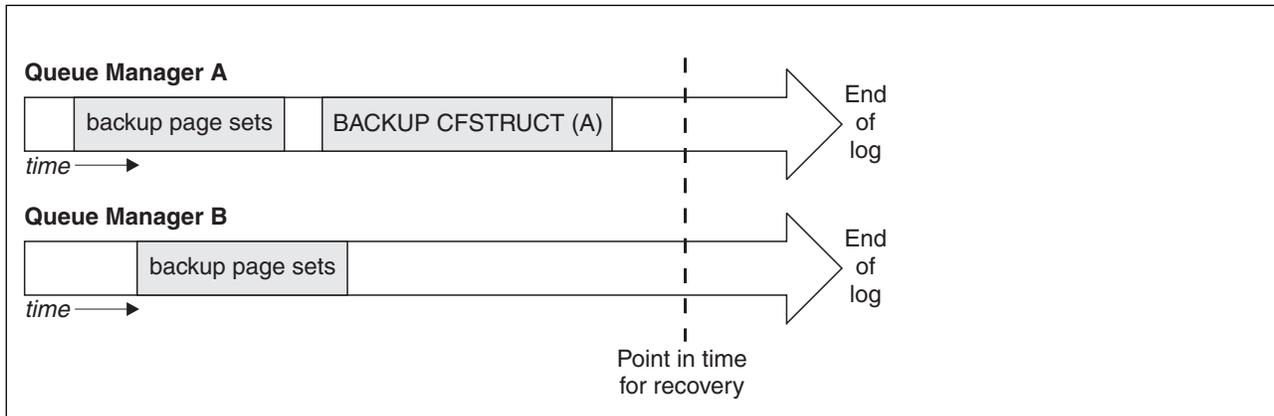


Figure 66. Point in time for recovery for 2 queue managers in a queue-sharing group

WebSphere MQ records information associated with the CF structure backups in a table in Db2. Depending on your requirements, you might want to coordinate the point in time for recovery of WebSphere MQ with that for Db2, or it might be sufficient to take a copy of the WebSphere MQ CSQ.ADMIN_B_STRBACKUP table after the BACKUP CFSTRUCT commands have finished.

To prepare for a recovery:

1. Create page set backups for each queue manager in the queue-sharing group.
2. Issue a BACKUP CFSTRUCT command for each CF structure with the RECOVER(YES) attribute. You can issue these commands from a single queue manager, or from different queue managers within the queue-sharing group to balance the workload.
3. Once all the backups have completed, issue an ARCHIVE LOG command to switch the active log and create copies of the logs and BSDSs of each queue manager in the queue-sharing group.
4. Transport the page set backups, the archived logs, the archived BSDSs of all the queue managers in the queue-sharing group, and your chosen Db2 backup information, off-site.

Recovering a queue-sharing group at the alternative site

Before you can recover the queue-sharing group, you need to prepare the environment:

1. If you have old information in your coupling facility from practice startups when you installed the queue-sharing group, you need to clean this out first (if you do not have old information in the coupling facility, you can omit this step:
 - a. Enter the following z/OS command to display the CF structures for this queue-sharing group:

```
D XCF,STRUCTURE,STRNAME=qsgname
```

- b. For all structures that start with the queue-sharing group name, use the z/OS command SETXCF FORCE CONNECTION to force the connection off those structures:

```
SETXCF FORCE,CONNECTION,STRNAME=strname,CONNAME=ALL
```

c. Delete all the CF structures using the following command for each structure:

```
SETXCF FORCE,STRUCTURE,STRNAME=strname
```

2. Restore Db2 systems and data-sharing groups.
3. Recover the CSQ.ADMIN_B_STRBACKUP table so that it contains information about the most recent structure backups taken at the prime site.

Note: It is important that the STRBACKUP table contains the most recent structure backup information. Older structure backup information might require data sets that you have discarded as a result of the information given by a recent DISPLAY USAGE TYPE(DATASET) command, which would mean that your recovered CF structure would not contain accurate information.

4. Run the ADD QMGR command of the CSQ5PQSG utility for every queue manager in the queue-sharing group. This will restore the XCF group entry for each queue manager.

To recover the queue managers in the queue-sharing group:

1. Define new page set data sets and load them with the data in the copies of the page sets from the primary site.
2. Define new active log data sets.
3. Define a new BSDS data set and use Access Method Services REPRO to copy the *most recent* archived BSDS into it.
4. Use the print log map utility CSQJU004 to print information from this most recent BSDS. At the time this BSDS was archived, the most recent archived log you have would have just been truncated as an active log, and does not appear as an archived log. Record the STARTRBA, STARTLRSN, ENDRBA, and ENDLRSN values of this log.
5. Use the change log inventory utility, CSQJU003, to register this latest archive log data set in the BSDS that you have just restored, using the values recorded in Step 4.
6. Use the DELETE option of CSQJU003 to remove all active log information from the BSDS.
7. Use the NEWLOG option of CSQJU003 to add active logs to the BSDS, do not specify STARTRBA or ENDRBA.
8. Calculate the *recoverylrsn* for the queue-sharing group. The *recoverylrsn* is the lowest of the ENDLRSNs across all queue managers in the queue-sharing group (as recorded in Step 4), minus 1. For example, if there are two queue managers in the queue-sharing group, and the ENDLRSN for one of them is B713 3C72 22C5, and for the other is B713 3D45 2123, the *recoverylrsn* is B713 3C72 22C4.
9. Use CSQJU003 to add a restart control record to the BSDS. Specify:

```
CRESTART CREATE,ENDLRN=recoverylrsn
```

where *recoverylrsn* is the value you recorded in Step 8 on page 336.

The BSDS now describes all active logs as being empty, all the archived logs you have available, and no checkpoints beyond the end of your logs.

You must add the CRESTART record to the BSDS for each queue manager within the queue-sharing group.

10. Restart each queue manager in the queue-sharing group with the START QMGR command. During initialization, an operator reply message such as the following is issued:

```
CSQJ245D +CSQ1 RESTART CONTROL INDICATES TRUNCATION AT RBA highrba.  
REPLY Y TO CONTINUE, N TO CANCEL
```

Reply Y to start the queue manager. The queue manager starts, and recovers data up to ENDRBA specified in the CRESTART statement.

At WebSphere MQ Version 7.0.1 and later, the first queue manager started can rebuild the admin structure partitions for other members of the queue sharing group as well as its own, and it is no longer necessary to restart each queue manager in the queue sharing group at this stage.

11. When the admin structure data for all queue managers has been rebuilt, issue a RECOVER CFSTRUCT command for each CF application structure.

If you issue the RECOVER CFSTRUCT command for all structures on a single queue manager, the log merge process is only performed once, so is quicker than issuing the command on a different queue manager for each CF structure, where each queue manager has to perform the log merge step.

When conditional restart processing is used in a queue sharing group, WebSphere MQ Version 7.0.1 and later queue managers, performing peer admin rebuild, check that peers BSDS contain the same CRESTART LRSN as their own. This is to ensure the integrity of the rebuilt admin structure. It is therefore important to restart other peers in the QSG, so they can process their own CRESTART information, before the next unconditional restart of any member of the group.

Using disk mirroring

Many installations now use disk mirroring technologies such as IBM's Metro Mirror (formerly PPRC) to make synchronous copies of data sets at an alternative site. In such situations, many of the steps detailed become unnecessary as the WebSphere MQ page sets and logs at the alternative site are effectively identical to those at the prime site. Where such technologies are used, the steps to restart a queue sharing group at an alternative site may be summarized as:

- Clear WebSphere MQ CF structures at the alternative site. (These often contain residual information from any previous disaster recovery exercise).
- Restore Db2 systems and all tables in the database used by the WebSphere MQ queue sharing group.
- Restart queue managers. Before WebSphere MQ Version 7.0.1, it is necessary to restart each queue manager defined in the queue sharing group as each queue manager recovers its own partition of the admin structure during queue manager restart. After each queue manager has been restarted, those not on their 'home' LPAR can be shut down again. For WebSphere MQ Version 7.0.1 and later, the first queue manager started rebuilds the admin structure partitions for other members of the queue sharing group as well as its own, and it is no longer necessary to restart each queue manager in the queue sharing group.
- After the admin structure has been rebuilt, recover the application structures.

Reinitializing a queue manager

If the queue manager has terminated abnormally you might not be able to restart it. This could be because your page sets or logs have been lost, truncated, or corrupted. If this has happened, you might have to reinitialize the queue manager (perform a cold start).

Attention

Only perform a cold start if you cannot restart the queue manager any other way. Performing a cold start enables you to recover your queue manager and your object definitions; you will **not** be able to recover your message data. Check that none of the other restart scenarios described in this topic work for you before you do this.

When you have restarted, all your WebSphere MQ objects are defined and available for use, but there is no message data.

Note: Do not reinitialize a queue manager while it is part of a cluster. You must first remove the queue manager from the cluster (using RESET CLUSTER commands on the other queue managers in the cluster), then reinitialize it, and finally reintroduce it to the cluster as a new queue manager.

This is because during reinitialization, the queue manager identifier (QMID) is changed, so any cluster object with the old queue manager identifier must be removed from the cluster.

For further information see the following sections:

- Reinitializing a queue manager that is not in a queue-sharing group
- Reinitializing queue managers in a queue-sharing group

Reinitializing a queue manager that is not in a queue-sharing group

To reinitialize a queue manager, follow this procedure:

1. Prepare the object definition statements that to be used when you restart the queue manager. To do this, either:
 - If page set zero is available, use the CSQUTIL SDEFS function (see  Producing a list of WebSphere MQ define commands (SDEFS) (*WebSphere MQ V7.1 Reference*)). You must get definitions for all object types (authentication information objects, CF structures, channels, namelists, processes, queues, and storage classes).
 - If page set zero is not available, use the definitions from the last time you backed up your object definitions.
2. Redefine your queue manager data sets (do not do this until you have completed step 1 on page 325).
3. Restart the queue manager using the newly defined and initialized log data sets, BSDS, and page sets. Use the object definition input statements that you created in step 1 on page 325 as input in the CSQINP2 initialization input data set.

Reinitializing queue managers in a queue-sharing group

In a queue-sharing group, reinitializing a queue manager is more complex. It might be necessary to reinitialize one or more queue managers because of page set or log problems, but there might also be problems with Db2 or the coupling facility to deal with. Because of this, there are a number of alternatives:

Cold start

Reinitializing the entire queue-sharing group involves forcing all the coupling facilities structures, clearing all object definitions for the queue-sharing group from Db2, deleting or redefining the logs and BSDS, and formatting page sets for all the queue managers in the queue-sharing group.

Shared definitions retained

Delete or redefine the logs and BSDS, format page sets for all queue managers in the queue-sharing group, and force all the coupling facilities structures. On restart, all messages will have been deleted. The queue managers re-create COPY objects that correspond to GROUP objects that still exist in the Db2 database. Any shared queues still exist and can be used.

Single queue manager reinitialized

Delete or redefine the logs and BSDS, and format page sets for the single queue manager (this deletes all its private objects and messages). On restart, the queue manager re-creates COPY objects that correspond to GROUP objects that still exist in the Db2 database. Any shared queues still exist, as do the messages on them, and can be used.

Point in time recovery of a queue-sharing group

This is the alternative site disaster recovery scenario.

Shared objects are recovered to the point in time achieved by Db2 recovery (described in “A Db2 system fails” on page 1374). Each queue manager can be recovered to a point in time achievable from the backup copies available at the alternative site.

Persistent messages can be used in queue-sharing groups, and can be recovered using the MQSC RECOVER CFSTRUCT command. Note that this command recovers to the time of failure. However, there is no recovery of nonpersistent shared queue messages; they are lost unless you have made backup copies independently using the COPY function of the CSQUTIL utility program.

It is not necessary to try to restore each queue manager to the same point in time because there are no interdependencies between the local objects on different queue managers (which are what is actually being recovered), and the queue manager resynchronization with Db2 on restart creates or deletes COPY objects as necessary on a queue manager by queue manager basis.

Using the z/OS Automatic Restart Manager (ARM)

Use this topic to understand how you can use ARM to automatically restart your queue managers.

This section contains information about the following topics:

- “What is the ARM?”
- “ARM policies” on page 340
- “Using ARM in a WebSphere MQ network” on page 342

What is the ARM?

The z/OS Automatic Restart Manager (ARM) is a z/OS recovery function that can improve the availability of your queue managers. When a job or task fails, or the system on which it is running fails, ARM can restart the job or task without operator intervention.

If a queue manager or a channel initiator has failed, ARM restarts it on the same z/OS image. If z/OS, and hence a whole group of related subsystems and applications have failed, ARM can restart all the failed systems automatically, in a predefined order, on another z/OS image within the sysplex. This is called a *cross-system restart*.

Restart the channel initiator by ARM only in exceptional circumstances. If the queue manager is restarted by ARM, restart the channel initiator from the CSQINP2 initialization data set (see “Using ARM in a WebSphere MQ network” on page 342).

You can use ARM to restart a queue manager on a different z/OS image within the sysplex in the event of z/OS failure. The network implications of WebSphere MQ ARM restart on a different z/OS image are described in “Using ARM in a WebSphere MQ network” on page 342.

To enable automatic restart:

- Set up an ARM couple data set.
- Define the automatic restart actions that you want z/OS to perform in an *ARM policy*.
- Start the ARM policy.

Also, WebSphere MQ must register with ARM at startup (this happens automatically).

Note: If you want to restart queue managers in different z/OS images automatically, you must define every queue manager as a subsystem in each z/OS image on which that queue manager might be restarted, with a sysplex wide unique four character subsystem name.

ARM couple data sets

Ensure that you define the couple data sets required for ARM, and that they are online and active before you start any queue manager for which you want ARM support. WebSphere MQ automatic ARM registration fails if the couple data sets are not available at queue manager startup. In this situation, WebSphere MQ assumes that the absence of the couple data set means that you do not want ARM support, and initialization continues.

See the *z/OS MVS Setting up a Sysplex* manual for information about ARM couple data sets.

ARM policies:

The Automatic Restart Manager policies are user-defined rules that control ARM functions that can control any restarts of a queue manager.

ARM functions are controlled by a user-defined *ARM policy*. Each z/OS image running a queue manager instance that is to be restarted by ARM must be connected to an ARM couple data set with an active ARM policy.

IBM provides a default ARM policy. You can define new policies, or override the policy defaults by using the *administrative data utility* (IXCMIAPU) provided with z/OS. The *z/OS MVS Setting up a Sysplex* manual describes this utility, and includes full details of how to define an ARM policy.

Figure 67 shows an example of an ARM policy. This sample policy restarts any queue manager within a sysplex, if either the queue manager failed, or a whole system failed.

```
//IXCMIAPU EXEC PGM=IXCMIAPU,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
  DATA TYPE(ARM)
DEFINE POLICY NAME(ARMPOL1) REPLACE(YES)
  RESTART_GROUP(DEFAULT)
    ELEMENT(*)
      RESTART_ATTEMPTS(0) /* Jobs not to be restarted by ARM */
  RESTART_GROUP(GROUP1)
    ELEMENT(SYSMQGRMQ*) /* These jobs to be restarted by ARM */
/*
```

Figure 67. Sample ARM policy

For more information see:

- Defining an ARM policy
- Activating an ARM policy
- Registering with ARM

Defining an ARM policy

Set up your ARM policy as follows:

- Define RESTART_GROUPS for each queue manager instance that also contain any CICS or IMS subsystems that connect to that queue manager instance. If you use a subsystem naming convention, you might be able to use the '?' and '*' wild-card characters in your element names to define RESTART_GROUPS with minimum definition effort.
- Specify TERMTYPE(ELEMTerm) for your channel initiators to indicate that they will be restarted only if the channel initiator has failed and the z/OS image has not failed.
- Specify TERMTYPE(ALLTERM) for your queue managers to indicate that they will be restarted if either the queue manager has failed or the z/OS image has failed.
- Specify RESTART_METHOD(BOTH, PERSIST) for both queue managers and channel initiators. This tells ARM to restart using the JCL it saved (after resolution of system symbols) during the last startup. It tells ARM to do this irrespective of whether the individual element failed, or the z/OS image failed.
- Accept the default values for all the other ARM policy options.

Activating an ARM policy

To start your automatic restart management policy, issue the following z/OS command:

```
SETXCF START,POLICY,TYPE=ARM,POLNAME=mypol
```

When the policy is started, all systems connected to the ARM couple data set use the same active policy.

Use the SETXCF STOP command to disable automatic restarts.

Registering with ARM

WebSphere MQ registers automatically as an *ARM element* during queue manager startup (subject to ARM availability). It deregisters during its shutdown phase, unless requested not to.

At startup, the queue manager determines whether ARM is available. If it is, WebSphere MQ registers using the name SYSMQMGR*ssid*, where *ssid* is the four character queue manager name, and SYSMQMGR is the element type.

The STOP QMGR MODE(QUIESCE) and STOP QMGR MODE(FORCE) commands deregister the queue manager from ARM (if it was registered with ARM at startup). This prevents ARM restarting this queue manager. The STOP QMGR MODE(RESTART) command does not deregister the queue manager from ARM, so it is eligible for immediate automatic restart.

Each channel initiator address space determines whether ARM is available, and if so registers with the element name SYSMQCH*ssid*, where *ssid* is the queue manager name, and SYSMQCH is the element type.

The channel initiator is always deregistered from ARM when it stops normally, and remains registered only if it ends abnormally. The channel initiator is always deregistered if the queue manager fails.

Using ARM in a WebSphere MQ network:

You can set up your queue manager so that the channel initiators and associated listeners are started automatically when the queue manager is restarted.

To ensure fully automatic queue manager restart on the same z/OS image for both LU 6.2 and TCP/IP communication protocols:

- Start your listeners automatically by adding the appropriate START LISTENER command to the CSQINPX data set.
- Start your channel initiator automatically by adding the appropriate START CHINIT command to the CSQINP2 data set.

For restarting a queue manager with TCP/IP or LU6.2, see

- “Restarting on a different z/OS image with TCP/IP”
- “Restarting on a different z/OS image with LU 6.2” on page 343

See  Task 13: Customize the initialization input data sets (*WebSphere MQ V7.1 Installing Guide*) for information about the CSQINP2 and CSQINPX data sets.

Restarting on a different z/OS image with TCP/IP

If you are using TCP/IP as your communication protocol, and you are using virtual IP addresses, you can configure these to recover on other z/OS images, allowing channels connecting to that queue manager to reconnect without any changes. Otherwise, you can reallocate a TCP/IP address after moving a queue manager to a different z/OS image only if you are using clusters or if you are connecting to a queue-sharing group using a WLM dynamic Domain Name System (DNS) logical group name.

- When using clustering
- When connecting to a queue-sharing group

When using clustering

z/OS ARM responds to a system failure by restarting the queue manager on a different z/OS image in the same sysplex; this system has a different TCP/IP address to the original z/OS image. The following explains how you can use WebSphere MQ clusters to reassign a queue manager's TCP/IP address after it has been moved by ARM restart to a different z/OS image.

When a client queue manager detects the queue manager failure (as a channel failure), it responds by reallocating suitable messages on its cluster transmission queue to a different server queue manager that hosts a different instance of the target cluster queue. However, it cannot reallocate messages that are bound to the original server by affinity constraints, or messages that are in doubt because the server queue manager failed during end-of-batch processing. To process these messages, do the following:

1. Allocate a different cluster-receiver channel name and a different TCP/IP port to each z/OS queue manager. Each queue manager needs a different port so that two systems can share a single TCP/IP stack on a z/OS image. One of these is the queue manager originally running on that z/OS image, and the other is the queue manager that ARM will restart on that z/OS image following a system failure. Configure each port on each z/OS image, so that ARM can restart any queue manager on any z/OS image.
2. Create a different channel initiator command input file (CSQINPX) for each queue manager and z/OS image combination, to be referenced during channel initiator startup.

Each CSQINPX file must include a START LISTENER PORT(port) command specific to that queue manager, and an ALTER CHANNEL command for a cluster-receiver channel specific to that queue manager and z/OS image combination. The ALTER CHANNEL command needs

to set the connection name to the TCP/IP name of the z/OS image on which it is restarted. It must include the port number specific to the restarted queue manager as part of the connection name.

The start-up JCL of each queue manager can have a fixed data set name for this CSQINPX file, and each z/OS image must have a different version of each CSQINPX file on a non-shared DASD volume.

If an ARM restart occurs, WebSphere MQ advertises the changed channel definition to the cluster repository, which in turn publishes it to all the client queue managers that have expressed an interest in the server queue manager.

The client queue manager treats the server queue manager failure as a channel failure, and tries to restart the failed channel. When the client queue manager learns the new server connection-name, the channel restart reconnects the client queue manager to the restarted server queue manager. The client queue manager can then resynchronize its messages, resolve any in-doubt messages on the client queue manager's transmission queue, and normal processing can continue.

When connecting to a queue-sharing group

When connecting to a queue-sharing group through a TCP/IP dynamic Domain Name System (DNS) logical group name, the connection name in your channel definition specifies the logical group name of your queue-sharing group, not the host name or IP address of a physical machine. When this channel starts, it connects to the dynamic DNS and is then connected to one of the

queue managers in the queue-sharing group. This process is explained in  [Setting up communication for WebSphere MQ for z/OS using queue-sharing groups](#) (*WebSphere MQ V7.1 Installing Guide*).

In the unlikely event of an image failure, one of the following occurs:

- The queue managers on the failing image de-register from the dynamic DNS running on your sysplex. The channel responds to the connection failure by entering RETRYING state and then connects to the dynamic DNS running on the sysplex. The dynamic DNS allocates the inbound request to one of the remaining members of the queue-sharing group that is still running on the remaining images.
- If no other queue manager in the queue-sharing group is active and ARM restarts the queue manager and channel initiator on a different image, the group listener registers with dynamic DNS from this new image. This means that the logical group name (from the connection name field of the channel) connects to the dynamic DNS and is then connected to the same queue manager, now running on a different image. No change was required to the channel definition.

For this type of recovery to occur, the following points must be noted:

- On z/OS, the dynamic DNS runs on one of the z/OS images in the sysplex. If this image were to fail, the dynamic DNS needs to be configured so that there is a secondary name server active in the sysplex, acting as an alternative to the primary name server. Information about primary and secondary dynamic DNS servers can be found in the *OS/390 SecureWay CS IP Configuration* manual.
- The TCP/IP group listener might have been started on a particular IP address that might not be available on this z/OS image. If so, the listener might need to be started on a different IP address on the new image. If you are using virtual IP addresses, you can configure these to recover on other z/OS images so that no change to the START LISTENER command is required.

Restarting on a different z/OS image with LU 6.2

If you use only LU 6.2 communication protocols, carry out the following procedure to enable network reconnect after automatic restart of a queue manager on a different z/OS image within the sysplex:

- Define each queue manager within the sysplex with a unique subsystem name.

- Define each channel initiator within the sysplex with a unique LUNAME. This is specified in both the queue manager attributes and in the START LISTENER command.

Note: The LUNAME names an entry in the APPC side table, which in turn maps this to the actual LUNAME.

- Set up a shared APPC side table, which is referenced by each z/OS image within the sysplex. This should contain an entry for each channel initiator's LUNAME. See the *MVS Planning: APPC/MVS Management* manual for information about this.
- Set up an APPCPMxx member of SYS1.PARMLIB for each channel initiator within the sysplex to contain an LUADD to activate the APPC side table entry for that channel initiator. These members should be shared by each z/OS image. The appropriate SYS1.PARMLIB member is activated by a z/OS command SET APPC=xx, which is issued automatically during ARM restart of the queue manager (and its channel initiator) on a different z/OS image, as described in the following text.
- Use the LU62ARM queue manager attribute to specify the xx suffix of this SYS1.PARMLIB member for each channel initiator. This causes the channel initiator to issue the required z/OS command SET APPC=xx to activate its LUNAME.

Define your ARM policy so that it restarts the channel initiator only if it fails while its z/OS image stays up; the user ID associated with the XCFAS address space must be authorized to issue the WebSphere MQ command START CHINIT. Do not restart the channel initiator automatically if its z/OS image also fails, instead use commands in the CSQINP2 and CSQINPX data sets to start the channel initiator and listeners.

Recovering units of work manually

You can manually recover units of work CICS, IMS, RRS, or other queue managers in a queue-sharing group. You can use queue manager commands to display the status of the units of work associated with each connection to the queue manager.

This topic contains information about the following subjects:

- “Displaying connections and threads”
- “Recovering CICS units of recovery manually” on page 345
- “Recovering IMS units of recovery manually” on page 349
- “Recovering RRS units of recovery manually” on page 351
- “Recovering units of recovery on another queue manager in the queue-sharing group” on page 352

Displaying connections and threads

You can use the  DISPLAY CONN (*WebSphere MQ V7.1 Reference*) command to get information about connections to queue managers and their associated units of work. You can display active units of work to see what is currently happening, or to see what needs to be terminated to allow the queue manager to shut down, and you can display unresolved units of work to help with recovery.

Active units of work

To display only active units of work, use
 DISPLAY CONN(*) WHERE(UOWSTATE EQ ACTIVE)

Unresolved units of work

An unresolved unit of work, also known as an "in-doubt thread", is one that is in the second pass of the two-phase commit operation. Resources are held in WebSphere MQ on its behalf. To display unresolved units of work, use
 DISPLAY CONN(*) WHERE(UOWSTATE EQ UNRESOLVED)

External intervention is needed to resolve the status of unresolved units of work. This might only involve starting the recovery coordinator (CICS, IMS, or RRS) or might involve more, as described in the following sections.

Recovering CICS units of recovery manually:

Use this topic to understand what happens when the CICS adapter restarts, and then explains how to deal with any unresolved units of recovery that arise.

What happens when the CICS adapter restarts

For background information, see  *Planning on z/OS (WebSphere MQ V7.1 Installing Guide)*.

Whenever a connection is broken, the adapter has to go through a *restart phase* during the *reconnect process*. The restart phase resynchronizes resources. Resynchronization between CICS and WebSphere MQ enables in-doubt units of work to be identified and resolved.

Resynchronization can be caused by:

- An explicit request from the distributed queuing component
- An implicit request when a connection is made to WebSphere MQ

If the resynchronization is caused by connecting to WebSphere MQ, the sequence of events is:

1. The connection process retrieves a list of in-doubt units of work (UOW) IDs from WebSphere MQ.
2. The UOW IDs are displayed on the console in CSQC313I messages.
3. The UOW IDs are passed to CICS.
4. CICS initiates a resynchronization task (CRSY) for each in-doubt UOW ID.
5. The result of the task for each in-doubt UOW is displayed on the console.

You need to check the messages that are displayed during the connect process:

CSQC313I

Shows that a UOW is in doubt.

CSQC400I

Identifies the UOW and is followed by one of these messages:

- CSQC402I or CSQC403I shows that the UOW was resolved successfully (committed or backed out).
- CSQC404E, CSQC405E, CSQC406E, or CSQC407E shows that the UOW was not resolved.

CSQC409I

Shows that all UOWs were resolved successfully.

CSQC408I

Shows that not all UOWs were resolved successfully.

CSQC314I

Warns that UOW IDs highlighted with a * are not resolved automatically. These UOWs must be resolved explicitly by the distributed queuing component when it is restarted.

Figure 68 on page 346 shows an example set of restart messages displayed on the z/OS console.

```

CSQ9022I +CSQ1 CSQYASCP ' START QMGR' NORMAL COMPLETION
+CSQC323I VICIC1 CSQCQCON CONNECT received from TERMID=PB62 TRANID=CKCN
+CSQC303I VICIC1 CSQCCON CSQCSERV loaded. Entry point is 850E8918
+CSQC313I VICIC1 CSQCCON UOWID=VICIC1.A6E5A6F0E2178D25 is in doubt
+CSQC313I VICIC1 CSQCCON UOWID=VICIC1.A6E5A6F055B2AC25 is in doubt
+CSQC313I VICIC1 CSQCCON UOWID=VICIC1.A6E5A6EFFD60D425 is in doubt
+CSQC313I VICIC1 CSQCCON UOWID=VICIC1.A6E5A6F07AB56D22 is in doubt
+CSQC307I VICIC1 CSQCCON Successful connection to subsystem VC2
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008BAD18) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008BAA10) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008BA708) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008CAE88) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008CAB80) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008CA878) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008CA570) connect
successful
+CSQC472I VICIC1 CSQCSERV Server subtask (TCB address=008CA268) connect
successful
+CSQC403I VICIC1 CSQCTRUE Resolved BACKOUT for
+CSQC400I VICIC1 CSQCTRUE UOWID=VICIC1.A6E5A6F0E2178D25
+CSQC403I VICIC1 CSQCTRUE Resolved BACKOUT for
+CSQC400I VICIC1 CSQCTRUE UOWID=VICIC1.A6E5A6F055B2AC25
+CSQC403I VICIC1 CSQCTRUE Resolved BACKOUT for
+CSQC400I VICIC1 CSQCTRUE UOWID=VICIC1.A6E5A6F07AB56D22
+CSQC403I VICIC1 CSQCTRUE Resolved BACKOUT for
+CSQC400I VICIC1 CSQCTRUE UOWID=VICIC1.A6E5A6EFFD60D425
+CSQC409I VICIC1 CSQCTRUE Resynchronization completed successfully

```

Figure 68. Example restart messages

The total number of CSQC313I messages should equal the total number of CSQC402I plus CSQC403I messages. If the totals are not equal, there are UOWs that the connection process cannot resolve. Those UOWs that cannot be resolved are caused by problems with CICS (for example, a cold start) or with WebSphere MQ, or by distributing queuing. When these problems have been fixed, you can initiate another resynchronization by disconnecting and then reconnecting.

Alternatively, you can resolve each outstanding UOW yourself using the RESOLVE INDOUBT command and the UOW ID shown in message CSQC400I. You must then initiate a disconnect and a connect to clean up the *unit of recovery descriptors* in CICS. You need to know the correct outcome of the UOW to resolve UOWs manually.

All messages that are associated with unresolved UOWs are locked by WebSphere MQ and no Batch, TSO, or CICS task can access them.

If CICS fails and an emergency restart is necessary, *do not* vary the GENERIC APPLID of the CICS system. If you do and then reconnect to WebSphere MQ, data integrity with WebSphere MQ cannot be guaranteed. This is because WebSphere MQ treats the new instance of CICS as a different CICS (because the APPLID is different). In-doubt resolution is then based on the wrong CICS log.

How to resolve CICS units of recovery manually

If the adapter ends abnormally, CICS and WebSphere MQ build in-doubt lists either dynamically or during restart, depending on which subsystem caused the abend.

Note: If you use the DFH\$INDB sample program to show units of work, you might find that it does not always show WebSphere MQ UOWs correctly.

When CICS connects to WebSphere MQ, there might be one or more units of recovery that have not been resolved.

One of the following messages is sent to the console:

- CSQC404E
- CSQC405E
- CSQC406E
- CSQC407E
- CSQC408I

For details of what these messages mean, see the  z/OS Messages and Codes (*WebSphere MQ V7.1 Reference*) documentation.

CICS retains details of units of recovery that were not resolved during connection startup. An entry is purged when it no longer appears on the list presented by WebSphere MQ.

Any units of recovery that CICS cannot resolve must be resolved manually using WebSphere MQ commands. This manual procedure is rarely used within an installation, because it is required only where operational errors or software problems have prevented automatic resolution. *Any inconsistencies found during in-doubt resolution must be investigated.*

To resolve the units of recovery:

1. Obtain a list of the units of recovery from WebSphere MQ using the following command:

```
+CSQ1 DISPLAY CONN(*) WHERE(UOWSTATE EQ UNRESOLVED)
```

You receive the following message:

```

CSQM201I +CSQ1 CSQMDRTC DISPLAY CONN DETAILS
CONN(BC85772CBE3E0001)
EXTCONN(C3E2D8C3C7D9F0F940404040404040)
TYPE(CONN)
CONNOPTS(
  MQCNO_STANDARD_BINDING
)
UOWLOGDA(2005-02-04)
UOWLOGTI(10.17.44)
UOWSTDA(2005-02-04)
UOWSTTI(10.17.44)
UOWSTATE(UNRESOLVED)
NID(IYRCSQ1 .BC8571519B60222D)
EXTURID(BC8571519B60222D)
QMURID(0000002BDA50)
URTYPE(CICS)
USERID(MQTEST)
APPLTAG(IYRCSQ1)
ASID(0000)
APPLTYPE(CICS)
TRANSID(GP02)
TASKNO(0000096)
END CONN DETAILS

```

For CICS connections, the NID consists of the CICS applid and a unique number provided by CICS at the time the syncpoint log entries are written. This unique number is stored in records written to both the CICS system log and the WebSphere MQ log at syncpoint processing time. This value is referred to in CICS as the *recovery token*.

2. Scan the CICS log for entries related to a particular unit of recovery.

Look for a PREPARE record for the task-related installation where the recovery token field (JCSRMTKN) equals the value obtained from the network ID. The network ID is supplied by WebSphere MQ in the DISPLAY CONN command output.

The PREPARE record in the CICS log for the units of recovery provides the CICS task number. All other entries on the log for this CICS task can be located using this number.

You can use the CICS journal print utility DFHJUP when scanning the log. For details of using this program, see the *CICS Operations and Utilities Guide*.

3. Scan the WebSphere MQ log for records with the NID related to a particular unit of recovery. Then use the URID from this record to obtain the rest of the log records for this unit of recovery.

When scanning the WebSphere MQ log, note that the WebSphere MQ startup message CSQJ001I provides the start RBA for this session.

The print log records program (CSQ1LOGP) can be used for that purpose.

4. If you need to, do in-doubt resolution in WebSphere MQ.

WebSphere MQ can be directed to take the recovery action for a unit of recovery using a WebSphere MQ RESOLVE INDOUBT command.

For information about RESOLVE INDOUBT, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) documentation.

To recover all threads associated with a specific *connection-name*, use the NID(*) option.

The command produces one of the following messages showing whether the thread is committed or backed out:

```
CSQV414I +CSQ1 THREAD network-id COMMIT SCHEDULED
CSQV415I +CSQ1 THREAD network-id ABORT SCHEDULED
```

When performing in-doubt resolution, CICS and the adapter are not aware of the commands to WebSphere MQ to commit or back out units of recovery, because only WebSphere MQ resources are affected. However, CICS keeps details about the in-doubt threads that could not be resolved by WebSphere MQ. This information is purged either when the list presented is empty, or when the list does not include a unit of recovery of which CICS has details.

Recovering IMS units of recovery manually:

Use this topic to understand what happens when the IMS adapter restarts, and then explains how to deal with any unresolved units of recovery that arise.

What happens when the IMS adapter restarts

For background information, see the  [Planning on z/OS \(WebSphere MQ V7.1 Installing Guide\)](#).

Whenever the connection to WebSphere MQ is restarted, either following a queue manager restart or an IMS /START SUBSYS command, IMS initiates the following resynchronization process:

1. IMS presents the list of unit of work (UOW) IDs that it believes are in doubt to the WebSphere MQ IMS adapter one at a time with a resolution parameter of Commit or Backout.
2. The IMS adapter passes the resolution request to WebSphere MQ and reports the result back to IMS.
3. Having processed all the IMS resolution requests, the IMS adapter gets from WebSphere MQ a list of all UOWs that WebSphere MQ still holds in doubt that were initiated by the IMS system. These are reported to the IMS master terminal in message CSQQ008I.

Note: While a UOW is in doubt, any associated WebSphere MQ message is locked by WebSphere MQ and is not available to any application.

How to resolve IMS units of recovery manually

When IMS connects to WebSphere MQ, WebSphere MQ might have one, or more in-doubt units of recovery that have not been resolved.

If WebSphere MQ has in-doubt units of recovery that IMS did not resolve, the following message is issued at the IMS master terminal:

```
CSQQ008I nn units of recovery are still in doubt in queue manager qmgr-name
```

If this message is issued, IMS was either cold-started or it was started with an incomplete log tape. This message can also be issued if WebSphere MQ or IMS terminates abnormally because of a software error or other subsystem failure.

After receiving the CSQQ008I message:

- The connection remains active.
- IMS applications can still access WebSphere MQ resources.
- Some WebSphere MQ resources remain locked out.

If the in-doubt thread is not resolved, IMS message queues can start to build up. If the IMS queues fill to capacity, IMS terminates. You must be aware of this potential difficulty, and you must monitor IMS until the in-doubt units of recovery are fully resolved.

Recovery procedure

Use the following procedure to recover the IMS units of work:

1. Force the IMS log closed, using /SWI OLDS, and then archive the IMS log. Use the utility, DFSERA10, to print the records from the previous IMS log tape. Type X'3730' log records indicate a phase-2 commit request and type X'38' log records indicate an abort request. Record the requested action for the last transaction in each dependent region.
2. Run the DL/I batch job to back out each PSB involved that has not reached a commit point. The process might take some time because transactions are still being processed. It might also lock up a number of records, which could affect the rest of the processing and the rest of the message queues.
3. Produce a list of the in-doubt units of recovery from WebSphere MQ using the following command:

```
+CSQ1 DISPLAY CONN(*) WHERE(UOWSTATE EQ UNRESOLVED)
```

You receive the following message:

```
CSQM201I +CSQ1 CSQMDRTC  DISPLAY CONN DETAILS
CONN(BC45A794C4290001)
EXTCONN(C3E2D8C3E2C5C3F240404040404040)
TYPE(CONN)
CONNOPTS(
  MQCNO_STANDARD_BINDING
)
UOWLOGDA(2005-02-15)
UOWLOGTI(16.39.43)
UOWSTDA(2005-02-15)
UOWSTTI(16.39.43)
UOWSTATE(UNRESOLVED)
NID(IM8F      .BC45A794D3810344)
EXTURID(
  0000052900000000
)
QMURID(00000354B76E)
URTYPE(IMS)
USERID(STCPI)
APPLTAG(IM8F)
ASID(0000)
APPLTYPE(IMS)
PSTID(0004)
PSBNAME(GP01MPP)
```

For IMS, the NID consists of the IMS connection name and a unique number provided by IMS. The value is referred to in IMS as the *recovery token*. For more information, see the *IMS Customization Guide*.

4. Compare the NIDs (IMSID plus OASN in hexadecimal) displayed in the DISPLAY THREAD messages with the OASNs (4 bytes decimal) shown in the DFSERA10 output. Decide whether to commit or back out.
5. Perform in-doubt resolution in WebSphere MQ with the RESOLVE INDOUBT command, as follows:

```
RESOLVE INDOUBT(connection-name)  
  ACTION(COMMIT|BACKOUT)  
  NID(network-id)
```

For information about the RESOLVE INDOUBT command, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) documentation.

To recover all threads associated with *connection-name*, use the NID(*) option. The command results in one of the following messages to indicate whether the thread is committed or backed out:

```
CSQV414I  THREAD network-id COMMIT SCHEDULED  
CSQV415I  THREAD network-id BACKOUT SCHEDULED
```

When performing in-doubt resolution, IMS and the adapter are not aware of the commands to WebSphere MQ to commit or back out in-doubt units of recovery because only WebSphere MQ resources are affected.

Recovering RRS units of recovery manually:

Use this topic to understand the how to determine if there are in-doubt RRS units of recovery, and how to manually resolve those units of recovery.

When RRS connects to WebSphere MQ, WebSphere MQ might have one, or more in-doubt units of recovery that have not been resolved. If WebSphere MQ has in-doubt units of recovery that RRS did not resolve, one of the following messages is issued at the z/OS console:

- CSQ3011I
- CSQ3013I
- CSQ3014I
- CSQ3016I

Both WebSphere MQ and RRS provide tools to display information about in-doubt units of recovery, and techniques for manually resolving them.

In WebSphere MQ, use the DISPLAY CONN command to display information about in-doubt WebSphere MQ threads. The output from the command includes RRS unit of recovery IDs for those WebSphere MQ threads that have RRS as a coordinator. This can be used to determine the outcome of the unit of recovery.

Use the RESOLVE INDOUBT command to resolve the WebSphere MQ in-doubt thread manually. This command can be used to either commit or back out the unit of recovery after you have determined what the correct decision is.

Recovering units of recovery on another queue manager in the queue-sharing group:

Use this topic to identify, and manually recover units of recovery on other queue managers in a queue-sharing group.

If a queue manager that is a member of a queue-sharing group fails and cannot be restarted, other queue managers in the group can perform peer recovery, and take over from it. However, the queue manager might have in-doubt units of recovery that cannot be resolved by peer recovery because the final disposition of that unit of recovery is known only to the failed queue manager. These units of recovery are resolved when the queue manager is eventually restarted, but until then, they remain in doubt.

This means that certain resources (for example, messages) might be locked, making them unavailable to other queue managers in the group. In this situation, you can use the DISPLAY THREAD command to display these units of work on the inactive queue manager. If you want to resolve these units of recovery manually to make the messages available to other queue managers in the group, you can use the RESOLVE INDOUBT command.

When you issue the DISPLAY THREAD command to display units of recovery that are in doubt, you can use the QMNAME keyword to specify the name of the inactive queue manager. For example, if you issue the following command:

```
+CSQ1 DISPLAY THREAD(*) TYPE(INDOUBT) QMNAME(QM01)
```

You receive the following messages:

```
CSQV436I +CSQ1 INDOUBT THREADS FOR QM01 -  
NAME      THREAD-XREF      URID  NID  
USER1     00000000000000000000000000000000 CSQ:0001.0  
USER2     00000000000000000000000000000000 CSQ:0002.0  
DISPLAY THREAD REPORT COMPLETE
```

If the queue manager specified is active, WebSphere MQ does not return information about in-doubt threads, but issues the following message:

```
CSQV435I CANNOT USE QMNAME KEYWORD, QM01 IS ACTIVE
```

Use the WebSphere MQ command RESOLVE INDOUBT to resolve the in-doubt threads manually. Use the QMNAME keyword to specify the name of the inactive queue manager in the command.

This command can be used to commit or back out the unit of recovery. The command resolves the shared portion of the unit of recovery only; any local messages are unaffected and remain locked until the queue manager restarts, or reconnects to CICS, IMS, or RRS batch.

WebSphere MQ and IMS

WebSphere MQ provides 2 components to interface with IMS, the WebSphere MQ-IMS adapter, and the WebSphere MQ-IMS bridge. These components are commonly called the IMS adapter, and the IMS bridge.

Operating the IMS adapter

Use this topic to understand how to operate the IMS adapter, which connects WebSphere MQ to IMS systems.

This topic describes how to operate the IMS adapter, which connects WebSphere MQ to IMS systems.

Note: The IMS adapter does not incorporate any operations and control panels.

This topic contains the following sections:

- “Controlling IMS connections”
- “Connecting from the IMS control region” on page 354
- “Displaying in-doubt units of recovery” on page 356
- “Controlling IMS dependent region connections” on page 358
- “Disconnecting from IMS” on page 360
- “Controlling the IMS trigger monitor” on page 361

Controlling IMS connections:

Use this topic to understand the IMS operator commands which control and monitor the connection to WebSphere MQ.

IMS provides the following operator commands to control and monitor the connection to WebSphere MQ:

/CHANGE SUBSYS

Deletes an in-doubt unit of recovery from IMS.

/DISPLAY OASN SUBSYS

Displays outstanding recovery elements.

/DISPLAY SUBSYS

Displays connection status and thread activity.

/START SUBSYS

Connects the IMS control region to a queue manager.

/STOP SUBSYS

Disconnects IMS from a queue manager.

/TRACE

Controls the IMS trace.

For more information about these commands, see the *IMS/ESA[®] Operator's Reference* manual for the level of IMS that you are using.

IMS command responses are sent to the terminal from which the command was issued. Authorization to issue IMS commands is based on IMS security.

Connecting from the IMS control region:

Use this topic to understand the mechanisms available to connect from IMS to WebSphere MQ.

IMS makes one connection from its control region to each queue manager that uses IMS. IMS must be enabled to make the connection in one of these ways:

- Automatically during either:
 - A cold start initialization.
 - A warm start of IMS, if the WebSphere MQ connection was active when IMS was shut down.
- In response to the IMS command:

```
/START SUBSYS sysid
```

where *sysid* is the queue manager name.

The command can be issued regardless of whether the queue manager is active.

The connection is not made until the first MQ API call to the queue manager is made. Until that time, the IMS command /DIS SUBSYS shows the status as 'NOT CONN'.

The order in which you start IMS and the queue manager is not significant.

IMS cannot re-enable the connection to the queue manager automatically if the queue manager is stopped with a STOP QMGR command, the IMS command /STOP SUBSYS, or an abnormal end. Therefore, you must make the connection by using the IMS command /START SUBSYS.

Initializing the adapter and connecting to the queue manager

The adapter is a set of modules loaded into the IMS control and dependent regions, using the IMS external Subsystem Attach Facility.

This procedure initializes the adapter and connects to the queue manager:

1. Read the subsystem member (SSM) from IMS.PROCLIB. The SSM chosen is an IMS EXEC parameter. There is one entry in the member for each queue manager to which IMS can connect. Each entry contains control information about a WebSphere MQ adapter.
2. Load the IMS adapter.

Note: IMS loads one copy of the adapter modules for each WebSphere MQ instance that is defined in the SSM member.

3. Attach the external subsystem task for WebSphere MQ.
4. Run the adapter with the CTL EXEC parameter (IMSID) as the connection name.

The process is the same whether the connection is part of initialization or a result of the IMS command /START SUBSYS.

If the queue manager is active when IMS tries to make the connection, the following messages are sent:

- to the z/OS console:

```
DFS3613I ESS TCB INITIALIZATION COMPLETE
```

- to the IMS master terminal:

```
CSQQ000I IMS/TM imsid connected to queue manager ssnm
```

When IMS tries to make the connection and *the queue manager is not active*, the following messages are sent to the IMS master terminal each time an application makes an MQI call:

```
CSQQ001I IMS/TM imsid not connected to queue manager ssnm.  
    Notify message accepted  
DFS3607I MQM1      SUBSYSTEM ID  EXIT FAILURE, FC = 0286, RC = 08,  
JOBNAME = IMSEMPR1
```

If you get DFS3607I messages when you start the connection to IMS or on system startup, this indicates that the queue manager is not available. To prevent a large number of messages being generated, you must do one of the following:

1. Start the relevant queue manager.
2. Issue the IMS command:

```
/STOP SUBSYS
```

so that IMS does not expect to connect to the queue manager.

If you do neither, a DFS3607I message and the associated CSQQ001I message are issued each time a job is scheduled in the region and each time a connection request to the queue manager is made by an application.

Thread attachment

In an MPP or IFP region, IMS makes a thread connection when the first application program is scheduled into that region, even if that application program does not make a WebSphere MQ call. In a BMP region, the thread connection is made when the application makes its first WebSphere MQ call (**MQCONN** or **MQCONNX**). This thread is retained for the duration of the region or until the connection is stopped.

For both the message driven and non-message driven regions, the recovery thread cross-reference identifier, *Thread-xref*, associated with the thread is:

```
PSTid + PSBname
```

where:

PSTid Partition specification table region identifier

PSBname

Program specification block name

You can use connection IDs as unique identifiers in WebSphere MQ commands, in which case WebSphere MQ automatically inserts these IDs into any operator message that it generates.

Displaying in-doubt units of recovery:

You can display and in-doubt of units of recovery and attempt to recover them.

The operational steps used to list and recover in-doubt units of recovery in this topic are for relatively simple cases only. If the queue manager ends abnormally while connected to IMS, IMS might commit or back out work without WebSphere MQ being aware of it. When the queue manager restarts, that work is termed *in doubt*. A decision must be made about the status of the work.

To display a list of in-doubt units of recovery, issue the command:

```
+CSQ1 DISPLAY CONN(*) WHERE(UOWSTATE EQ UNRESOLVED)
```

WebSphere MQ responds with a message like the following:

```
CSQM201I +CSQ1 CSQMDRTC DIS CONN DETAILS
CONN(BC0F6125F5A30001)
EXTCONN(C3E2D8C3C3E2D8F140404040404040)
TYPE(CONN)
CONNOPTS(
  MQCNO_STANDARD_BINDING
)
UOWLOGDA(2004-11-02)
UOWLOGTI(12.27.58)
UOWSTDA(2004-11-02)
UOWSTTI(12.27.58)
UOWSTATE(UNRESOLVED)
NID(CSQ1CHIN.BC0F5F1C86FC0766)
EXTURID(0000000000000001F000000007472616E5F6964547565204E6F762020...)
QMURID(000000026232)
URTYPE(XA)
USERID( )
APPLTAG(CSQ1CHIN)
ASID(0000)
APPLTYPE(CHINIT)
CHANNEL( )
CONNAME( )
END CONN DETAILS
```

For an explanation of the attributes in this message, see the description of the DISPLAY CONN command in  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*).

Recovering in-doubt units of recovery

To recover in-doubt units of recovery, issue this command:

```
+CSQ1 RESOLVE INDOUBT(connection-name) ACTION(COMMIT|BACKOUT)
NID(net-node.number)
```

where:

connection-name

The IMS system ID.

ACTION

Indicates whether to commit (COMMIT) or back out (BACKOUT) this unit of recovery.

net-node.number

The associated net-node.number.

When you have issued the RESOLVE INDOUBT command, one of the following messages is displayed:

```
CSQV414I +CSQ1 THREAD network-id COMMIT SCHEDULED
CSQV415I +CSQ1 THREAD network-id BACKOUT SCHEDULED
```

Resolving residual recovery entries

At given times, IMS builds a list of residual recovery entries (RREs). RREs are units of recovery about which WebSphere MQ might be in doubt. They arise in several situations:

- If the queue manager is not active, IMS has RREs that cannot be resolved until the queue manager is active. These RREs are not a problem.
- If the queue manager is active and connected to IMS, and if IMS backs out the work that WebSphere MQ has committed, the IMS adapter issues message CSQQ010E. If the data in the two systems must be consistent, there is a problem. For information about resolving this problem, see “Recovering IMS units of recovery manually” on page 349.
- If the queue manager is active and connected to IMS, there might still be RREs even though no messages have informed you of this problem. After the WebSphere MQ connection to IMS has been established, you can issue the following IMS command to find out if there is a problem:

```
/DISPLAY OASN SUBSYS sysid
```

To purge the RRE, issue one of the following IMS commands:

```
/CHANGE SUBSYS sysid RESET
/CHANGE SUBSYS sysid RESET OASN nnnn
```

where *nnnn* is the originating application sequence number listed in response to your +CSQ1 DISPLAY command. This is the schedule number of the program instance, giving its place in the sequence of invocations of that program since the last IMS cold start. IMS cannot have two in-doubt units of recovery with the same schedule number.

These commands reset the status of IMS; they do not result in any communication with WebSphere MQ.

Controlling IMS dependent region connections:

You can control, monitor, and, when necessary, terminate connections between IMS and WebSphere MQ.

Controlling IMS dependent region connections involves the following activities:

- Connecting from dependent regions
- Region error options
- Monitoring the activity on connections
- Disconnecting from dependent regions

Connecting from dependent regions

The IMS adapter used in the control region is also loaded into dependent regions. A connection is made from each dependent region to WebSphere MQ. This connection is used to coordinate the commitment of WebSphere MQ and IMS work. To initialize and make the connection, IMS does the following:

1. Reads the subsystem member (SSM) from IMS.PROCLIB.
A subsystem member can be specified on the dependent region EXEC parameter. If it is not specified, the control region SSM is used. If the region is never likely to connect to WebSphere MQ, to avoid loading the adapter, specify a member with no entries.
2. Loads the WebSphere MQ adapter.
For a batch message program, the load is not done until the application issues its first messaging command. At that time, IMS tries to make the connection.
For a message-processing program region or IMS fast-path region, the attempt is made when the region is initialized.

Region error options

If the queue manager is not active, or if resources are not available when the first messaging command is sent from application programs, the action taken depends on the error option specified on the SSM entry. The options are:

- R** The appropriate return code is sent to the application.
- Q** The application ends abnormally with abend code U3051. The input message is re-queued.
- A** The application ends abnormally with abend code U3047. The input message is discarded.

Monitoring the activity on connections

A thread is established from a dependent region when an application makes its first successful WebSphere MQ request. You can display information about connections and the applications currently using them by issuing the following command from WebSphere MQ:

```
+CSQ1 DISPLAY CONN(*)
```

The command produces a message like the following:

```

CONN(BC45A794C4290001)
EXTCONN(C3E2D8C3C3E2D8F140404040404040)
TYPE(CONN)
CONNOPTS(
  MQCNO_STANDARD_BINDING
)
UOWLOGDA(2004-12-15)
UOWLOGTI(16.39.43)
UOWSTDA(2004-12-15)
UOWSTTI(16.39.43)
UOWSTATE(ACTIVE)
NID( )
EXTURID(
  0000052900000000
)
QMURID(00000354B76E)
URTYPE(IMS)
USERID(STCPI)
APPLTAG(IM8F)
ASID(0049)
APPLTYPE(IMS)
PSTID(0004)
PSBNAME(GP01MPP)

```

For the control region, *thread-xref* is the special value CONTROL. For dependent regions, it is the PSTid concatenated with the PSBname. *auth-id* is either the user field from the job card, or the ID from the z/OS started procedures table.

For an explanation of the displayed list, see the description of message CSQV402I in the  z/OS Messages and Codes (*WebSphere MQ V7.1 Reference*) documentation.

IMS provides a display command to monitor the connection to WebSphere MQ. It shows which program is active on each dependent region connection, the LTERM user name, and the control region connection status. The command is:

```
/DISPLAY SUBSYS name
```

The status of the connection between IMS and WebSphere MQ is shown as one of:

```

CONNECTED
NOT CONNECTED
CONNECT IN PROGRESS
STOPPED
STOP IN PROGRESS
INVALID SUBSYSTEM NAME=name
SUBSYSTEM name NOT DEFINED BUT RECOVERY OUTSTANDING

```

The thread status from each dependent region is one of the following:

```
CONN  
CONN, ACTIVE (includes LTERM of user)
```

Disconnecting from dependent regions

To change values in the SSM member of IMS.PROCLIB, you disconnect a dependent region. To do this, you must:

1. Issue the IMS command:

```
/STOP REGION
```

2. Update the SSM member.
3. Issue the IMS command:

```
/START REGION
```

Disconnecting from IMS:

The connection is ended when either IMS or the queue manager terminates. Alternatively, the IMS master terminal operator can explicitly break the connection.

To terminate the connection between IMS and WebSphere MQ, use the following IMS command:

```
/STOP SUBSYS sysid
```

The command sends the following message to the terminal that issued it, typically the master terminal operator (MTO):

```
DFS058I STOP COMMAND IN PROGRESS
```

The IMS command:

```
/START SUBSYS sysid
```

is required to reestablish the connection.

Note: The IMS command /STOP SUBSYS is not completed if an IMS trigger monitor is running.

Controlling the IMS trigger monitor:

You can use the CSQQTRMN transaction to stop, and start the IMS trigger monitor.

The IMS trigger monitor (the CSQQTRMN transaction) is described in the  Setting up the IMS trigger monitor (*WebSphere MQ V7.1 Installing Guide*).

To control the IMS trigger monitor see:

- Starting CSQQTRMN
- Stopping CSQQTRMN

Starting CSQQTRMN

1. Start a batch-oriented BMP that runs the program CSQQTRMN for each initiation queue you want to monitor.
2. Modify your batch JCL to add a DDname of CSQQUT1 that points to a data set containing the following information:

QMGRNAME=q_manager_name	Comment: queue manager name
INITQUEUEUENAME=init_q_name	Comment: initiation queue name
LTERM=lterm	Comment: LTERM to remove error messages
CONSOLEMESSAGES=YES	Comment: Send error messages to console

where:

q_manager_name	The name of the queue manager (if this is blank, the default nominated in CSQQDEFV is assumed)
init_q_name	The name of the initiation queue to be monitored
lterm	The IMS LTERM name for the destination of error messages (if this is blank, the default value is MASTER).
CONSOLEMESSAGES=YES	Requests that messages sent to the nominated IMS LTERM are also sent to the z/OS console. If this parameter is omitted or misspelled, the default is NOT to send messages to the console.

3. Add a DD name of CSQQUT2 if you want a printed report of the processing of CSQQUT1 input.

Note:

1. The data set CSQQUT1 is defined with LRECL=80. Other DCB information is taken from the data set. The DCB for data set CSQQUT2 is RECFM=VBA and LRECL=125.
2. You can put only one keyword on each record. The keyword value is delimited by the first blank following the keyword; this means that you can include comments. An asterisk in column 1 means that the whole input record is a comment.
3. If you misspell either of the QMGRNAME or LTERM keywords, CSQQTRMN uses the default for that keyword.
4. Ensure that the subsystem is started in IMS (by the /START SUBSYS command) before submitting the trigger monitor BMP job. If it is not started, your trigger monitor job terminates with abend code U3042.

Stopping CSQQTRMN

Once started, CSQQTRMN runs until either the connection between WebSphere MQ and IMS is broken due to one of the following events:

- the queue manager ending

- IMS ending
- or a z/OS STOP **jobname** command is entered.

Controlling the IMS bridge

Use this topic to understand the IMS commands that you can use to control the IMS bridge.

There are no WebSphere MQ commands to control the WebSphere MQ-IMS bridge. However, you can stop messages being delivered to IMS in the following ways:

- For non-shared queues, by using the ALTER QLOCAL(*xxx*) GET(DISABLED) command for all bridge queues.
- For clustered queues, by using the SUSPEND QMGR CLUSTER(*xxx*) command. This is effective only when another queue manager is also hosting the clustered bridge queue.
- For clustered queues, by using the SUSPEND QMGR FACILITY(IMSBRIDGE) command. No further messages are sent to IMS, but the responses for any outstanding transactions are received from IMS. To start sending messages to IMS again, issue the RESUME QMGR FACILITY(IMSBRIDGE) command.

You can also use the MQSC command DISPLAY SYSTEM to display whether the bridge is suspended.

See  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for details of these commands.

For further information see:

- “Starting and stopping the IMS bridge”
- “Controlling IMS connections” on page 363
- Controlling bridge queues
- “Resynchronizing the IMS bridge” on page 364
- Working with tpipe names
- Deleting messages from IMS
- Deleting tpipes
- “IMS Transaction Expiration” on page 366

Starting and stopping the IMS bridge

Start the WebSphere MQ bridge by starting OTMA. Either use the IMS command:

```
/START OTMA
```

or start it automatically by specifying OTMA=YES in the IMS system parameters. If OTMA is already started, the bridge starts automatically when queue manager startup has completed. A WebSphere MQ event message is produced when OTMA is started.

Use the IMS command:

```
/STOP OTMA
```

to stop OTMA communication. When this command is issued, a WebSphere MQ event message is produced.

Controlling IMS connections

IMS provides these operator commands to control and monitor the connection to WebSphere MQ:

/DEQUEUE TMEMBER *tmember* **TPIPE** *tpipe*

Removes messages from a Tpipe. Specify PURGE to remove all messages or PURGE1 to remove the first message only.

/DISPLAY OTMA

Displays summary information about the OTMA server and clients, and client status.

/DISPLAY TMEMBER *name*

Displays information about an OTMA client.

/DISPLAY TRACE TMEMBER *name*

Displays information about what is being traced.

/SECURE OTMA

Sets security options.

/START OTMA

Enables communications through OTMA.

/START TMEMBER *tmember* **TPIPE** *tpipe*

Starts the named Tpipe.

/STOP OTMA

Stops communications through OTMA.

/STOP TMEMBER *tmember* **TPIPE** *tpipe*

Stops the named Tpipe.

/TRACE

Controls the IMS trace.

For more information about these commands, see the *IMS/ESA Operators Reference* manual for the level of IMS that you are using.

IMS command responses are sent to the terminal from which the command was issued. Authorization to issue IMS commands is based on IMS security.

Controlling bridge queues

To stop communicating with the queue manager with XCF member name *tmember* through the bridge, issue the following IMS command:

```
/STOP TMEMBER tmember TPIPE ALL
```

To resume communication, issue the following IMS command:

```
/START TMEMBER tmember TPIPE ALL
```

The Tpipes for a queue can be displayed using the MQ DISPLAY QUEUE command.

To stop communication with the queue manager on a single Tpipe, issue the following IMS command:

```
/STOP TMEMBER tmember TPIPE tpipe
```

One or two Tpipes are created for each active bridge queue, so issuing this command stops communication with the WebSphere MQ queue. To resume communication, use the following IMS command :

```
/START TMEMBER tmember TPIPE tpipe
```

Alternatively, you can alter the attributes of the WebSphere MQ queue to make it get inhibited.

Resynchronizing the IMS bridge

The IMS bridge is automatically restarted whenever the queue manager, IMS, or OTMA are restarted.

The first task undertaken by the IMS bridge is to resynchronize with IMS. This involves WebSphere MQ and IMS checking sequence numbers on every synchronized Tpipe. A synchronized Tpipe is used when persistent messages are sent to IMS from a WebSphere MQ-IMS bridge queue using commit mode zero (commit-then-send).

If the bridge cannot resynchronize with IMS, the IMS sense code is returned in message CSQ2023E and the connection to OTMA is stopped. If the bridge cannot resynchronize with an individual IMS Tpipe, the IMS sense code is returned in message CSQ2025E and the Tpipe is stopped. If a Tpipe has been cold started, the recoverable sequence numbers are automatically reset to 1.

If the bridge discovers mismatched sequence numbers when resynchronizing with a Tpipe, message CSQ2020E is issued. Use the WebSphere MQ command RESET TPIPE to initiate resynchronization with the IMS Tpipe. You need to provide the XCF group and member name, and the name of the Tpipe; this information is provided by the message.

You can also specify:

- A new recoverable sequence number to be set in the Tpipe for messages sent by WebSphere MQ, and to be set as the partner's receive sequence number. If you do not specify this, the partner's receive sequence number is set to the current WebSphere MQ send sequence number.
- A new recoverable sequence number to be set in the Tpipe for messages received by WebSphere MQ, and to be set as the partner's send sequence number. If you do not specify this, the partner's send sequence number is set to the current WebSphere MQ receive sequence number.

If there is an unresolved unit of recovery associated with the Tpipe, this is also notified in the message. Use the WebSphere MQ command RESET TPIPE to specify whether to commit the unit of recovery, or back it out. If you commit the unit of recovery, the batch of messages has already been sent to IMS, and is deleted from the bridge queue. If you back the unit of recovery out, the messages are returned to the bridge queue, to be later sent to IMS.

Commit mode 1 (send-then-commit) Tpipes are not synchronized.

Considerations for Commit mode 1 transactions

In IMS, commit mode 1 (CM1) transactions send their output replies before sync point.

A CM1 transaction might not be able to send its reply, for example because:

- The Tpipe on which the reply is to be sent is stopped
- OTMA is stopped
- The OTMA client (that is, the queue manager) has gone away
- The reply-to queue and dead-letter queue are unavailable

For these reasons, the IMS application sending the message pseudo-abends with code U0119. The IMS transaction and program are not stopped in this case.

These reasons often prevent messages being sent into IMS, as well as replies being delivered from IMS. A U0119 abend can occur if:

- The Tpipe, OTMA, or the queue manager is stopped while the message is in IMS
- IMS replies on a different Tpipe to the incoming message, and that Tpipe is stopped
- IMS replies to a different OTMA client, and that client is unavailable.

Whenever a U0119 abend occurs, both the incoming message to IMS and the reply messages to WebSphere MQ are lost. If the output of a CM0 transaction cannot be delivered for any of these reasons, it is queued on the Tpipe within IMS.

Working with tpipe names

Many of the commands used to control the WebSphere MQ - IMS bridge require the *tpipe* name. Use this topic to understand how you can find further details of the tpipe name.

You need *tpipe* names for many of the commands that control the WebSphere MQ - IMS bridge. You can get the tpipe names from DISPLAY QUEUE command and note the following points:

- tpipe names are assigned when a local queue is defined
- a local queue is given two tpipe names, one for sync and one for non-sync
- tpipe names will not be known to IMS until after some communication between IMS and WebSphere MQ specific to that particular local queue takes place
- For a tpipe to be available for use by the WebSphere MQ IMS Bridge its associated queue must be assigned to a Storage Class that has the correct XCF group and member name fields completed

Deleting messages from IMS

A message that is destined for WebSphere MQ through the IMS bridge can be deleted if the Tmember/Tpipe is stopped. To delete one message for the queue manager with XCF member name *tmember*, issue the following IMS command:

```
/DEQUEUE TMEMBER tmember TPIPE tpipe PURGE1
```

To delete all the messages on the Tpipe, issue the following IMS command:

```
/DEQUEUE TMEMBER tmember TPIPE tpipe PURGE
```

Deleting tpipes

You cannot delete IMS tpipes yourself. They are deleted by IMS at the following times:

- Synchronized tpipes are deleted when IMS is cold started.
- Non-synchronized tpipes are deleted when IMS is restarted.

IMS Transaction Expiration

An expiration time is associated with a transaction; any WebSphere MQ message can have an expiration time associated with it. The expiration interval is passed from the application, to WebSphere MQ, using the MQMD.Expiry field. The time is the duration of a message before it expires, expressed as a value in tenths of a second. An attempt to perform the MQGET of a message, later than it has expired, results in the message being removed from the queue and expiry processing performed. The expiration time decreases as a message flows between queue managers on a WebSphere MQ network. When an IMS message is passed across the IMS Bridge to OTMA, the remaining message expiry time is passed to OTMA as a transaction expiration time.

If a transaction has an expiration time specified, OTMA expires the input transactions in three different places in IMS:

- input message receiving from XCF
- input message enqueueing time
- application GU time

No expiration is performed after the GU time.

The transaction EXPRTIME can be provided by:

- IMS transaction definition
- IMS OTMA message header
- IMS DFSINSX0 user exit
- IMS CREATE or UPDATE TRAN commands

IMS indicates that it has expired a transaction by abending a transaction with 0243, and issuing a message. The message issued is either DFS555I in the non-shared-queues environment, or DFS2224I in the shared-queues environment.

Security

Security is an important consideration for both developers of WebSphere MQ applications, and for system administrators configuring WebSphere MQ authorities.

Security overview

This collection of topics introduces the WebSphere MQ security concepts.

Security concepts and mechanisms, as they apply to any computer system, are presented first, followed by a discussion of those security mechanisms as they are implemented in IBM WebSphere MQ.

Security concepts and mechanisms

This collection of topics describes aspects of security to consider in your IBM WebSphere MQ installation.

When you are planning a WebSphere MQ implementation, consider which security mechanisms you require to implement those aspects of security that are important to you. For information about what to consider after you have read these topics, see “Planning for your security requirements” on page 414.

Related concepts:

“Connecting two queue managers using SSL or TLS” on page 761

“Working with SSL or TLS” on page 616

Identification and authentication

Identification is the ability to identify uniquely a user of a system or an application that is running in the system. *Authentication* is the ability to prove that a user or application is genuinely who that person or what that application claims to be.

For example, consider a user who logs on to a system by entering a user ID and password. The system uses the user ID to identify the user. The system authenticates the user at the time of logon by checking that the supplied password is correct.

Non-repudiation

The *non-repudiation* service can be viewed as an extension to the identification and authentication service. In general, non-repudiation applies when data is transmitted electronically; for example, an order to a stock broker to buy or sell stock, or an order to a bank to transfer funds from one account to another.

The overall goal of the non-repudiation service is to be able to prove that a particular message is associated with a particular individual.

The non-repudiation service can contain more than one component, where each component provides a different function. If the sender of a message ever denies sending it, the non-repudiation service with *proof of origin* can provide the receiver with undeniable evidence that the message was sent by that particular individual. If the receiver of a message ever denies receiving it, the non-repudiation service with *proof of delivery* can provide the sender with undeniable evidence that the message was received by that particular individual.

In practice, proof with virtually 100% certainty, or undeniable evidence, is a difficult goal. In the real world, nothing is fully secure. Managing security is more concerned with managing risk to a level that is acceptable to the business. In such an environment, a more realistic expectation of the non-repudiation service is to be able to provide evidence that is admissible, and supports your case, in a court of law.

Non-repudiation is a relevant security service in a IBM WebSphere MQ environment because IBM WebSphere MQ is a means of transmitting data electronically. For example, you might require contemporaneous evidence that a particular message was sent or received by an application associated with a particular individual.

IBM WebSphere MQ and IBM WebSphere MQ Advanced Message Security do not provide a non-repudiation service as part of their base function. However, this product documentation does contain suggestions on how you might provide your own non-repudiation service within a WebSphere MQ environment by writing your own exit programs.

Related concepts:

“Identification and authentication in IBM WebSphere MQ” on page 383

Authorization

Authorization protects critical resources in a system by limiting access only to authorized users and their applications. It prevents the unauthorized use of a resource or the use of a resource in an unauthorized manner.

Related concepts:

“Authorization in IBM WebSphere MQ” on page 384

Auditing

Auditing is the process of recording and checking events to detect whether any unexpected or unauthorized activity has taken place, or whether any attempt has been made to perform such activity.

Related concepts:

“Auditing in IBM WebSphere MQ” on page 384

Confidentiality

The *confidentiality* service protects sensitive information from unauthorized disclosure.

When sensitive data is stored locally, access control mechanisms might be sufficient to protect it on the assumption that the data cannot be read if it cannot be accessed. If a greater level of security is required, the data can be encrypted.

Encrypt sensitive data when it is transmitted over a communications network, especially over an insecure network such as the Internet. In a networking environment, access control mechanisms are not effective against attempts to intercept the data, such as wiretapping.

Related concepts:

“Confidentiality in IBM WebSphere MQ” on page 385

Data integrity

The *data integrity* service detects whether there has been unauthorized modification of data.

There are two ways in which data might be altered: accidentally, through hardware and transmission errors, or because of a deliberate attack. Many hardware products and transmission protocols have mechanisms to detect and correct hardware and transmission errors. The purpose of the data integrity service is to detect a deliberate attack.

The data integrity service aims only to detect whether data has been modified. It does not aim to restore data to its original state if it has been modified.

Access control mechanisms can contribute to data integrity insofar as data cannot be modified if access is denied. But, as with confidentiality, access control mechanisms are not effective in a networking environment.

Related concepts:

“Data integrity in IBM WebSphere MQ” on page 385

Cryptographic concepts

This collection of topics describes the concepts of cryptography applicable to WebSphere MQ.

The term *entity* is used to refer to a queue manager, a WebSphere MQ MQI client, an individual user, or any other system capable of exchanging messages.

Related concepts:

“Cryptography in IBM WebSphere MQ” on page 385

Cryptography:

Cryptography is the process of converting between readable text, called *plaintext*, and an unreadable form, called *ciphertext*.

This occurs as follows:

1. The sender converts the plaintext message to ciphertext. This part of the process is called *encryption* (sometimes *encipherment*).
2. The ciphertext is transmitted to the receiver.
3. The receiver converts the ciphertext message back to its plaintext form. This part of the process is called *decryption* (sometimes *decipherment*).

The conversion involves a sequence of mathematical operations that change the appearance of the message during transmission but do not affect the content. Cryptographic techniques can ensure confidentiality and protect messages against unauthorized viewing (eavesdropping), because an encrypted message is not understandable. Digital signatures, which provide an assurance of message integrity, use encryption techniques. See “Digital signatures in SSL and TLS” on page 381 for more information.

Cryptographic techniques involve a general algorithm, made specific by the use of keys. There are two classes of algorithm:

- Those that require both parties to use the same secret key. Algorithms that use a shared key are known as *symmetric* algorithms. Figure 69 on page 370 illustrates symmetric key cryptography.
- Those that use one key for encryption and a different key for decryption. One of these must be kept secret but the other can be public. Algorithms that use public and private key pairs are known as *asymmetric* algorithms. Figure 70 on page 370 illustrates asymmetric key cryptography, which is also known as *public key cryptography*.

The encryption and decryption algorithms used can be public but the shared secret key and the private key must be kept secret.

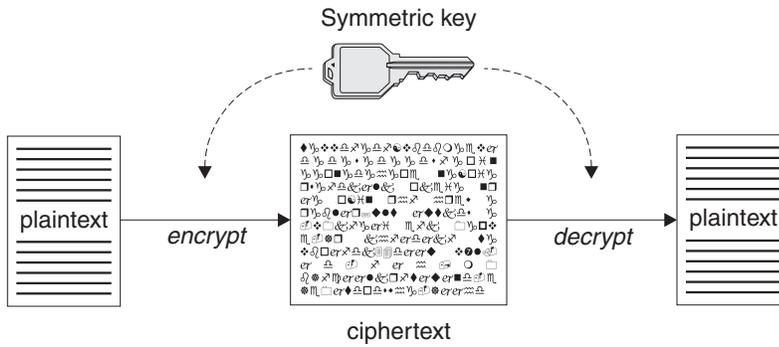


Figure 69. Symmetric key cryptography

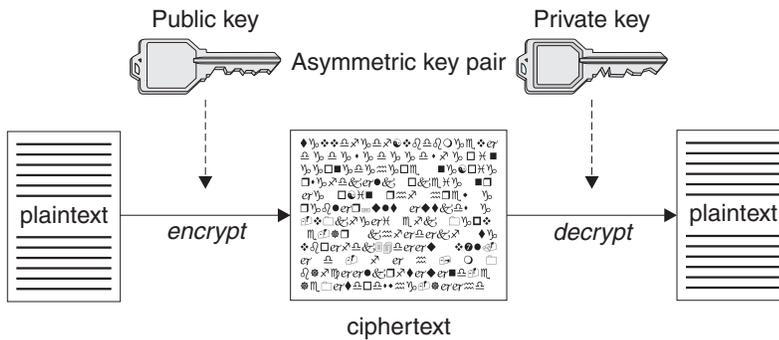


Figure 70. Asymmetric key cryptography

Figure 70 shows plaintext encrypted with the receiver's public key and decrypted with the receiver's private key. Only the intended receiver holds the private key for decrypting the ciphertext. Note that the sender can also encrypt messages with a private key, which allows anyone that holds the sender's public key to decrypt the message, with the assurance that the message must have come from the sender.

With asymmetric algorithms, messages are encrypted with either the public or the private key but can be decrypted only with the other key. Only the private key is secret, the public key can be known by anyone. With symmetric algorithms, the shared key must be known only to the two parties. This is called the *key distribution problem*. Asymmetric algorithms are slower but have the advantage that there is no key distribution problem.

Other terminology associated with cryptography is:

Strength

The strength of encryption is determined by the key size. Asymmetric algorithms require large keys, for example:

- 1024 bits Low-strength asymmetric key
- 2048 bits Medium-strength asymmetric key
- 4096 bits High-strength asymmetric key

Symmetric keys are smaller: 256 bit keys give you strong encryption.

Block cipher algorithm

These algorithms encrypt data by blocks. For example, the RC2 algorithm from RSA Data Security Inc. uses blocks 8 bytes long. Block algorithms are typically slower than stream algorithms.

Stream cipher algorithm

These algorithms operate on each byte of data. Stream algorithms are typically faster than block algorithms.

Message digests and digital signatures:

A message digest is a fixed size numeric representation of the contents of a message, computed by a hash function. A message digest can be encrypted, forming a digital signature.

Messages are inherently variable in size. A message digest is a fixed size numeric representation of the contents of a message. A message digest is computed by a hash function, which is a transformation that meets two criteria:

- The hash function must be one way. It must not be possible to reverse the function to find the message corresponding to a particular message digest, other than by testing all possible messages.
- It must be computationally infeasible to find two messages that hash to the same digest.

The message digest is sent with the message itself. The receiver can generate a digest for the message and compare it with the digest of the sender. The integrity of the message is verified when the two message digests are the same. Any tampering with the message during transmission almost certainly results in a different message digest.

A message digest created using a secret symmetric key is known as a Message Authentication Code (MAC), because it can provide assurance that the message has not been modified.

The sender can also generate a message digest and then encrypt the digest using the private key of an asymmetric key pair, forming a digital signature. The signature must then be decrypted by the receiver, before comparing it with a locally generated digest.

Digital certificates:

Digital certificates protect against impersonation, certifying that a public key belongs to a specified entity. They are issued by a Certificate Authority.

Digital certificates provide protection against impersonation, because a digital certificate binds a public key to its owner, whether that owner is an individual, a queue manager, or some other entity. Digital certificates are also known as public key certificates, because they give you assurances about the ownership of a public key when you use an asymmetric key scheme. A digital certificate contains the public key for an entity and is a statement that the public key belongs to that entity:

- When the certificate is for an individual entity, the certificate is called a *personal certificate* or *user certificate*.
- When the certificate is for a Certificate Authority, the certificate is called a *CA certificate* or *signer certificate*.

If public keys are sent directly by their owner to another entity, there is a risk that the message could be intercepted and the public key substituted by another. This is known as a *man in the middle attack*. The solution to this problem is to exchange public keys through a trusted third party, giving you a strong assurance that the public key really belongs to the entity with which you are communicating. Instead of sending your public key directly, you ask the trusted third party to incorporate it into a digital certificate. The trusted third party that issues digital certificates is called a Certificate Authority (CA), as described in "Certificate Authorities" on page 373.

What is in a digital certificate:

Digital certificates contain specific pieces of information, as determined by the X.509 standard.

Digital certificates used by WebSphere MQ comply with the X.509 standard, which specifies the information that is required and the format for sending it. X.509 is the Authentication framework part of the X.500 series of standards.

Digital certificates contain at least the following information about the entity being certified:

- The owner's public key
- The owner's Distinguished Name
- The Distinguished Name of the CA that issued the certificate
- The date from which the certificate is valid
- The expiry date of the certificate
- The version number of the certificate data format as defined in X.509. The current version of the X.509 standard is Version 3, and most certificates conform to that version.
- A serial number. This is a unique identifier assigned by the CA which issued the certificate. The serial number is unique within the CA which issued the certificate: no two certificates signed by the same CA certificate have the same serial number.

An X.509 Version 2 certificate also contains an Issuer Identifier and a Subject Identifier, and an X.509 Version 3 certificate can contain a number of extensions. Some certificate extensions, such as the Basic Constraint extension, are *standard*, but others are implementation-specific. An extension can be *critical*, in which case a system must be able to recognize the field; if it does not recognize the field, it must reject the certificate. If an extension is not critical, the system can ignore it if it does not recognize it.

The digital signature in a personal certificate is generated using the private key of the CA which signed that certificate. Anyone who needs to verify the personal certificate can use the CA's public key to do so. The CA's certificate contains its public key.

Digital certificates do not contain your private key. You must keep your private key secret.

Related concepts:

“Distinguished Names” on page 373

Requirements for personal certificates:

WebSphere MQ supports digital certificates that comply with the X.509 standard. It requires the client authentication option.

Because IBM WebSphere MQ is a peer to peer system, it is viewed as client authentication in SSL terminology. Therefore, any personal certificate used for SSL authentication needs to allow a key usage of client authentication. Not all server certificates have this option enabled, so the certificate provider might need to enable client authentication on the root CA for the secure certificate.

In addition to the standards which specify the data format for a digital certificate, there are also standards for determining whether a certificate is valid. These standards have been updated over time in order to prevent certain types of security breach. For example, older X.509 version 1 and 2 certificates did not indicate whether the certificate could be legitimately used to sign other certificates. It was therefore possible for a malicious user to obtain a personal certificate from a legitimate source and create new certificates designed to impersonate other users.

When using X.509 version 3 certificates, the BasicConstraints and KeyUsage certificate extensions are used to specify which certificates can legitimately sign other certificates. The IETF RFC 5280 standard specifies

a series of certificate validation rules which compliant application software must implement in order to prevent impersonation attacks. A set of certificate rules is known as a certificate validation policy.

For more information about certificate validation policies in IBM WebSphere MQ, see “Certificate validation policies in WebSphere MQ” on page 402.

Certificate Authorities:

A Certificate Authority (CA) is a trusted third party that issues digital certificates to provide you with an assurance that the public key of an entity truly belongs to that entity.

The roles of a CA are:

- On receiving a request for a digital certificate, to verify the identity of the requestor before building, signing and returning the personal certificate
- To provide the CA's own public key in its CA certificate
- To publish lists of certificates that are no longer trusted in a Certificate Revocation List (CRL). For more information, see “Working with revoked certificates” on page 695
- To provide access to certificate revocation status by operating an OCSP responder server

Distinguished Names:

The Distinguished Name (DN) uniquely identifies an entity in an X.509 certificate.

The following attribute types are commonly found in the DN:

Attribute	Description
SERIALNUMBER	Certificate serial number
MAIL	Email address
E	Email address (Deprecated in preference to MAIL)
UID or USERID	User identifier
CN	Common Name
T	Title
OU	Organizational Unit name
DC	Domain component
O	Organization name
STREET	Street / First line of address
L	Locality name
ST (or SP or S)	State or Province name
PC	Postal code / zip code
C	Country
UNSTRUCTUREDNAME	Host name
UNSTRUCTUREDADDRESS	IP address
DNQ	Distinguished name qualifier

The X.509 standard defines other attributes that do not typically form part of the DN but can provide optional extensions to the digital certificate.

The X.509 standard provides for a DN to be specified in a string format. For example:

CN=John Smith, OU=Test, O=IBM, C=GB

The Common Name (CN) can describe an individual user or any other entity, for example a web server.

The DN can contain multiple OU and DC attributes. Only one instance of each of the other attributes is permitted. The order of the OU entries is significant: the order specifies a hierarchy of Organizational Unit names, with the highest-level unit first. The order of the DC entries is also significant.

IBM WebSphere MQ tolerates certain malformed DNs. For more information, see  WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*).

Obtaining personal certificates from a certificate authority:

You can obtain a certificate from a trusted external certificate authority (CA).

You obtain a digital certificate by sending information to a CA, in the form of a certificate request. The X.509 standard defines a format for this information, but some CAs have their own format. Certificate requests are typically generated by the certificate management tool your system uses, for example the iKeyman tool on UNIX, Linux, and Windows systems and RACF on z/OS. The information contains your Distinguished Name and your public key. When your certificate management tool generates your certificate request, it also generates your private key, which you must keep secure. Never distribute your private key.

When the CA receives your request, the authority verifies your identity before building the certificate and returning it to you as a personal certificate.

Figure 71 illustrates the process of obtaining a digital certificate from a CA.

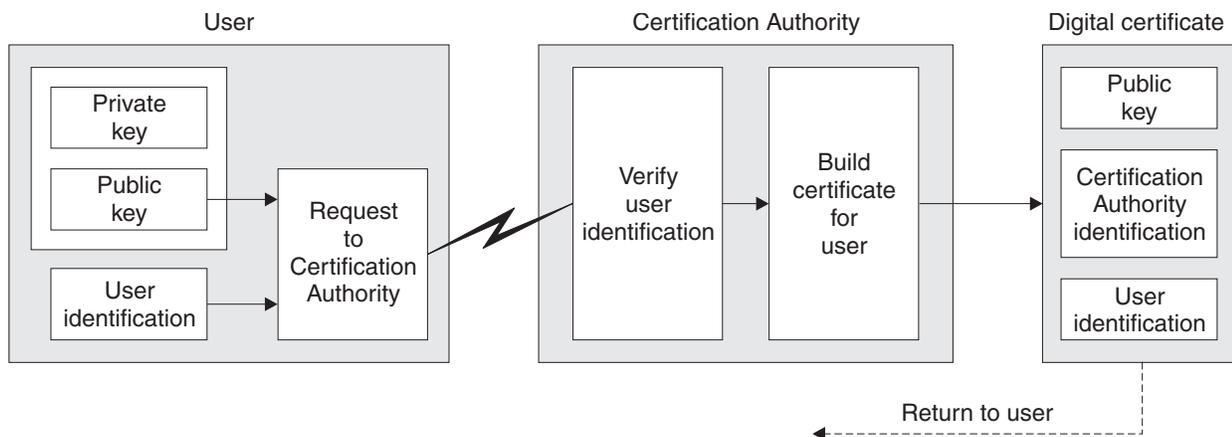


Figure 71. Obtaining a digital certificate

In the diagram:

- "User identification" includes your Subject Distinguished Name.
- "Certification Authority identification" includes the Distinguished Name of the CA that is issuing the certificate.
-

Digital certificates contain additional fields other than those shown in the diagram. For more information about the other fields in a digital certificate, see "What is in a digital certificate" on page 372.

How certificate chains work:

When you receive the certificate for another entity, you might need to use a *certificate chain* to obtain the *root CA* certificate.

The certificate chain, also known as the *certification path*, is a list of certificates used to authenticate an entity. The chain, or path, begins with the certificate of that entity, and each certificate in the chain is signed by the entity identified by the next certificate in the chain. The chain terminates with a root CA certificate. The root CA certificate is always signed by the certificate authority (CA) itself. The signatures of all certificates in the chain must be verified until the root CA certificate is reached.

Figure 72 illustrates a certification path from the certificate owner to the root CA, where the chain of trust begins.

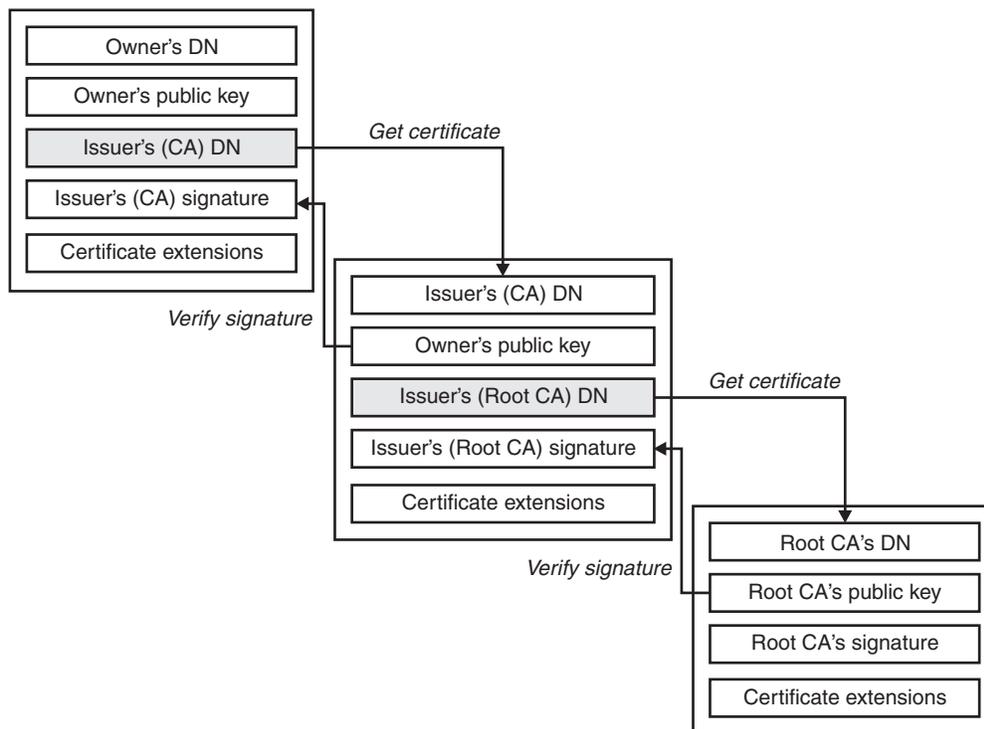


Figure 72. Chain of trust

Each certificate can contain one or more extensions. A certificate belonging to a CA typically contains a BasicConstraints extension with the isCA flag set to indicate that it is allowed to sign other certificates.

When certificates are no longer valid:

Digital certificates can expire or be revoked.

Digital certificates are issued for a fixed period and are not valid after their expiry date.

See the  Glossary (*WebSphere MQ V7.1 Product Overview Guide*) for a definition of certificate expiration.

Certificates can be revoked for various reasons, including:

- The owner has moved to a different organization.
- The private key is no longer secret.

WebSphere MQ can check whether a certificate is revoked by sending a request to an Online Certificate Status Protocol (OCSP) responder (on UNIX, Linux and Windows systems only). Alternatively, they can access a CRL on an LDAP server. The OCSP revocation and CRL information is published by a Certificate Authority. For more information, see “Working with revoked certificates” on page 695.

Public Key Infrastructure (PKI):

A Public Key Infrastructure (PKI) is a system of facilities, policies, and services that supports the use of public key cryptography for authenticating the parties involved in a transaction.

There is no single standard that defines the components of a Public Key Infrastructure, but a PKI typically comprises certificate authorities (CAs) and Registration Authorities (RAs). CAs provide the following services::

- Issuing digital certificates
- Validating digital certificates
- Revoking digital certificates
- Distributing public keys

The X.509 standards provide the basis for the industry standard Public Key Infrastructure.

Refer to “Digital certificates” on page 371 for more information about digital certificates and certificate authorities (CAs). RAs verify that the information provided when digital certificates are requested. If the RA verifies that information, the CA can issue a digital certificate to the requester.

A PKI might also provide tools for managing digital certificates and public keys. A PKI is sometimes described as a *trust hierarchy* for managing digital certificates, but most definitions include additional services. Some definitions include encryption and digital signature services, but these services are not essential to the operation of a PKI.

Cryptographic security protocols: SSL and TLS

Cryptographic protocols provide secure connections, enabling two parties to communicate with privacy and data integrity. The Transport Layer Security (TLS) protocol evolved from that of the Secure Sockets Layer (SSL). IBM WebSphere MQ supports both SSL and TLS.

The primary goals of both protocols is to provide confidentiality, (sometimes referred to as *privacy*), data integrity, identification, and authentication using digital certificates.

Although the two protocols are similar, the differences are sufficiently significant that SSL 3.0 and the various versions of TLS do not interoperate.

Related concepts:

“Security protocols in WebSphere MQ” on page 386

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) concepts:

The SSL and TLS protocols enable two parties to identify and authenticate each other and communicate with confidentiality and data integrity. The TLS protocol evolved from the Netscape SSL 3.0 protocol but TLS and SSL do not interoperate.

The SSL and TLS protocols provide communications security over the internet, and allow client/server applications to communicate in a way that is confidential and reliable. The protocols have two layers: a Record Protocol and a Handshake Protocol, and these are layered above a transport protocol such as TCP/IP. They both use asymmetric and symmetric cryptography techniques.

An SSL or TLS connection is initiated by an application, which becomes the SSL or TLS client. The application which receives the connection becomes the SSL or TLS server. Every new session begins with a handshake, as defined by the SSL or TLS protocols.

A full list of CipherSpecs supported by IBM WebSphere MQ is provided at “Specifying CipherSpecs” on page 774.

For more information about the SSL protocol, see the information provided at <http://www.mozilla.org/projects/security/pki/nss/ssl/draft302.txt>. For more information about the TLS protocol, see the information provided by the TLS Working Group on the website of the Internet Engineering Task Force at <http://www.ietf.org>

Related concepts:

“Cryptographic concepts” on page 369

An overview of the SSL or TLS handshake:

The SSL or TLS handshake enables the SSL or TLS client and server to establish the secret keys with which they communicate.

This section provides a summary of the steps that enable the SSL or TLS client and server to communicate with each other:

- Agree on the version of the protocol to use.
- Select cryptographic algorithms.
- Authenticate each other by exchanging and validating digital certificates.
- Use asymmetric encryption techniques to generate a shared secret key, which avoids the key distribution problem. SSL or TLS then uses the shared key for the symmetric encryption of messages, which is faster than asymmetric encryption.

For more information about cryptographic algorithms and digital certificates, refer to the related information.

This section does not attempt to provide full details of the messages exchanged during the SSL handshake. In overview, the steps involved in the SSL handshake are as follows:

1. The SSL or TLS client sends a “client hello” message that lists cryptographic information such as the SSL or TLS version and, in the client's order of preference, the CipherSuites supported by the client. The message also contains a random byte string that is used in subsequent computations. The protocol allows for the “client hello” to include the data compression methods supported by the client.
2. The SSL or TLS server responds with a “server hello” message that contains the CipherSuite chosen by the server from the list provided by the client, the session ID, and another random byte string. The server also sends its digital certificate. If the server requires a digital certificate for client authentication, the server sends a “client certificate request” that includes a list of the types of certificates supported and the Distinguished Names of acceptable Certification Authorities (CAs).
3. The SSL or TLS client verifies the server's digital certificate. For more information, see “How SSL and TLS provide identification, authentication, confidentiality, and integrity” on page 378.
4. The SSL or TLS client sends the random byte string that enables both the client and the server to compute the secret key to be used for encrypting subsequent message data. The random byte string itself is encrypted with the server's public key.
5. If the SSL or TLS server sent a “client certificate request”, the client sends a random byte string encrypted with the client's private key, together with the client's digital certificate, or a “no digital certificate alert”. This alert is only a warning, but with some implementations the handshake fails if client authentication is mandatory.

6. The SSL or TLS server verifies the client's certificate. For more information, see "How SSL and TLS provide identification, authentication, confidentiality, and integrity."
7. The SSL or TLS client sends the server a "finished" message, which is encrypted with the secret key, indicating that the client part of the handshake is complete.
8. The SSL or TLS server sends the client a "finished" message, which is encrypted with the secret key, indicating that the server part of the handshake is complete.
9. For the duration of the SSL or TLS session, the server and client can now exchange messages that are symmetrically encrypted with the shared secret key.

Figure 73 illustrates the SSL or TLS handshake.

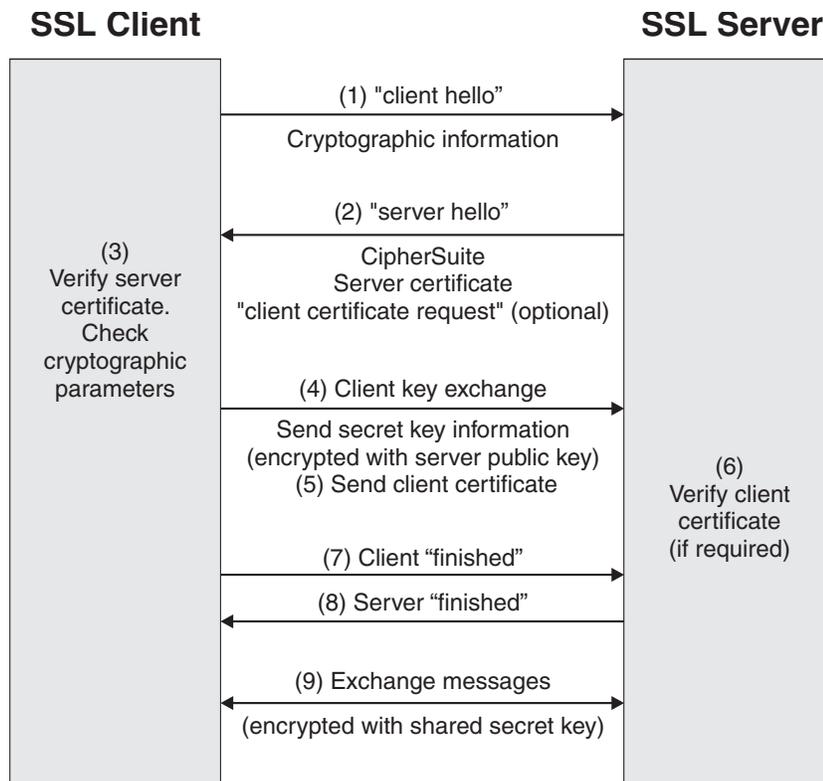


Figure 73. Overview of the SSL or TLS handshake

Related concepts:

"CipherSpecs and CipherSuites" on page 380

"Digital certificates" on page 371

How SSL and TLS provide identification, authentication, confidentiality, and integrity:

During both client and server authentication there is a step that requires data to be encrypted with one of the keys in an asymmetric key pair and decrypted with the other key of the pair. A message digest is used to provide integrity.

How SSL and TLS provide authentication

For server authentication, the client uses the server's public key to encrypt the data that is used to compute the secret key. The server can generate the secret key only if it can decrypt that data with the correct private key.

For client authentication, the server uses the public key in the client certificate to decrypt the data the client sends during step 5 on page 377 of the handshake. The exchange of finished messages that are encrypted with the secret key (steps 7 on page 378 and 8 on page 378 in the overview) confirms that authentication is complete.

If any of the authentication steps fail, the handshake fails and the session terminates.

The exchange of digital certificates during the SSL or TLS handshake is part of the authentication process. For more information about how certificates provide protection against impersonation, refer to the related information. The certificates required are as follows, where CA X issues the certificate to the SSL or TLS client, and CA Y issues the certificate to the SSL or TLS server:

For server authentication only, the SSL or TLS server needs:

- The personal certificate issued to the server by CA Y
- The server's private key

and the SSL or TLS client needs:

- The CA certificate for CA Y

If the SSL or TLS server requires client authentication, the server verifies the client's identity by verifying the client's digital certificate with the public key for the CA that issued the personal certificate to the client, in this case CA X. For both server and client authentication, the server needs:

- The personal certificate issued to the server by CA Y
- The server's private key
- The CA certificate for CA X

and the client needs:

- The personal certificate issued to the client by CA X
- The client's private key
- The CA certificate for CA Y

Both the SSL or TLS server and client might need other CA certificates to form a certificate chain to the root CA certificate. For more information about certificate chains, refer to the related information.

What happens during certificate verification

As noted in steps 3 on page 377 and 6 on page 378 of the overview, the SSL or TLS client verifies the server's certificate, and the SSL or TLS server verifies the client's certificate. There are four aspects to this verification:

1. The digital signature is checked (see "Digital signatures in SSL and TLS" on page 381).
2. The certificate chain is checked you should have intermediate CA certificates (see "How certificate chains work" on page 375).
3. The expiry and activation dates and the validity period are checked.
4. The revocation status of the certificate is checked (see "Working with revoked certificates" on page 695).

Secret key reset

During an SSL or TLS handshake a *secret key* is generated to encrypt data between the SSL or TLS client and server. The secret key is used in a mathematical formula that is applied to the data to transform plaintext into unreadable ciphertext, and ciphertext into plaintext.

The secret key is generated from the random text sent as part of the handshake and is used to encrypt plaintext into ciphertext. The secret key is also used in the MAC (Message Authentication Code)

algorithm, which is used to determine whether a message has been altered. See “Message digests and digital signatures” on page 371 for more information.

If the secret key is discovered, the plaintext of a message could be deciphered from the ciphertext, or the message digest could be calculated, allowing messages to be altered without detection. Even for a complex algorithm, the plaintext can eventually be discovered by applying every possible mathematical transformation to the ciphertext. To minimize the amount of data that can be deciphered or altered if the secret key is broken, the secret key can be renegotiated periodically. When the secret key has been renegotiated, the previous secret key can no longer be used to decrypt data encrypted with the new secret key.

How SSL and TLS provide confidentiality

SSL and TLS use a combination of symmetric and asymmetric encryption to ensure message privacy. During the SSL or TLS handshake, the SSL or TLS client and server agree an encryption algorithm and a shared secret key to be used for one session only. All messages transmitted between the SSL or TLS client and server are encrypted using that algorithm and key, ensuring that the message remains private even if it is intercepted. SSL supports a wide range of cryptographic algorithms. Because SSL and TLS use asymmetric encryption when transporting the shared secret key, there is no key distribution problem. For more information about encryption techniques, refer to “Cryptography” on page 369.

How SSL and TLS provide integrity

SSL and TLS provide data integrity by calculating a message digest. For more information, see “Data integrity of messages” on page 781.

Use of SSL or TLS does ensure data integrity, provided that the CipherSpec in your channel definition uses a hash algorithm as described in the table in “Specifying CipherSpecs” on page 774.

In particular, if data integrity is a concern, you should avoid choosing a CipherSpec whose hash algorithm is listed as “None”. Use of MD5 is also strongly discouraged as this is now very old and no longer secure for most practical purposes.

Related concepts:

“Digital certificates” on page 371

“How certificate chains work” on page 375

CipherSpecs and CipherSuites:

Cryptographic security protocols must agree the algorithms used by a secure connection. CipherSpecs and CipherSuites define specific combinations of algorithms.

A CipherSpec identifies a combination of encryption algorithm and MAC algorithm. Both ends of an SSL or TLS connection must agree the same CipherSpec to be able to communicate.

For more information about CipherSpecs, see the related information.

A CipherSuite is a suite of cryptographic algorithms used by an SSL or TLS connection. A suite comprises three distinct algorithms:

- The key exchange and authentication algorithm, used during the handshake
- The encryption algorithm, used to encipher the data
- The MAC (Message Authentication Code) algorithm, used to generate the message digest

There are several options for each component of the suite, but only certain combinations are valid when specified for an SSL or TLS connection. The name of a valid CipherSuite defines the combination of algorithms used. For example, the CipherSuite `SSL_RSA_WITH_RC4_128_MD5` specifies:

- The RSA key exchange and authentication algorithm
- The RC4 encryption algorithm, using a 128-bit key
- The MD5 MAC algorithm

Several algorithms are available for key exchange and authentication, but the RSA algorithm is currently the most widely used. There is more variety in the encryption algorithms and MAC algorithms that are used.

Related concepts:

“CipherSpecs and CipherSuites in IBM WebSphere MQ” on page 399

Digital signatures in SSL and TLS:

A digital signature is formed by encrypting a representation of a message. The encryption uses the private key of the signatory and, for efficiency, usually operates on a message digest rather than the message itself.

Digital signatures vary with the data being signed, unlike handwritten signatures, which do not depend on the content of the document being signed. If two different messages are signed digitally by the same entity, the two signatures differ, but both signatures can be verified with the same public key, that is, the public key of the entity that signed the messages.

The steps of the digital signature process are as follows:

1. The sender computes a message digest and then encrypts the digest using the sender's private key, forming the digital signature.
2. The sender transmits the digital signature with the message.
3. The receiver decrypts the digital signature using the sender's public key, regenerating the sender's message digest.
4. The receiver computes a message digest from the message data received and verifies that the two digests are the same.

Figure 74 illustrates this process.

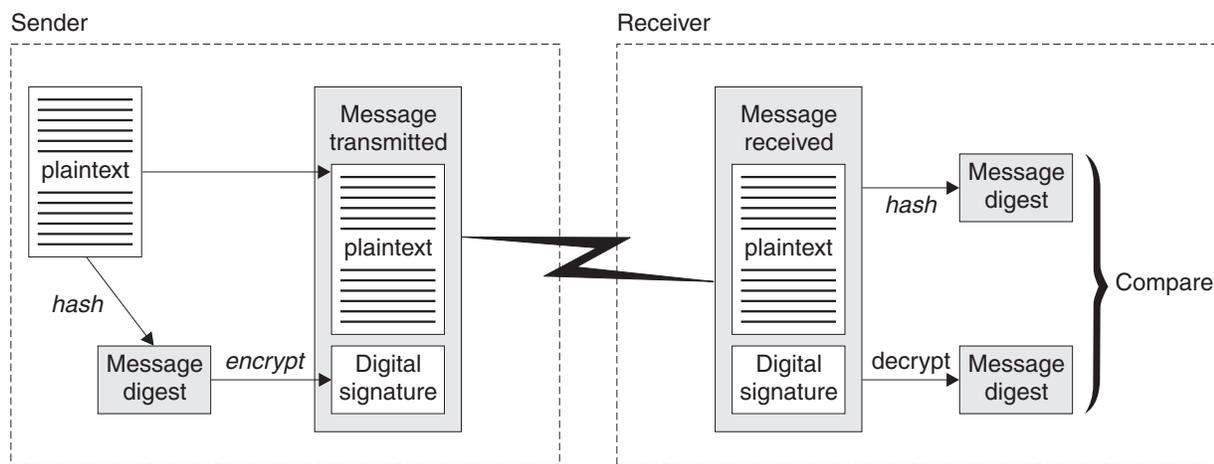


Figure 74. The digital signature process

If the digital signature is verified, the receiver knows that:

- The message has not been modified during transmission.
- The message was sent by the entity that claims to have sent it.

Digital signatures are part of integrity and authentication services. Digital signatures also provide proof of origin. Only the sender knows the private key, which provides strong evidence that the sender is the originator of the message.

Note: You can also encrypt the message itself, which protects the confidentiality of the information in the message.

Related concepts:

“Message digests and digital signatures” on page 371

Federal Information Processing Standards:

The US government produces technical advice on IT systems and security, including data encryption. The National Institute for Standards and Technology (NIST) is an important body concerned with IT systems and security. NIST produces recommendations and standards, including the Federal Information Processing Standards (FIPS).

A significant one of these standards is FIPS 140-2, which requires the use of strong cryptographic algorithms. FIPS 140-2 also specifies requirements for hashing algorithms to be used to protect packets against modification in transit.

IBM WebSphere MQ provides FIPS 140-2 support when it has been configured to do so.

Over time, analysts develop attacks against existing encryption and hashing algorithms. New algorithms are adopted to resist those attacks. FIPS 140-2 is periodically updated to take account of these changes.

Related concepts:

“Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows” on page 394

National Security Agency (NSA) Suite B Cryptography:

The government of the United States of America produces technical advice on IT systems and security, including data encryption. The US National Security Agency (NSA) recommends a set of interoperable cryptographic algorithms in its Suite B standard.

The Suite B standard specifies a mode of operation in which only a specific set of secure cryptographic algorithms are used. The Suite B standard specifies:

- The encryption algorithm (AES)
- The key exchange algorithm (Elliptic Curve Diffie-Hellman, also known as ECDH)
- The digital signature algorithm (Elliptic Curve Digital Signature Algorithm, also known as ECDSA)
- The hashing algorithms (SHA-256 or SHA-384)

Additionally, the IETF RFC 6460 standard specifies Suite B compliant profiles which define the detailed application configuration and behavior necessary to comply with the Suite B standard. It defines two profiles:

1. A Suite B compliant profile for use with TLS version 1.2. When configured for Suite B compliant operation, only the restricted set of cryptographic algorithms listed above will be used.
2. A transitional profile for use with TLS version 1.0 or TLS version 1.1. This profile enables interoperability with non-Suite B compliant servers. When configured for Suite B transitional operation, additional encryption and hashing algorithms may be used.

The Suite B standard is conceptually similar to FIPS 140-2, because it restricts the set of enabled cryptographic algorithms in order to provide an assured level of security.

On Windows, UNIX and Linux systems, WebSphere MQ, can be configured to conform to the Suite B compliant TLS 1.2 profile, but does not support the Suite B transitional profile. For further information, see “NSA Suite B Cryptography in IBM WebSphere MQ” on page 399.

Related information:

“Federal Information Processing Standards” on page 382

Common Criteria

Common Criteria is a scheme for independent assessment, analysis, and testing of IT products to a set of security requirements.

The Common Criteria Scheme provides consumers with an impartial security assurance of a product to predefined levels. These levels range from EAL1 to EAL7. Each assurance level places increased demands on the developer for evidence of testing, and provides increased assurance within the product.

Under the Common Criteria Recognition Arrangement (CCRA), countries agree to recognize Common Criteria certificates that are produced by any certificate authorizing participant, in accordance with the terms laid out in the CCRA. Currently, the CCRA is composed of 26 member nations: Australia, Austria, Canada, Denmark, Finland, France, Germany, Greece, Hungary, India, Israel, Italy, Japan, Malaysia, New Zealand, Norway, Pakistan, the Republic of Singapore, South Korea, Spain, Sweden, the Czech Republic, the Netherlands, Turkey, the United Kingdom and the United States.

For more information about the Common Criteria scheme, see  <https://www.niap-ccavs.org/>.

Related concepts:

“Common Criteria in WebSphere MQ” on page 384

IBM WebSphere MQ security mechanisms

This collection of topics explains how you can implement the various security concepts in IBM WebSphere MQ.

IBM WebSphere MQ provides mechanisms to implement all the security concepts introduced in “Security concepts and mechanisms” on page 367. These are discussed in more detail in the following sections.

Identification and authentication in IBM WebSphere MQ

In IBM WebSphere MQ, you can implement identification and authentication using message context information and mutual authentication.

Here are some examples of the identification and authentication in a IBM WebSphere MQ environment:

- Every message can contain *message context* information. This information is held in the message descriptor. It can be generated by the queue manager when a message is put on a queue by an application. Alternatively, the application can supply the information if the user ID associated with the application is authorized to do so.
The context information in a message allows the receiving application to find out about the originator of the message. It contains, for example, the name of the application that put the message and the user ID associated with the application.
- When a message channel starts, it is possible for the message channel agent (MCA) at each end of the channel to authenticate its partner. This technique is known as *mutual authentication*. For the sending MCA, it provides assurance that the partner it is about to send messages to is genuine. For the receiving MCA, there is a similar assurance that it is about to receive messages from a genuine partner.

Related concepts:

“Identification and authentication” on page 367

Authorization in IBM WebSphere MQ

You can use authorization to limit what particular individuals or applications can do in your IBM WebSphere MQ environment.

Here are some examples of authorization in a IBM WebSphere MQ environment:

- Allowing only an authorized administrator to issue commands to manage IBM WebSphere MQ resources.
- Allowing an application to connect to a queue manager only if the user ID associated with the application is authorized to do so.
- Allowing an application to open only those queues that are necessary for its function.
- Allowing an application to subscribe only to those topics that are necessary for its function.
- Allowing an application to perform only those operations on a queue that are necessary for its function. For example, an application might need only to browse messages on a particular queue, and not to put or get messages.

Related concepts:

“Authorization” on page 368

Auditing in IBM WebSphere MQ

IBM WebSphere MQ can issue event messages to record that unusual activity has taken place.

Here are some examples of auditing in a IBM WebSphere MQ environment:

- An application attempts to open a queue that it is not authorized to open. An instrumentation event message is issued. By inspecting the event message, you discover that this attempt occurred and can decide what action is necessary.
- An application attempts to open a channel, but the attempt fails because SSL does not allow the connection. An instrumentation event message is issued. By inspecting the event message, you discover that this attempt occurred and can decide what action is necessary.

Related concepts:

“Auditing” on page 368

Common Criteria in WebSphere MQ

Common Criteria is a scheme for independent assessment, analysis, and testing of IT products to a set of security requirements.

The Common Criteria Scheme provides consumers with an impartial security assurance of a product to predefined levels. These levels range from EAL1 to EAL7. Each assurance level places increased demands on the developer for evidence of testing, and provides increased assurance within the product.

IBM WebSphere MQ v7.1.0.2 is evaluated to Common Criteria EAL2. This evaluation provides assurance that the product has been structurally tested.

Related concepts:

“Common Criteria” on page 383

“Common Criteria environmental considerations” on page 445

“Configuring IBM WebSphere MQ for Common Criteria” on page 685

Confidentiality in IBM WebSphere MQ

You can implement confidentiality in IBM WebSphere MQ by encrypting messages.

Here are some examples of how confidentiality can be ensured in a IBM WebSphere MQ environment:

- After a sending MCA gets a message from a transmission queue, IBM WebSphere MQ uses SSL or TLS to encrypt the message before it is sent over the network to the receiving MCA. At the other end of the channel, the message is decrypted before the receiving MCA puts it on its destination queue.
- While messages are stored on a local queue, the access control mechanisms provided by IBM WebSphere MQ might be considered sufficient to protect their contents against unauthorized disclosure. However, for a greater level of security, you can use IBM WebSphere MQ Advanced Message Security to encrypt the messages stored in the queues.

Related concepts:

“Confidentiality” on page 368

“Security protocols in WebSphere MQ” on page 386

“WebSphere MQ Advanced Message Security” on page 439

Data integrity in IBM WebSphere MQ

You can use a data integrity service to detect whether a message has been modified.

Here are some examples of how data integrity can be ensured in a IBM WebSphere MQ environment:

- You can use SSL or TLS to detect whether the contents of a message have been deliberately modified while it was being transmitted over a network. In SSL and TLS, the message digest algorithm provides detection of modified messages in transit. All IBM WebSphere MQ CipherSpecs provide a message digest algorithm, except for TLS_RSA_WITH_NULL_NULL which does not provide message data integrity.
- While messages are stored on a local queue, the access control mechanisms provided by IBM WebSphere MQ might be considered sufficient to prevent deliberate modification of the contents of the messages. However, for a greater level of security, you can use IBM WebSphere MQ Advanced Message Security to detect whether the contents of a message have been deliberately modified between the time the message was put on the queue and the time it was retrieved from the queue.

Related concepts:

“Data integrity” on page 368

“Security protocols in WebSphere MQ” on page 386

“WebSphere MQ Advanced Message Security” on page 439

Cryptography in IBM WebSphere MQ

IBM WebSphere MQ provides cryptography by using the Secure sockets Layer (SSL) and Transport Security Layer (TLS) protocols.

For more information see “Security protocols in WebSphere MQ” on page 386.

Related concepts:

“Cryptographic concepts” on page 369

Security protocols in WebSphere MQ

WebSphere MQ supports both the Transport Layer Security (TLS) and the Secure Sockets Layer (SSL) protocols to provide link level security for message channels and MQI channels.

Message channels and MQI channels can use the SSL or TLS protocol to provide link level security. A caller MCA is an SSL or TLS client and a responder MCA is an SSL or TLS server. WebSphere MQ supports Version 3.0 of the SSL protocol and Version 1.0 and Version 1.2 of the Transport Layer Security (TLS) protocol. You specify the cryptographic algorithms that are used by the SSL or protocol by supplying a CipherSpec as part of the channel definition.

At each end of a message channel, and at the server end of an MQI channel, the MCA acts on behalf of the queue manager to which it is connected. During the SSL or TLS handshake, the MCA sends the digital certificate of the queue manager to its partner MCA at the other end of the channel. The WebSphere MQ code at the client end of an MQI channel acts on behalf of the user of the WebSphere MQ client application. During the SSL or TLS handshake, the WebSphere MQ code sends the user's digital certificate to the MCA at the server end of the MQI channel.

Queue managers and WebSphere MQ client users are not required to have personal digital certificates associated with them when they are acting as SSL or TLS clients, unless SSLCAUTH(REQUIRED) is specified at the server side of the channel.

Digital certificates are stored in a *key repository*. The queue manager attribute *SSLKeyRepository* specifies the location of the key repository that holds the queue manager's digital certificate. On a WebSphere MQ client system, the MQSSLKEYR environment variable specifies the location of the key repository that holds the user's digital certificate. Alternatively, a WebSphere MQ client application can specify its location in the *KeyRepository* field of the SSL and TLS configuration options structure, MQSCO, on an MQCONN call. See the related topics for more information about key repositories and how to specify where they are located.

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

Related reference:

“WebSphere MQ support for SSL and TLS”

WebSphere MQ support for SSL and TLS:

WebSphere MQ supports both the Secure Sockets Layer (SSL) protocol and the Transport Layer Security (TLS) protocol.

For more information about the SSL and TLS protocols, refer to the related information.

WebSphere MQ provides the following support for SSL Version 3.0 and TLS 1.0 and TLS 1.2:

IBM i SSL and TLS support is integral to the IBM i operating system.

See  Using TLS Version 1.2 for further information on using this protocol.

Java and JMS clients

These clients use the JVM to provide SSL and TLS support.

Windows, UNIX and Linux, and HP Integrity NonStop Server systems

For UNIX, Linux, HP Integrity NonStop Server, and Windows systems, the SSL and TLS support is installed with WebSphere MQ.

z/OS SSL and TLS support is integral to the z/OS operating system. The SSL and TLS support on z/OS is known as *System SSL*.

For information about any prerequisites for WebSphere MQ SSL and TLS support, see [System Requirements for IBM WebSphere MQ](#).

The following topics describe the provisions in WebSphere MQ that enable you to use and control the SSL and TLS support:

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

CipherSpec values supported in WebSphere MQ:

The set of default CipherSpecs allows only the following values:

TLS 1.0

- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA

TLS 1.2

- ECDHE_ECDSA_AES_128_CBC_SHA256
- ECDHE_ECDSA_AES_256_CBC_SHA384
- ECDHE_ECDSA_AES_128_GCM_SHA256
- ECDHE_ECDSA_AES_256_GCM_SHA384
- ECDHE_RSA_AES_128_CBC_SHA256
- ECDHE_RSA_AES_256_CBC_SHA384
- ECDHE_RSA_AES_128_GCM_SHA256
- ECDHE_RSA_AES_256_GCM_SHA384
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384

The default set of CipherSpecs for WebSphere MQ for IBM i allows only the following values:

- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

Enabling deprecated CipherSpecs

By default, you are not allowed to specify a deprecated CipherSpec on a channel definition. If you attempt to specify a deprecated CipherSpec, you receive message AMQ9788 in the error log for the queue manager.

It is possible for you to re-enable the deprecated CipherSpecs by editing the `qm.ini` file. Within the SSL stanza of the `qm.ini` file, add the following line:

SSL:
 AllowWeakCipherSpec=Yes

You can also re-enable one or more of the deprecated CipherSpecs at runtime on the server by setting the environment variable `AMQ_SSL_WEAK_CIPHER_ENABLE` to any value. This environment variable enables the CipherSpecs regardless of the value that is specified in the `qm.ini` file.

Deprecated CipherSpecs:

A list of deprecated CipherSpecs that you are able to use with WebSphere MQ if necessary.

See “CipherSpec values supported in WebSphere MQ” on page 387 for more information on how you can enable deprecated CipherSpecs.

Deprecated CipherSpecs that you can use with WebSphere MQ TLS support are listed in the following table:

Platform support ¹	CipherSpec name	Protocol used	Data integrity	Encryption algorithm	Encryption bits	FIPS ²	Suite B
IBM i	AES_SHA_US	SSL 3.0	SHA-1	AES	128	No	No
All	DES_SHA_EXPORT ³	SSL 3.0	SHA-1	DES	56	No	No
UNIX, Linux, and Windows	DES_SHA_EXPORT1024 ⁴	SSL 3.0	SHA-1	DES	56	No	No
UNIX, Linux, and Windows	FIPS_WITH_DES_CBC_SHA	SSL 3.0	SHA-1	DES	56	No ⁶	No
UNIX, Linux, and Windows	FIPS_WITH_3DES_EDE_CBC_SHA	SSL 3.0	SHA-1	3DES	168	No ⁷	No
All	NULL_MD5	SSL 3.0	MD5	None	0	No	No
All	NULL_SHA	SSL 3.0	SHA-1	None	0	No	No
All	RC2_MD5_EXPORT ³	SSL 3.0	MD5	RC2	40	No	No
All	RC4_MD5_EXPORT ³	SSL 3.0	MD5	RC4	40	No	No
All	RC4_MD5_US	SSL 3.0	MD5	RC4	128	No	No
All	RC4_SHA_US	SSL 3.0	SHA-1	RC4	128	No	No
UNIX, Linux, Windows	RC4_56_SHA_EXPORT1024 ⁴	SSL 3.0	SHA-1	RC4	56	No	No
All	TRIPLE_DES_SHA_US	SSL 3.0	SHA-1	3DES	168	No	No
IBM i	TLS_RSA_EXPORT_WITH_RC2_40_MD5	TLS 1.0	MD5	RC2	40	No	No
IBM i	TLS_RSA_EXPORT_WITH_RC4_40_MD5 ³	TLS 1.0	MD5	RC4	40	No	No
All	TLS_RSA_WITH_DES_CBC_SHA	TLS 1.0	SHA-1	DES	56	No ⁵	No

Platform support ¹	CipherSpec name	Protocol used	Data integrity	Encryption algorithm	Encryption bits	FIPS ²	Suite B
IBM i	TLS_RSA_WITH_NULL_MD5	TLS 1.0	MD5	None	0	No	No
IBM i	TLS_RSA_WITH_NULL_SHA	TLS 1.0	SHA-1	None	0	No	No
IBM i	TLS_RSA_WITH_RC4_128_MD5	TLS 1.0	MD5	RC4	128	No	No
Unix, Linux, and Windows	ECDHE_ECDSA_NULL_SHA256	TLS 1.2	SHA-1	None	0	No	No
UNIX, Linux, and Windows	ECDHE_ECDSA_RC4_128_SHA256	TLS 1.2	SHA-1	RC4	128	No	No
UNIX, Linux, and Windows	ECDHE_RSA_NULL_SHA256	TLS 1.2	SHA-1	None	0	No	No
UNIX, Linux, and Windows	ECDHE_RSA_RC4_128_SHA256	TLS 1.2	SHA-1	RC4	128	No	No
UNIX, Linux, and Windows	TLS_RSA_WITH_NULL_NULL	TLS 1.2	None	None	0	No	No
All	TLS_RSA_WITH_NULL_SHA256	TLS 1.2	SHA-256	None	0	No	No
UNIX, Linux, and Windows	TLS_RSA_WITH_RC4_128_SHA256	TLS 1.2	SHA-1	RC4	128	No	No
All	TLS_RSA_WITH_3DES_EDE_CBC_SHA ⁸	TLS 1.0	SHA-1	3DES	168	Yes	No
UNIX, Linux, and Windows	ECDHE_ECDSA_3DES_EDE_CBC_SHA256 ⁸	TLS 1.2	SHA-1	3DES	168	Yes	No
UNIX, Linux, and Windows	ECDHE_RSA_3DES_EDE_CBC_SHA256 ⁸	TLS 1.2	SHA-1	3DES	168	Yes	No

Platform support ¹	CipherSpec name	Protocol used	Data integrity	Encryption algorithm	Encryption bits	FIPS ²	Suite B
<p>Notes:</p> <ol style="list-style-type: none"> If no specific platform is noted, the CipherSpec is available on all platforms. Specifies whether the CipherSpec is FIPS-certified on a FIPS-certified platform. The maximum handshake key size is 512 bits. If either of the certificates exchanged during the SSL handshake has a key size greater than 512 bits, a temporary 512-bit key is generated for use during the handshake. The handshake key size is 1024 bits. This CipherSpec was FIPS 140-2 certified before 19 May 2007. This CipherSpec was FIPS 140-2 certified before 19 May 2007. The name FIPS_WITH_DES_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended. The name FIPS_WITH_3DES_EDE_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. The use of this CipherSpec is deprecated. This CipherSpec can be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, either avoid using triple DES, or enable secret key reset when using this CipherSpec. 							

The SSL or TLS key repository:

A mutually authenticated SSL or TLS connection requires a key repository (which can be known by different names on different platforms) at each end of the connection. The key repository includes digital certificates and private keys.

This information uses the general term *key repository* to describe the store for digital certificates and their associated private keys. The specific store names used on the platforms and environments that support SSL and TLS are:

IBM i	certificate store
Java and JMS	keystore and trust store
Windows, UNIX and Linux systems	key database file
z/OS	keyring

For more information, refer to “Digital certificates” on page 371 and “Secure Sockets Layer (SSL) and Transport Layer Security (TLS) concepts” on page 376.

A mutually authenticated SSL or TLS connection requires a key repository at each end of the connection. The key repository may contain:

- A number of CA certificates from various Certification Authorities that allow the queue manager or client to verify certificates that it receives from its partner at the remote end of the connection. Individual certificates might be in a certificate chain.
- One or more personal certificates received from a Certification Authority. You associate a separate personal certificate with each queue manager or WebSphere MQ MQI client. Personal certificates are essential on an SSL or TLS client if mutual authentication is required. If mutual authentication is not required, personal certificates are not needed on the client. The key repository might also contain the private key corresponding to each personal certificate.
- Certificate requests which are waiting to be signed by a trusted CA certificate.

For more information about protecting your key repository, see “Protecting WebSphere MQ key repositories.”

The location of the key repository depends on the platform you are using:

IBM i On IBM i, the key repository is a certificate store. The default system certificate store is located at /QIBM/UserData/ICSS/Cert/Server/Default in the integrated file system (IFS). On IBM i, WebSphere MQ stores the password for the certificate store in a *password stash file*. For example, the stash file for queue manager QM1 is /QIBM/UserData/mqm/qmgrs/QM1/ssl/Stash.sth.

Alternatively, you can specify that the IBM i system certificate store is to be used instead. To do this you change the value of the queue manager's SSLKEYR attribute to *SYSTEM. This value indicates that the queue manager must use the system certificate store, and the queue manager is registered for use as an application with Digital Certificate Manager (DCM).

On IBM i the certificate store also contains the private key for the queue manager.

Windows, UNIX and Linux systems

On Windows, UNIX and Linux systems the key repository is a key database file. The name of the key database file must have a file extension of .kdb. For example, on UNIX and Linux, the default key database file for queue manager QM1 is /var/mqm/qmgrs/QM1/ssl/key.kdb. If WebSphere MQ is installed in the default location, the equivalent path on Windows is C:\Program Files\IBM\WebSphere MQ\Qmgrs\QM1\ssl\key.kdb.

On Windows, UNIX and Linux systems, each key database file has an associated password stash file. This file holds encoded passwords that allow programs to access the key database. The password stash file must be in the same directory and have the same file stem as the key database, and must end with the suffix .sth, for example /var/mqm/qmgrs/QM1/ssl/key.sth

Note: On Windows, UNIX and Linux systems, PKCS #11 cryptographic hardware cards can contain the certificates and keys that are otherwise held in a key database file. When certificates and keys are held on PKCS #11 cards, WebSphere MQ still requires access to both a key database file and a password stash file.

On Windows and UNIX systems, the key database also contains the private key for the personal certificate associated with the queue manager or WebSphere MQ MQI client.

z/OS Certificates are held in a keyring in RACF.

Other external security managers (ESMs) also use keyrings for storing certificates.

On z/OS, private keys are managed by RACF.

Protecting WebSphere MQ key repositories:

The key repository for WebSphere MQ is a file. Ensure that only the intended user can access the key repository file. This prevents an intruder or other unauthorized user copying the key repository file to another system, and then setting up an identical user ID on that system to impersonate the intended user.

The permissions on the files depend on the user's umask and which tool is used. On Windows, WebSphere MQ accounts require permission BypassTraverseChecking which means the permissions of the folders in the file path have no effect.

Check the file permissions of key repository files and make sure that the files and containing folder are not world readable, preferably not even group readable.

Making the keystore read-only is good practice, on whichever system you use, with only the administrator being permitted to enable write operations in order to perform maintenance.

In practice, you must protect all the keystores, whatever the location and whether they are password protected or not; protect the key repositories.

Refreshing the queue manager's key repository:

When you change the contents of a key repository, the queue manager does not immediately pick up the new contents. For a queue manager to use the new key repository contents, you must issue the REFRESH SECURITY TYPE(SSL) command.

This process is intentional, and prevents the situation where multiple running channels could use different versions of a key repository. As a security control, only one version of a key repository can be loaded by the queue manager at any time.

For more information about the REFRESH SECURITY TYPE(SSL) command, see  REFRESH SECURITY (WebSphere MQ V7.1 Reference).

You can also refresh a key repository using PCF commands or the WebSphere MQ Explorer. For more information, see  Refresh Security (WebSphere MQ V7.1 Reference) and the topic *Refreshing SSL or TLS Security* in the WebSphere MQ Explorer section of this product documentation.

Related concepts:

“Refreshing a client's view of the SSL key repository contents and SSL settings”

Refreshing a client's view of the SSL key repository contents and SSL settings:

To update the client application with the refreshed contents of the key repository, you must stop and restart the client application.

You cannot refresh security on a WebSphere MQ client; there is no equivalent of the REFRESH SECURITY TYPE(SSL) command for clients (see  REFRESH SECURITY (WebSphere MQ V7.1 Reference)) for more information.

You must stop and restart the application, whenever you change the security certificate, to update the client application with the refreshed contents of the key repository.

If restarting the channel refreshes the configurations, and if your application has reconnection logic, it is possible for you to refresh security at the client by issuing the STOP CHL STATUS(INACTIVE) command.

Related concepts:

“Refreshing the queue manager's key repository”

Digital Certificate Manager (DCM):

Use the DCM to manage digital certificates and private keys on IBM i.

The Digital Certificate Manager (DCM) enables you to manage digital certificates and to use them in secure applications on the IBM i server. With Digital Certificate Manager, you can request and process digital certificates from Certificate Authorities (CAs) or other third-parties. You can also act as a local Certificate Authority to create and manage digital certificates for your users.

DCM also supports using Certificate Revocation Lists (CRLs) to provide a stronger certificate and application validation process. You can use DCM to define the location where a specific Certificate Authority CRL resides on an LDAP server so that WebSphere MQ can verify that a specific certificate has not been revoked.

DCM supports and can automatically detect certificates in a variety of formats. When DCM detects a PKCS #12 encoded certificate, or a PKCS #7 certificate that contains encrypted data, it automatically

prompts the user to enter the password that was used to encrypt the certificate. DCM does not prompt for PKCS #7 certificates that do not contain encrypted data.

DCM provides a browser-based user interface that you can use to manage digital certificates for your applications and users. The user interface is divided into two main frames: a navigation frame and a task frame.

You use the navigation frame to select the tasks to manage certificates or the applications that use them. Some individual tasks are shown directly in the main navigation frame, but most tasks in the navigation frame are organized into categories. For example, Manage Certificates is a task category that contains various individual guided tasks, such as View certificate, Renew certificate, and Import certificate. If an item in the navigation frame is a category that contains more than one task, an arrow is displayed to the left of it. The arrow indicates that when you select the category link, an expanded list of tasks displays, enabling you to choose which task to perform.

For important information about DCM, see the following IBM Redbooks® publications:

- *IBM i Wired Network Security: OS/400® V5R1 DCM and Cryptographic Enhancements*, SG24-6168. Specifically, see the appendixes for essential information about setting up your IBM i system as a local CA.
- *AS/400 Internet Security: Developing a Digital Certificate Infrastructure*, SG24-5659. Specifically, see “Chapter 5. Digital Certificate Manager for AS/400”, which explains the AS/400 DCM.

Federal Information Processing Standards (FIPS):

This topic introduces the Federal Information Processing Standards (FIPS) Cryptomodule Validation Program of the US National Institute of Standards and Technology and the cryptographic functions which can be used on SSL or TLS channels, for Windows, UNIX and Linux, and z/OS systems.

The FIPS 140-2 compliance of a WebSphere MQ SSL or TLS connection on UNIX, Linux, and Windows systems is found here “Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows” on page 394.

The FIPS 140-2 compliance of a WebSphere MQ SSL or TLS connection on z/OS is found here “Federal Information Processing Standards (FIPS) for z/OS” on page 396.

If cryptographic hardware is present, the cryptographic modules used by WebSphere MQ can be configured to be those provided by the hardware manufacturer. If this is done, the configuration is only FIPS-compliant if those cryptographic modules are FIPS-certified.

Over time, the Federal Information Processing Standards are updated to reflect new attacks against encryption algorithms and protocols. For example, some CipherSpecs may cease to be FIPS certified. When such changes occur, WebSphere MQ is also updated to implement the latest standard. As a result, you might see changes in behavior after applying maintenance. The WebSphere MQ 7.1 readme lists the version of FIPS enforced by each product maintenance level. If you configure WebSphere MQ to enforce FIPS compliance, always consult the readme when planning to apply maintenance. The readme is located at <http://www-01.ibm.com/support/docview.wss?rs=171&uid=swg27006097>

Related concepts:

“Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client” on page 614

“Using iKeyman, iKeycmd, runmqakm, and runmqckm” on page 631

Related information:

 Enabling SSL in WebSphere MQ classes for Java (*WebSphere MQ V7.1 Programming Guide*)

 SSL properties of JMS objects (*WebSphere MQ V7.1 Reference*)

 Using Secure Sockets Layer (SSL) with WebSphere MQ classes for JMS (*WebSphere MQ V7.1 Programming Guide*)

“Federal Information Processing Standards” on page 382

Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows:

When cryptography is required on an SSL or TLS channel on Windows, UNIX and Linux systems, WebSphere MQ uses a cryptography package called IBM Crypto for C (ICC). On the Windows, UNIX and Linux platforms, the ICC software has passed the Federal Information Processing Standards (FIPS) Cryptomodule Validation Program of the US National Institute of Standards and Technology, at level 140-2.

The FIPS 140-2 compliance of a WebSphere MQ SSL or TLS connection on Windows, UNIX and Linux systems is as follows:

- For all WebSphere MQ message channels (except CLNTCONN channel types), the connection is FIPS-compliant if the following conditions are met:
 - The installed GSKit ICC version has been certified FIPS 140-2 compliant on the installed operating system version and hardware architecture.
 - The queue manager's SSLFIPS attribute has been set to YES.
 - All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.
- For all WebSphere MQ MQI client applications, the connection uses GSKit and is FIPS-compliant if the following conditions are met:
 - The installed GSKit ICC version has been certified FIPS 140-2 compliant on the installed operating system version and hardware architecture.
 - You have specified that only FIPS-certified cryptography is to be used, as described in the related topic for the MQI client.
 - All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.
- For WebSphere MQ classes for Java applications using client mode, the connection uses the JRE's SSL and TLS implementations and is FIPS-compliant if the following conditions are met:
 - The Java runtime environment used to run the application is FIPS-compliant on the installed operating system version and hardware architecture.
 - You have specified that only FIPS-certified cryptography is to be used, as described in the related topic for the Java client.
 - All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.
- For WebSphere MQ classes for JMS applications using client mode, the connection uses the JRE's SSL and TLS implementations and is FIPS-compliant if the following conditions are met:
 - The Java runtime environment used to run the application is FIPS-compliant on the installed operating system version and hardware architecture.
 - You have specified that only FIPS-certified cryptography is to be used, as described in the related topic for the JMS client.

- All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.
- For unmanaged .NET client applications, the connection uses GSKit and is FIPS-compliant if the following conditions are met:
 - The installed GSKit ICC version has been certified FIPS 140-2 compliant on the installed operating system version and hardware architecture.
 - You have specified that only FIPS-certified cryptography is to be used, as described in the related topic for the .NET client.
 - All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.
- For unmanaged XMS .NET client applications, the connection uses GSKit and is FIPS-compliant if the following conditions are met:
 - The installed GSKit ICC version has been certified FIPS 140-2 compliant on the installed operating system version and hardware architecture.
 - You have specified that only FIPS-certified cryptography is to be used, as described in the XMS .NET documentation.
 - All key repositories have been created and manipulated using only FIPS-compliant software, such as **runmqakm** with the **-fips** option.

All supported AIX, Linux, HP-UX, Solaris, Windows, and z/OS platforms are FIPS 140-2 certified except as noted in the readme file included with each fix pack or refresh pack.

For SSL and TLS connections using GSKit, the component which is FIPS 140-2 certified is named *ICC*. It is the version of this component which determines GSKit FIPS compliance on any given platform. To determine the ICC version currently installed, run the **dspmqrver -p 64 -v** command.

Here is an example extract of the **dspmqrver -p 64 -v** output relating to ICC:

```

ICC
=====
@(#)CompanyName:      IBM Corporation
@(#)LegalTrademarks:  IBM
@(#)FileDescription:  IBM Crypto for C-language
@(#)FileVersion:      8.0.0.0
@(#)LegalCopyright:   Licensed Materials - Property of IBM
@(#)                  ICC
@(#)                  (C) Copyright IBM Corp. 2002, 2019
@(#)                  All Rights Reserved. US Government Users
@(#)                  Restricted Rights - Use, duplication or disclosure
@(#)                  restricted by GSA ADP Schedule Contract with IBM Corp.
@(#)ProductName:      icc_8.0 (GoldCoast Build) 100415
@(#)ProductVersion:   8.0.0.0
@(#)ProductInfo:      10/04/15.03:32:19.10/04/15.18:41:51
@(#)CMVCInfo:

```

The NIST certification statement for GSKit ICC 8 (included in GSKit 8) can be found at the following address: <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2010.htm#1433>.

If cryptographic hardware is present, the cryptographic modules used by WebSphere MQ can be configured to be those provided by the hardware manufacturer. If this is done, the configuration is only FIPS-compliant if those cryptographic modules are FIPS-certified.

Note: 32-bit Solaris x86 SSL and TLS clients configured for FIPS 140-2 compliant operation fail when running on Intel systems. This failure occurs because the FIPS 140-2 compliant GSKit-Crypto Solaris x86

32-bit library file does not load on the Intel chipset. On affected systems, error AMQ9655 is reported in the client error log. To resolve this issue, disable FIPS 140-2 compliance or recompile the client application 64 bit, because 64-bit code is not affected.

Triple DES restrictions enforced when operating in compliance with FIPS 140-2

When WebSphere MQ is configured to operate in compliance with FIPS 140-2, additional restrictions are enforced in relation to Triple DES (3DES) CipherSpecs. These restrictions enable compliance with the US NIST SP800-67 recommendation.

1. All parts of the Triple DES key must be unique.
2. No part of the Triple DES key can be a Weak, Semi-Weak, or Possibly-Weak key according to the definitions in NIST SP800-67.
3. No more than 32 GB of data can be transmitted over the connection before a secret key reset must occur. By default, WebSphere MQ does not reset the secret session key so this reset must be configured. Failure to enable secret key reset when using a Triple DES CipherSpec and FIPS 140-2 compliance results in the connection closing with error AMQ9288 after the maximum byte count is exceeded. For information about how to configure secret key reset, see “Resetting SSL and TLS secret keys” on page 778.

WebSphere MQ generates Triple DES session keys which already comply with rules 1 and 2. However, to satisfy the third restriction you must enable secret key reset when using Triple DES CipherSpecs in a FIPS 140-2 configuration. Alternatively, you can avoid using Triple DES.

Related concepts:

“Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client” on page 614

“Using iKeyman, iKeycmd, runmqakm, and runmqckm” on page 631

Related information:

 Enabling SSL in WebSphere MQ classes for Java (*WebSphere MQ V7.1 Programming Guide*)

 SSL properties of JMS objects (*WebSphere MQ V7.1 Reference*)

 Using Secure Sockets Layer (SSL) with WebSphere MQ classes for JMS (*WebSphere MQ V7.1 Programming Guide*)

“Federal Information Processing Standards” on page 382

Federal Information Processing Standards (FIPS) for z/OS:

When cryptography is required on an SSL or TLS channel on z/OS, WebSphere MQ uses a service called System SSL. The objective of System SSL is to provide the capability to execute securely in a mode designed to adhere to the Federal Information Processing Standards (FIPS) Cryptomodule Validation Program of the US National Institute of Standards and Technology, at level 140-2.

When implementing a FIPS 140-2 compliant connection with WebSphere MQ SSL or TLS connection there are a number of points to consider:

- Ensure the conditions outlined below are met to enable WebSphere MQ message channels for FIPS-compliance:
 - System SSL Security Level 3 FMID (JCPT3B1) is installed and configured.
 - System SSL modules are validated.
 - The queue manager’s SSLFIPS attribute has been set to **YES**.

When executing in FIPS mode, System SSL exploits CP Assist for Cryptographic Function (CPACF) when available. Cryptographic functions performed by ICSF-supported hardware when running in non-FIPS mode continue to be exploited when executing in FIPS mode, with the exception of RSA signature generation which must be performed in software.

Table 12. Differences between FIPS mode and non-FIPS mode algorithm support.

Algorithm	Non-FIPS		FIPS	
	Sizes	Hardware	Sizes	Hardware
RC2	40 and 128			
RC4	40 and 128			
DES	56	x		
TDES	168	x	168	x
AES	128 and 256	x	128 and 256	x
MD5	48			
SHA-1	160	x	160	x
SHA-2	224, 256, 384 and 512	x	224, 256, 384 and 512	x
RSA	512-4096	x	1024-4096	x
DSA	512-1024		1024	
DH	512-2048		2048	

In FIPS mode, System SSL can only use certificates that use the algorithms and key sizes shown in Table 1. During X.509 certificate validation if an algorithm that is incompatible with FIPS mode is encountered, then the certificate cannot be used and is treated as not valid.

For WebSphere MQ classes applications using client mode within WebSphere Application Server, refer to the [Federal Information Processing Standards-approved](#) section of the IBM WebSphere Application Server product documentation.

For information on System SSL module configuration, see [System SSL Module Verification Setup](#).

Related information:

“Federal Information Processing Standards” on page 382

SSL and TLS on the WebSphere MQ MQI client:

WebSphere MQ supports SSL and TLS on clients. You can tailor the use of SSL or TLS in various ways.

WebSphere MQ provides SSL and TLS support for WebSphere MQ MQI clients on Windows, UNIX and Linux systems. If you are using WebSphere MQ classes for Java, see [Using WebSphere MQ classes for Java \(WebSphere MQ V7.1 Programming Guide\)](#) and if you are using WebSphere MQ classes for JMS, see [Using WebSphere MQ classes for JMS \(WebSphere MQ V7.1 Programming Guide\)](#). The rest of this section does not apply to the Java or JMS environments.

You can specify the key repository for a WebSphere MQ MQI client either with the MQSSLKEYR value in your WebSphere MQ client configuration file, or when your application makes an MQCONN call. You have three options for specifying that a channel uses SSL:

- Using a channel definition table
- Using the SSL configuration options structure, MQSCO, on an MQCONN call
- Using the Active Directory (on Windows systems)

You cannot use the MQSERVER environment variable to specify that a channel uses SSL.

You can continue to run your existing WebSphere MQ MQI client applications without SSL, as long as SSL is not specified at the other end of the channel.

If changes are made on a client machine to the contents of the SSL Key Repository, the location of the SSL Key Repository, the Authentication Information, or the Cryptographic hardware parameters, you need to end all the SSL connections in order to reflect these changes in the client-connection channels that the application is using to connect to the queue manager. Once all the connections have ended, restart the SSL channels. All the new SSL settings are used. These settings are analogous to those refreshed by the REFRESH SECURITY TYPE(SSL) command on queue manager systems.

When your WebSphere MQ MQI client runs on a Windows, UNIX and Linux system with cryptographic hardware, you configure that hardware with the MQSSLCRYP environment variable. This variable is equivalent to the SSLCRYP parameter on the ALTER QMGR MQSC command. Refer to  ALTER QMGR (*WebSphere MQ V7.1 Reference*) for a description of the SSLCRYP parameter on the ALTER QMGR MQSC command. If you use the GSK_PCS11 version of the SSLCRYP parameter, the PKCS #11 token label must be specified entirely in lower-case.

SSL secret key reset and FIPS are supported on WebSphere MQ MQI clients. For more information, see “Resetting SSL and TLS secret keys” on page 778 and “Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows” on page 394.

See “Setting up WebSphere MQ MQI client security” on page 613 for more information about the SSL support for WebSphere MQ MQI clients.

Related concepts:

 [Configuring a client using a configuration file \(*WebSphere MQ V7.1 Installing Guide*\)](#)

Specifying that an MQI channel uses SSL:

For an MQI channel to use SSL, the value of the *SSLCipherSpec* attribute of the client-connection channel must be the name of a CipherSpec that is supported by WebSphere MQ on the client platform.

You can define a client-connection channel with a value for this attribute in the following ways. They are listed in order of decreasing precedence.

1. When a PreConnect exit provides a channel definition structure to use.
A PreConnect exit can provide the name of a CipherSpec in the *SSLCipherSpec* field of a channel definition structure, MQCD. This structure is returned in the **ppMQCDArrayPtr** field of the MQNXP exit parameter structure used by the PreConnect exit.
2. When a WebSphere MQ MQI client application issues an MQCONN call.
The application can specify the name of a CipherSpec in the *SSLCipherSpec* field of a channel definition structure, MQCD. This structure is referenced by the connect options structure, MQCNO, which is a parameter on the MQCONN call.
3. Using a client channel definition table (CCDT).
One or more entries in a client channel definition table can specify the name of a CipherSpec. For example, if you create an entry by using the DEFINE CHANNEL MQSC command, you can use the SSLCIPH parameter on the command to specify the name of a CipherSpec.
4. Using Active Directory on Windows.
On Windows systems, you can use the **setmqscp** control command to publish the client-connection channel definitions in Active Directory. One or more of these definitions can specify the name of a CipherSpec.

For example, if a client application provides a client-connection channel definition in an MQCD structure on an MQCONN call, this definition is used in preference to any entries in a client channel definition table that can be accessed by the WebSphere MQ client.

You cannot use the MQSERVER environment variable to provide the channel definition at the client end of an MQI channel that uses SSL.

To check whether a client certificate has flowed, display the channel status at the server end of a channel for the presence of a peer name parameter value.

CipherSpecs and CipherSuites in IBM WebSphere MQ:

WebSphere MQ supports both SSL and TLS CipherSpecs, and RSA and Diffie-Hellman algorithms.

WebSphere MQ supports SSL V3 and TLS V1.0 and V1.2 CipherSpecs.

WebSphere MQ supports the RSA and Diffie-Hellman key exchange and authentication algorithms. The size of the key used during the SSL handshake can depend on the digital certificate you use, but some CipherSpecs include a specification of the handshake key size. Larger handshake key sizes provide stronger authentication. With smaller key sizes, the handshake is faster.

Related concepts:

“CipherSpecs and CipherSuites” on page 380

“An overview of the SSL or TLS handshake” on page 377

NSA Suite B Cryptography in IBM WebSphere MQ:

This topic provides information about how to configure WebSphere MQ on Windows, Linux, and UNIX systems to conform to the Suite B compliant TLS 1.2 profile.

Over time, the NSA Cryptography Suite B Standard is updated to reflect new attacks against encryption algorithms and protocols. For example, some CipherSpecs might cease to be Suite B certified. When such changes occur, IBM WebSphere MQ is also updated to implement the latest standard. As a result, you might see changes in behavior after applying maintenance. The IBM WebSphere MQ 7.1 readme file lists the version of Suite B enforced by each product maintenance level. If you configure WebSphere MQ to enforce Suite B compliance, always consult the readme file when planning to apply maintenance. The readme file is at <http://www-01.ibm.com/support/docview.wss?rs=171&uid=swg27006097>.

On Windows, UNIX, and Linux systems, WebSphere MQ can be configured to conform to the Suite B compliant TLS 1.2 profile at the security levels shown in Table 1.

Table 13. Suite B security levels with allowed CipherSpecs and digital signature algorithms

Security level	Allowed CipherSpecs	Allowed digital signature algorithms
128-bit	ECDHE_ECDSA_AES_128_GCM_SHA256 ECDHE_ECDSA_AES_256_GCM_SHA384	ECDSA with SHA-256 ECDSA with SHA-384
192-bit	ECDHE_ECDSA_AES_256_GCM_SHA384	ECDSA with SHA-384
Both ¹	ECDHE_ECDSA_AES_128_GCM_SHA256 ECDHE_ECDSA_AES_256_GCM_SHA384	ECDSA with SHA-256 ECDSA with SHA-384

1. It is possible to configure both the 128-bit and 192-bit security levels concurrently. Since the Suite B configuration determines the minimum acceptable cryptographic algorithms, configuring both security levels is equivalent to configuring only the 128-bit security level. The cryptographic algorithms of the 192-bit security level are stronger than the minimum required for the 128-bit security level, so they are permitted for the 128-bit security level even if the 192-bit security level is not enabled.

Note: The naming conventions used for the Security level do not necessarily represent the elliptic curve size or the key size of the AES encryption algorithm.

CipherSpec conformation to Suite B

Although the default behavior of IBM WebSphere MQ is not to comply with the Suite B standard, IBM WebSphere MQ can be configured to conform to either, or both security levels on Windows, UNIX and Linux systems. Following the successful configuration of IBM WebSphere MQ to use Suite B, any attempt to start an outbound channel using a CipherSpec not conforming to Suite B results in the error AMQ9282. This activity also results in the MQI client returning the reason code MQRC_CIPHER_SPEC_NOT_SUITE_B. Similarly, attempting to start an inbound channel using a CipherSpec not conforming to the Suite B configuration results in the error AMQ9616.

For more information about WebSphere MQ CipherSpecs, see “Specifying CipherSpecs” on page 774

Suite B and digital certificates

Suite B restricts the digital signature algorithms which can be used to sign digital certificates. Suite B also restricts the type of public key which certificates can contain. Therefore WebSphere MQ must be configured to use certificates whose digital signature algorithm and public key type are allowed by the configured Suite B security level of the remote partner. Digital certificates which do not comply with the security level requirements are rejected and the connection fails with error AMQ9633 or AMQ9285.

For the 128-bit Suite B security level, the public key of the certificate subject is required to use either the NIST P-256 elliptic curve or the NIST P-384 elliptic curve and to be signed with either the NIST P-256 elliptic curve or the NIST P-384 elliptic curve. At the 192-bit Suite B security level, the public key of the certificate subject is required to use the NIST P-384 elliptic curve and to be signed with the NIST P-384 elliptic curve.

To obtain a certificate suitable for Suite B compliant operation, use the **runmqakm** command and specify the **-sig_alg** parameter to request a suitable digital signature algorithm. The **EC_ecdsa_with_SHA256** and **EC_ecdsa_with_SHA384** **-sig_alg** parameter values correspond to elliptic curve keys signed by the allowed Suite B digital signature algorithms.

For more information about the **runmqakm** command, see  **runmqckm** and **runmqakm** options (*WebSphere MQ V7.1 Reference*).

Note: The **iKeycmd** and **iKeyman** tools do not support the creation of digital certificates for Suite B compliant operation.

Creating and requesting digital certificates

To create a self-signed digital certificate for Suite B testing, see “Creating a self-signed personal certificate on UNIX, Linux, and Windows systems” on page 639

To request a CA-signed digital certificate for Suite B production use, see “Requesting a personal certificate on UNIX, Linux, and Windows systems” on page 641.

Note: The certificate authority being used must generate digital certificates which satisfy the requirements described in IETF RFC 6460.

FIPS 140-2 and Suite B

The Suite B standard is conceptually similar to FIPS 140-2, as it restricts the set of enabled cryptographic algorithms in order to provide an assured level of security. The Suite B CipherSpecs currently supported can be used when IBM WebSphere MQ is configured for FIPS 140-2 compliant operation. It is therefore possible to configure WebSphere MQ for both FIPS and Suite B compliance simultaneously, in which case both sets of restrictions apply.

The following diagram illustrates the relationship between these subsets:



Configuring WebSphere MQ for Suite B compliant operation

For information about how to configure IBM WebSphere MQ on Windows, UNIX and Linux for Suite B compliant operation, see “Configuring WebSphere MQ for Suite B.”

IBM WebSphere MQ does not support Suite B compliant operation on the IBM i and z/OS platforms. The WebSphere MQ Java and JMS clients also do not support Suite B compliant operation.

Related concepts:

“Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client” on page 614

Configuring WebSphere MQ for Suite B:

IBM WebSphere MQ can be configured to operate in compliance with the NSA Suite B standard on Windows, UNIX and Linux platforms.

Suite B restricts the set of enabled cryptographic algorithms in order to provide an assured level of security. IBM WebSphere MQ can be configured to operate in compliance with Suite B to provide an enhanced level of security. For further information on Suite B, see “National Security Agency (NSA) Suite B Cryptography” on page 382. For more information about Suite B configuration and its effect on SSL and TLS channels, see “NSA Suite B Cryptography in IBM WebSphere MQ” on page 399.

Queue manager

For a queue manager, use the command **ALTER QMGR** with the parameter **SUITEB** to set the values appropriate for your required level of security. For further information see  **ALTER QMGR** (*WebSphere MQ V7.1 Reference*).

You can also use the PCF **MQCMD_CHANGE_Q_MGR** command with the **MQIA_SUITE_B_STRENGTH** parameter to configure the queue manager for Suite B compliant operation

MQI client

By default, MQI clients do not enforce Suite B compliance. You can enable the MQI client for Suite B compliance by executing one of the options below:

1. By setting the **EncryptionPolicySuiteB** field in the MQSCO structure on an MQCONN call to one or more of the values below:
 - MQ_SUITE_B_NONE
 - MQ_SUITE_B_128_BIT
 - MQ_SUITE_B_192_BITUsing MQ_SUITE_B_NONE with any other value is invalid.
2. By setting the MQSUITEB environment variable to one or more of the values below:
 - NONE
 - 128_BIT

- 192_BIT

You can specify multiple values using a comma separated list. Using the value NONE with any other value is invalid.

3. By setting the **EncryptionPolicySuiteB** attribute in the SSL stanza of the MQI client configuration file to one or more of the values below:

- NONE
- 128_BIT
- 192_BIT

You can specify multiple values using a comma separated list. Using NONE with any other value is invalid.

Note: The MQI client settings are listed in order of priority. The MSCO structure on the MQCONN call overrides the setting on the MQSUIEB environment variable, which overrides the attribute in the SSL stanza.

For full details of the MQSCO structure, see  MQSCO – SSL configuration options (*WebSphere MQ V7.1 Reference*).

For more information about the use of Suite B in the client configuration file, see  SSL stanza of the client configuration file (*WebSphere MQ V7.1 Installing Guide*).

For further information on the use of the MQSUIEB environment variable, see  Environment Variables (*WebSphere MQ V7.1 Reference*).

.NET

For .NET unmanaged clients, the property **MQC. ENCRYPTION_POLICY_SUITE_B** indicates the type of Suite B security required.

For information about the using Suite B in IBM WebSphere MQ classes for .NET, see  MQEnvironment .NET class (*WebSphere MQ V7.1 Reference*).

Certificate validation policies in WebSphere MQ:

The certificate validation policy determines how strictly the certificate chain validation conforms to industry security standards.

The certificate validation policy depends upon the platform and environment as follows:

- For Java and JMS applications on all platforms, the certificate validation policy depends on the JSSE component of the Java runtime environment. For more information about the certificate validation policy, see the documentation for your JRE.
- For IBM i systems, the certificate validation policy depends on the secure sockets library provided by the operating system. For more information about the certificate validation policy, see the documentation for the operating system.
- For z/OS systems, the certificate validation policy depends on the System SSL component provided by the operating system. For more information about the certificate validation policy, see the documentation for the operating system.
- For UNIX, Linux, and Windows systems, the certificate validation policy is supplied by GSKit and can be configured. Two different certificate validation policies are supported:

- A legacy certificate validation policy, used for maximum backwards compatibility and interoperability with old digital certificates that do not comply with the current IETF certificate validation standards. This policy is known as the Basic policy.
- A strict, standards-compliant certificate validation policy which enforces the RFC 5280 standard. This policy is known as the Standard policy.

For information about how to configure the certificate validation policy on UNIX, Linux, and Windows systems, see “Configuring certificate validation policies in WebSphere MQ.” For more information about the differences between the Basic and Standard certificate validation policies, see “Certificate validation and trust policy design on UNIX, Linux, and Windows systems” on page 660 (*WebSphere MQ V7.1 Reference*).

Configuring certificate validation policies in WebSphere MQ:

You can specify which SSL/TLS certificate validation policy is used to validate digital certificates received from remote partner systems in four ways.

On the queue manager, the certificate validation policy can be set in the following ways:

- Using the queue manager attribute *CERTVPOL*. For more information about setting this attribute, see  *ALTER QMGR (WebSphere MQ V7.1 Reference)*.

This attribute can be used only on queue managers at command level 711, or higher. For more information about command levels and enabling fix pack function, see  *New function in maintenance level upgrades (WebSphere MQ V7.1 Installing Guide)*.

On the client, there are several methods that can be used to set the certificate validation policy. If more than one method is used to set the policy, the client uses the settings in the following priority order:

1. Using the *CertificateValPolicy* field in the client MQSCO structure. For more information about using this field, see  *MQSCO – SSL configuration options (WebSphere MQ V7.1 Reference)*.
2. Using the client environment variable, *MQCERTVPOL*. For more information about using this variable, see  *MQCERTVPOL (WebSphere MQ V7.1 Installing Guide)*.
3. Using the client SSL stanza tuning parameter setting, *CertificateValPolicy*. For more information about using this setting, see  *SSL stanza of the client configuration file (WebSphere MQ V7.1 Installing Guide)*.

For more information about certificate validation policies, see “Certificate validation policies in WebSphere MQ” on page 402.

Digital certificates and CipherSpec compatibility in IBM WebSphere MQ:

This topic provides information on how to choose appropriate CipherSpecs and digital certificates for your security policy, by outlining the relationship between CipherSpecs and digital certificates in IBM WebSphere MQ.

In previous releases of IBM WebSphere MQ, all supported SSL and TLS CipherSpecs used the RSA algorithm for digital signatures and key agreement. All of the supported types of digital certificate were compatible with all of the supported CipherSpecs, so it was possible to change the CipherSpec for any channel without needing to change digital certificates.

In IBM WebSphere MQ Version 7.1 only a subset of the supported CipherSpecs can be used with all of the supported types of digital certificate. It is therefore necessary to choose an appropriate CipherSpec for your digital certificate. Similarly, if your organization's security policy requires that you use a particular CipherSpec then you must obtain an appropriate digital certificate for that CipherSpec.

The MD5 digital signature algorithm and TLS 1.2

Digital certificates signed using the MD5 algorithm are rejected when the TLS 1.2 protocol is used. This is because the MD5 algorithm is now considered weak by many cryptographic analysts and its use is generally discouraged. If you wish to use newer CipherSpecs based on the TLS 1.2 protocol, ensure that the digital certificates do not use the MD5 algorithm in their digital signatures. Older CipherSpecs which use the SSL 3.0 and TLS 1.0 protocols are not subject to this restriction and can continue to use certificates with MD5 digital signatures.

To view the digital signature algorithm for a particular certificate, you can use the **runmqakm** command:

```
runmqakm -cert -details -db key.kdb -pw password -label cert_label
```

where `cert_label` is the certificate label of the digital signature algorithm you need to display.

Note: Although the **iKeycmd (runmqckm)** tool and the **iKeyman (strmqikm)** GUI can be used to view a selection of digital signature algorithms, the **runmqakm** tool provides a wider range.

The execution of the **runmqakm** command will produce output displaying the use of the signature algorithm specified:

```
Label : ibmwebspheremqexample
Key Size : 1024
Version : X509 V3
Serial : 4e4e93f1
Issuer : CN=Old Certificate Authority,OU=Test,O=Example,C=US
Subject : CN=Example Queue Manager,OU=Test,O=Example,C=US
Not Before : August 19, 2011 5:48:49 PM GMT+01:00
Not After : August 18, 2012 5:48:49 PM GMT+01:00
Public Key
 30 81 9F 30 0D 06 09 2A 86 48 86 F7 0D 01 01 01
 05 00 03 81 8D 00 30 81 89 02 81 81 00 98 5A 7A
 F0 18 21 EE E4 8A 6E DE C8 01 4B 3A 1E 41 90 3D
 CE 01 3F E6 32 30 6C 23 59 F0 FE 78 6D C2 80 EF
 BC 83 54 7A EB 60 80 62 6B F1 52 FE 51 9D C1 61
 80 A5 1C D4 F0 76 C7 15 6D 1F 0D 4D 31 3E DC C6
 A9 20 84 6E 14 A1 46 7D 4C F5 79 4D 37 54 0A 3B
 A9 74 ED E7 8B 0F 80 31 63 1A 0B 20 A5 99 EE 0A
 30 A6 B6 8F 03 97 F6 99 DB 6A 58 89 7F 27 34 DE
 55 08 29 D8 A9 6B 46 E6 02 17 C3 13 D3 02 03 01
 00 01
Public Key Type : RSA (1.2.840.113549.1.1.1)
Fingerprint : SHA1 :
 09 4E 4F F2 1B CB C1 F4 4F 15 C9 2A F7 32 0A 82
 DA 45 92 9F
Fingerprint : MD5 :
 44 54 81 7C 58 68 08 3A 5D 75 96 40 D5 8C 7A CB
Fingerprint : SHA256 :
 3B 47 C6 E7 7B B0 FF 85 34 E7 48 BE 11 F2 D4 35
 B7 9A 79 53 2B 07 F5 E7 65 E8 F7 84 E0 2E 82 55
Signature Algorithm : MD5WithRSASignature (1.2.840.113549.1.1.4)
Value
 3B B9 56 E6 F2 77 94 69 5B 3F 17 EA 7B 19 D0 A2
 D7 10 38 F1 88 A4 44 1B 92 35 6F 3B ED 99 9B 3A
 A5 A4 FC 72 25 5A A9 E3 B1 96 88 FC 1E 9F 9B F1
 C5 E8 8E CF C4 8F 48 7B 0E A6 BB 13 AE 2B BD D8
 63 2C 03 38 EF DC 01 E1 1F 7A 6F FB 2F 65 74 D0
```

```

FD 99 94 BA B2 3A D5 B4 89 6C C1 2B 43 6D E2 39
66 6A 65 CB C3 C4 E2 CC F5 49 39 A3 8B 93 5A DD
B0 21 0B A8 B2 59 5B 24 59 50 44 89 DC 78 19 51
Trust Status : Enabled

```

The Signature Algorithm line shows that the MD5WithRSASignature algorithm is used. This algorithm is based upon MD5 and so this digital certificate cannot be used with the TLS 1.2 CipherSpecs.

Interoperability of Elliptic Curve and RSA CipherSpecs

Not all CipherSpecs can be used with all digital certificates. There are three types of CipherSpec, denoted by the CipherSpec name prefix. Each type of CipherSpec imposes different restrictions upon the type of digital certificate which may be used. These restrictions apply to all WebSphere MQ SSL and TLS connections, but are particularly relevant to users of Elliptic Curve cryptography.

The relationships between CipherSpecs and digital certificates are summarized in the following table:

Table 14. Relationships between CipherSpecs and digital certificates

Type	CipherSpec Name Prefix	Description	Required public key type	Digital signature encryption algorithm	Secret key establishment method
1	ECDHE_ECDSA_	CipherSpecs which use Elliptic Curve public keys, Elliptic Curve secret keys, and Elliptic Curve digital signature algorithms.	Elliptic Curve	ECDSA	ECDHE
2	ECDHE_RSA_	CipherSpecs which use RSA public keys, Elliptic Curve secret keys, and Elliptic Curve digital signature algorithms.	RSA	RSA	ECDHE
3	(All others)	CipherSpecs which use RSA public keys and RSA digital signature algorithms.	RSA	RSA	RSA

Note: Type 1 and 2 CipherSpecs are only supported by WebSphere MQ queue managers and MQI clients on the UNIX, Linux, and Windows platforms.

The required public key type column shows the type of public key which the personal certificate must have when using each type of CipherSpec. The personal certificate is the end-entity certificate which identifies the queue manager or client to its remote partner.

The digital signature encryption algorithm refers to the encryption algorithm used to validate the peer. The encryption algorithm is used along with a hash algorithm such as MD5, SHA-1 or SHA-256 to compute the digital signature. There are various digital signature algorithms which can be used, for example "RSA with MD5" or "ECDSA with SHA-256". In the table, ECDSA refers to the set of digital

signature algorithms which use ECDSA; RSA refers to the set of digital signature algorithms which use RSA. Any supported digital signature algorithm in the set may be used, provided it is based upon the stated encryption algorithm.

Type 1 CipherSpecs require that the personal certificate must have an Elliptic Curve public key. When these CipherSpecs are used, Elliptic Curve Diffie Hellman Ephemeral key agreement is used to establish the secret key for the connection.

Type 2 CipherSpecs require that the personal certificate has an RSA public key. When these CipherSpecs are used, Elliptic Curve Diffie Hellman Ephemeral key agreement is used to establish the secret key for the connection.

Type 3 CipherSpecs require that the the personal certificate must have an RSA public key. When these CipherSpecs are used, RSA key exchange is used to establish the secret key for the connection.

This list of restrictions is not exhaustive: depending on the configuration, there might be additional restrictions which can further affect the ability to interoperate. For example, if WebSphere MQ is configured to comply with the FIPS 140-2 or NSA Suite B standards then this will also limit the range of allowable configurations. Refer to the following section for more information.

A WebSphere MQ queue manager can only use a single personal certificate to identify itself. This means all channels on the queue manager will use the same digital certificate and therefore each queue manager may only use one type of CipherSpec at a time. Similarly, a WebSphere MQ client application can only use a single personal certificate to identify itself. This means that all SSL and TLS connections within a single application process will use the same digital certificate and therefore each client application process may only use one type of CipherSpec at a time.

The three types of CipherSpec do not interoperate directly: this is a limitation of the current SSL and TLS standards. For example, suppose you have chosen to use the ECDHE_ECDSA_AES_128_CBC_SHA256 CipherSpec for a receiver channel named TO.QM1 on a queue manager named QM1. ECDHE_ECDSA_AES_128_CBC_SHA256 is a Type 1 CipherSpec, so QM1 must have a personal certificate with an Elliptic Curve key and an ECDSA-based digital signature. All clients and other queue managers which communicate directly with QM1 must therefore have digital certificates which satisfy the Type 1 CipherSpec requirements. Other channels connecting to queue manager QM1 may use other CipherSpecs (for example ECDHE_ECDSA_3DES_EDE_CBC_SHA256), but they may only use Type 1 CipherSpecs to communicate with QM1.

When planning your WebSphere MQ networks, consider carefully which channels require SSL or TLS and ensure that all of the clients and queue managers which need to interoperate use the same type of CipherSpecs and appropriate digital certificates. The IETF standards RFC 4492, RFC 5246 and RFC 6460 describe the detailed usage of Elliptic Curve CipherSpecs in TLS 1.2.

To view the digital signature algorithm and public key type for a digital certificate you can use the **runmqakm** command:

```
runmqakm -cert -details -db key.kdb -pw password -label cert_label
```

where `cert_label` is the label of the certificate whose digital signature algorithm you need to display.

The execution of the **runmqakm** command will produce output displaying the Public Key Type:

```
Label : ibmwebspheremqexample  
Key Size : 384  
Version : X509 V3  
Serial : 9ad5eeef5d756f41  
Issuer : CN=Example Certificate Authority,OU=Test,O=Example,C=US  
Subject : CN=Example Queue Manager,OU=Test,O=Example,C=US
```

```

Not Before : 21 August 2011 13:10:24 GMT+01:00
Not After  : 21 August 2012 13:10:24 GMT+01:00
Public Key
  30 76 30 10 06 07 2A 86 48 CE 3D 02 01 06 05 2B
  81 04 00 22 03 62 00 04 3E 6F A9 06 B6 C3 A0 11
  F8 D6 22 78 FE EF 0A FE 34 52 C0 8E AB 5E 81 73
  D0 97 3B AB D6 80 08 E7 31 E9 18 3F 6B DE 06 A7
  15 D6 9D 5B 6F 56 3B 7F 72 BB 6F 1E C9 45 1C 46
  60 BE F2 DC 1B AD AC EC 64 4C 0E 06 65 6E ED 93
  B8 F5 95 E0 F9 2A 05 D6 21 02 BD FB 06 63 A1 CC
  66 C6 8A 0A 5C 3F F7 D3
Public Key Type : EC_ecPublicKey (1.2.840.10045.2.1)
Fingerprint : SHA1 :
  3C 34 58 04 5B 63 5F 5C C9 7A E7 67 08 2B 84 43
  3D 43 7A 79
Fingerprint : MD5 :
  49 13 13 E1 B2 AC 18 9A 31 41 DC 8C B4 D6 06 68
Fingerprint : SHA256 :
  6F 76 78 68 F3 70 F1 53 CE 39 31 D9 05 C5 C5 9F
  F2 B8 EE 21 49 16 1D 90 64 6D AC EB 0C A7 74 17
Signature Algorithm : EC_ecdsa_with_SHA384 (1.2.840.10045.4.3.3)
Value
  30 65 02 30 0A B0 2F 72 39 9E 24 5A 22 FE AC 95
  0D 0C 6D 6C 2F B3 E7 81 F6 C1 36 1B 9A B0 6F 07
  59 2A A1 4C 02 13 7E DD 06 D6 FE 4B E4 03 BC B1
  AC 49 54 1E 02 31 00 90 0E 46 2B 04 37 EE 2C 5F
  1B 9C 69 E5 99 60 84 84 10 71 1A DA 63 88 33 E2
  22 CC E6 1A 4E F4 61 CC 51 F9 EE A0 8E F4 DC B5
  0B B9 72 58 C3 C7 A4
Trust Status : Enabled

```

The Public Key Type line in this case shows that the certificate has an Elliptic Curve public key. The Signature Algorithm line in this case shows that the EC_ecdsa_with_SHA384 algorithm is in use: this is based upon the ECDSA algorithm. This certificate is therefore only suitable for use with Type 1 CipherSpecs.

You can also use the **iKeycmd (runmqckm)** tool with the same parameters. Also the **iKeyman (strmqikm)** GUI can be used to view digital signature algorithms if you open the key repository and double-click the label of the certificate. However, you are advised to use the **runmqakm** tool to view digital certificates because it supports a wider range of algorithms.

Elliptic Curve CipherSpecs and NSA Suite B

When WebSphere MQ is configured to conform to the Suite B compliant TLS 1.2 profile, the permitted CipherSpecs and digital signature algorithms are restricted as described in “NSA Suite B Cryptography in IBM WebSphere MQ” on page 399. Additionally, the range of acceptable Elliptic Curve keys is reduced according to the configured security levels.

At the 128-bit Suite B security level, the certificate subject's public key is required to use either the NIST P-256 or NIST P-384 elliptic curve and to be signed with either the NIST P-256 elliptic curve or the NIST P-384 elliptic curve. The **runmqakm** command can be used to request digital certificates for this security level using a **-sig_alg** parameter of EC_ecdsa_with_SHA256, or EC_ecdsa_with_SHA384.

At the 192-bit Suite B security level, the certificate subject's public key is required to use the NIST P-384 elliptic curve and to be signed with the NIST P-384 elliptic curve. The **runmqakm** command can be used to request digital certificates for this security level using a **-sig_alg** parameter of EC_ecdsa_with_SHA384.

The supported NIST elliptic curves are as follows:

Table 15. Supported NIST elliptic curves

NIST FIPS 186-3 curve name	RFC 4492 curve name	Elliptic Curve key size (bits)
P-256	secp256r1	256
P-384	secp384r1	384
P-521	secp521r1	521

Note: The NIST P-521 elliptic curve cannot be used for Suite B compliant operation.

Related concepts:

“Specifying CipherSpecs” on page 774

“Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client” on page 614

“NSA Suite B Cryptography in IBM WebSphere MQ” on page 399

“National Security Agency (NSA) Suite B Cryptography” on page 382

Channel authentication records

To exercise more precise control over the access granted to connecting systems at a channel level, you can use channel authentication records.

You might find that clients attempt to connect to your queue manager using a blank user ID or a high-level user ID that would allow the client to perform undesirable actions. You can block access to these clients using channel authentication records. Alternatively, a client might assert a user ID that is valid on the client platform but is unknown or of an invalid format on the server platform. You can use a channel authentication record to map the asserted user ID to a valid user ID.

You might find a client application that connects to your queue manager and behaves badly in some way. To protect the server from the issues this application is causing, it needs to be temporarily blocked using the IP address the client application is on until such time as the firewall rules are updated or the client application is corrected. You can use a channel authentication record to block the IP address from which the client application connects.

If you have set up an administration tool such as the IBM WebSphere MQ Explorer, and a channel for that specific use, you might want to ensure that only specific client computers can use it. You can use a channel authentication record to allow the channel to be used only from certain IP addresses.

If you are just getting started with some sample applications running as clients, see  Preparing and running the sample programs (*WebSphere MQ V7.1 Programming Guide*) for an example of setting up the queue manager securely using channel authentication records.

Use of channel authentication records to control inbound channels is enabled using the **ALTER QMGR CHLAUTH(ENABLED)** switch.

Channel authentication records can be created to perform the following functions:

- To block connections from specific IP addresses.
- To block connections from specific user IDs.
- To set an MCAUSER value to be used for any channel connecting from a specific IP address.
- To set an MCAUSER value to be used for any channel asserting a specific user ID.
- To set an MCAUSER value to be used for any channel having a specific SSL or TLS Distinguished Name (DN).
- To set an MCAUSER value to be used for any channel connecting from a specific queue manager.
- To block connections claiming to be from a certain queue manager unless the connection is from a specific IP address.

- To block connections presenting a certain SSL or TLS certificate unless the connection is from a specific IP address.

These uses are explained further in the following sections.

You create, modify, or remove channel authentication records using the MQSC command **SET CHLAUTH** or the PCF command **Set Channel Authentication Record**.

Blocking IP addresses

It is normally the role of a firewall to prevent access from certain IP addresses. However, there might be occasions where you experience connection attempts from an IP address that should not have access to your WebSphere MQ system and must temporarily block the address before the firewall can be updated. These connection attempts might not even be coming from WebSphere MQ channels, but from other socket applications that are misconfigured to target your WebSphere MQ listener. Block IP addresses by setting a channel authentication record of type **BLOCKADDR**. You can specify one or more single addresses, ranges of addresses, or patterns including wildcards.

Whenever an inbound connection is refused because the IP address is blocked in this manner, an event message **MQRChannelBlocked** with reason qualifier **MQRChannelBlockedAddress** is issued, provided that channel events are enabled and the queue manager is running. Additionally, the connection is held open for 30 seconds prior to returning the error to ensure the listener is not flooded with repeated attempts to connect that are blocked.

To block IP addresses only on specific channels, or to avoid the delay before the error is reported, set a channel authentication record of type **ADDRESSMAP** with the **USERSRC(NOACCESS)** parameter.

Whenever an inbound connection is refused for this reason, an event message **MQRChannelBlocked** with reason qualifier **MQRChannelBlockedNoAccess** is issued, provided that channel events are enabled and the queue manager is running.

See “Blocking specific IP addresses” on page 733 for an example.

Blocking user IDs

To prevent certain user IDs from connecting over a client channel, set a channel authentication record of type **BLOCKUSER**. This type of channel authentication record applies only to client channels, not to message channels. You can specify one or more individual user IDs to be blocked, but you cannot use wildcards.

Whenever an inbound connection is refused for this reason, an event message **MQRChannelBlocked** with reason qualifier **MQRChannelBlockedUserID** is issued, provided that channel events are enabled.

See “Blocking specific user IDs” on page 735 for an example.

You can also block any access for specified user IDs on certain channels by setting a channel authentication record of type **USERMAP** with the **USERSRC(NOACCESS)** parameter.

Whenever an inbound connection is refused for this reason, an event message **MQRChannelBlocked** with reason qualifier **MQRChannelBlockedNoAccess** is issued, provided that channel events are enabled and the queue manager is running.

See “Blocking access for a client asserted user ID” on page 738 for an example.

Blocking queue manager names

To specify that any channel connecting from a specified queue manager is to have no access, set a channel authentication record of type QMGRMAP with the USERSRC(NOACCESS) parameter. You can specify a single queue manager name or a pattern including wildcards. There is no equivalent of the BLOCKUSER function to block access from queue managers.

Whenever an inbound connection is refused for this reason, an event message MQRQ_CHANNEL_BLOCKED with reason qualifier MQRQ_CHANNEL_BLOCKED_NOACCESS is issued, provided that channel events are enabled and the queue manager is running.

See “Blocking access from a remote queue manager” on page 737 for an example.

Blocking SSL or TLS DNs

To specify that any user presenting an SSL or TLS personal certificate containing a specified DN is to have no access, set a channel authentication record of type SSLPEERMAP with the USERSRC(NOACCESS) parameter. You can specify a single distinguished name or a pattern including wildcards. There is no equivalent of the BLOCKUSER function to block access for DNs.

Whenever an inbound connection is refused for this reason, an event message MQRQ_CHANNEL_BLOCKED with reason qualifier MQRQ_CHANNEL_BLOCKED_NOACCESS is issued, provided that channel events are enabled and the queue manager is running.

See “Blocking access for an SSL Distinguished Name” on page 738 for an example.

Mapping IP addresses to user IDs to be used

To specify that any channel connecting from a specified IP address is to use a specific MCAUSER, set a channel authentication record of type ADDRESSMAP. You can specify a single address, a range of addresses, or a pattern including wildcards.

If you use a port forwarder, DMZ session break, or any other setup which changes the IP address presented to the queue manager, then mapping IP addresses is not necessarily suitable for your use.

See “Mapping an IP address to an MCAUSER user ID” on page 739 for an example.

Mapping queue manager names to user IDs to be used

To specify that any channel connecting from a specified queue manager is to use a specific MCAUSER, set a channel authentication record of type QMGRMAP. You can specify a single queue manager name or a pattern including wildcards.

See “Mapping a remote queue manager to an MCAUSER user ID” on page 735 for an example.

Mapping user IDs asserted by a client to user IDs to be used

To specify that if a certain user ID is used by a connection from a WebSphere MQ MQI client, a different, specified MCAUSER is to be used, set a channel authentication record of type USERMAP. User ID mapping does not use wildcards.

See “Mapping a client asserted user ID to an MCAUSER user ID” on page 736 for an example.

Mapping SSL or TLS DNs to user IDs to be used

To specify that any user presenting an SSL/TLS personal certificate containing a specified DN is to use a specific MCAUSER, set a channel authentication record of type SSLPEERMAP. You can specify a single distinguished name or a pattern including wildcards.

See “Mapping an SSL or TLS Distinguished Name to an MCAUSER user ID” on page 737 for an example.

Mapping queue managers, clients, or SSL or TLS DNs according to IP address

In some circumstances it might be possible for a third party to spoof a queue manager name. An SSL or TLS certificate or key database file might also be stolen and reused. To protect against these threats, you can specify that a connection from a certain queue manager or client, or using a certain DN must be connecting from a specified IP address. Set a channel authentication record of type USERMAP, QMGRMAP or SSLPEERMAP and specify the permitted IP address, or pattern of IP addresses, using the ADDRESS parameter.

See “Mapping a remote queue manager to an MCAUSER user ID” on page 735 for an example.

Interaction between channel authentication records

It is possible that a channel attempting to make a connection matches more than one channel authentication record, and that these have contradictory effects. For example, a channel might assert a user ID which is blocked by a BLOCKUSER channel authentication record, but with an SSL or TLS certificate that matches an SSLPEERMAP record that sets a different user ID. In addition, if channel authentication records use wildcards, a single IP address, queue manager name, or SSL or TLS DN might match several patterns. For example, the IP address 192.0.2.6 matches the patterns 192.0.2.0-24, 192.0.2.*, and 192.0.*.6. The action taken is determined as follows.

- The channel authentication record used is selected as follows:
 - A channel authentication record explicitly matching the channel name takes priority over a channel authentication record matching the channel name by using a wildcard.
 - A channel authentication record using an SSL or TLS DN takes priority over a record using a user ID, queue manager name, or IP address.
 - A channel authentication record using a user ID or queue manager name takes priority over a record using an IP address.
- If a matching channel authentication record is found and it specifies an MCAUSER, this MCAUSER is assigned to the channel.
- If a matching channel authentication record is found and it specifies that the channel has no access, an MCAUSER value of *NOACCESS is assigned to the channel. This value can later be changed by a security exit program.
- If no matching channel authentication record is found, or a matching channel authentication record is found and it specifies that the user ID of the channel is to be used, the MCAUSER field is inspected.
 - If the MCAUSER field is blank, the client user ID is assigned to the channel.
 - If the MCAUSER field is not blank, it is assigned to the channel.
- Any security exit program is run. This exit program might set the channel user ID or determine that access is to be blocked.
- If the connection is blocked or the MCAUSER is set to *NOACCESS, the channel ends.
- If the connection is not blocked, for any channel except a client channel, the channel user ID determined in the previous steps is checked against the list of blocked users.
 - If the user ID is in the list of blocked users, the channel ends.
 - If the user ID is not in the list of blocked users, the channel runs.

Where a number of channel authentication records match a channel name, IP address, queue manager name, or SSL or TLS DN, the most specific match is used. The match considered to be most specific is determined as follows.

- For a channel name:
 - The most specific match is a name without wildcards, for example A.B.C.
 - The most generic match is a single asterisk (*), which matches all channel names.
 - A pattern with an asterisk in the left-most position is more generic than a pattern with a defined value in the left-most position. Thus *.B.C is more generic than A.*.
 - A pattern with an asterisk in the second position is more generic than a pattern with a defined value in the second position, and similarly for each subsequent position. Thus A.*.C is more generic than A.B.*
 - Where two or more patterns have an asterisk in the same position, the one with fewer nodes following the asterisk is more generic. Thus A.* is more generic than A.*.C
- For an IP address:
 - The most specific match is a name without wildcards, for example 192.0.2.6.
 - The most generic match is a single asterisk (*), which matches all channel names.
 - A pattern with an asterisk in the left-most position is more generic than a pattern with a defined value in the left-most position. Thus *.0.2.6 is more generic than 192.*.
 - A pattern with an asterisk in the second position is more generic than a pattern with a defined value in the second position, and similarly for each subsequent position. Thus 192.*.2.6 is more generic than 192.0.*.
 - Where two or more patterns have an asterisk in the same position, the one with fewer nodes following the asterisk is more generic. Thus 192.* is more generic than 192.*.2.*.
 - A range indicated with a hyphen (-), is more specific than an asterisk. Thus 192.0.2.0-24 is more specific than 192.0.2.*.
 - A range that is a subset of another is more specific than the larger range. Thus 192.0.2.5-15 is more specific than 192.0.2.0-24.
 - Overlapping ranges are not permitted. For example, you cannot have channel authentication records for both 192.0.2.0-15 and 192.0.2.10-20.
 - A pattern cannot have fewer than the required number of parts, unless the pattern ends with a single trailing asterisk. For example 192.0.2 is invalid, but 192.0.2.* is valid.
 - A trailing asterisk must be separated from the rest of the address by the appropriate part separator (a dot (.) for IPv4, a colon (:)) for IPv6). For example, 192.0.* is not valid because the asterisk is not in a part of its own.
 - A pattern may contain additional asterisks provided that no asterisk is adjacent to the trailing asterisk. For example, 192.*.2.* is valid, but 192.0.*.* is not valid.
 - An IPv6 address pattern cannot contain a double colon and a trailing asterisk, because the resulting address would be ambiguous. For example, 2001::.* could expand to 2001:0000:.*, 2001:0000:0000:.* and so on
- For a queue manager name:
 - The most specific match is a name without wildcards, for example 192.0.2.6.
 - The most generic match is a single asterisk (*), which matches all channel names.
 - A pattern with an asterisk in the left-most position is more generic than a pattern with a defined value in the left-most position. Thus *QUEUEMANAGER is more generic than QUEUEMANAGER*.
 - A pattern with an asterisk in the second position is more generic than a pattern with a defined value in the second position, and similarly for each subsequent position. Thus Q*MANAGER is more generic than QUEUE*.
 - Where two or more patterns have an asterisk in the same position, the one with fewer characters following the asterisk is more generic. Thus Q* is more generic than Q*MGR.

- For an SSL or TLS Distinguished Name (DN), the precedence order of substrings is as follows:

Order	DN substring	Name
1	SERIALNUMBER=	Certificate serial number
2	MAIL=	Email address
3	E=	Email address (Deprecated in preference to MAIL)
4	UID=, USERID=	User identifier
5	CN=	Common name
6	T=	Title
7	OU=	Organizational unit
8	DC=	Domain component
9	O=	Organization
10	STREET=	Street / First line of address
11	L=	Locality
12	ST=, SP=, S=	State or province name
13	PC=	Postal code / zip code
14	C=	Country
15	UNSTRUCTUREDNAME=	Host name
16	UNSTRUCTUREDADDRESS=	IP address
17	DNQ=	Distinguished name qualifier

Thus, if an SSL or TLS certificate is presented with a DN containing the substrings O=IBM and C=UK, WebSphere MQ uses a channel authentication record for O=IBM in preference to one for C=UK, if both are present.

A DN can contain multiple OUs, which must be specified in hierarchical order with the large organizational units specified first. If two DNs are equal in all respects except for their OU values, the more specific DN is determined as follows:

1. If they have different numbers of OU attributes then the DN with the most OU values is more specific. This is because the DN with more Organizational Units fully qualifies the DN in more detail and provides more matching criteria. Even if its top-level OU is a wildcard (OU=*), the DN with more OUs is still regarded as more specific overall.
2. If they have the same number of OU attributes then the corresponding pairs of OU values are compared in sequence left-to-right, where the left-most OU is the highest-level (least specific), according to the following rules.
 - a. An OU with no wildcard values is the most specific because it can only match exactly one string.
 - b. An OU with a single wildcard at either the beginning or end (for example, OU=ABC* or OU=*ABC) is next most specific.
 - c. An OU with two wildcards for example, OU=*ABC*) is next most specific.
 - d. An OU consisting only of an asterisk (OU=*) is the least specific.
3. If the string comparison is tied between two attribute values of the same specificity then whichever attribute string is longer is more specific.
4. If the string comparison is tied between two attribute values of the same specificity and length then the result is determined by a case-insensitive string comparison of the portion of the DN excluding any wildcards.

If two DNs are equal in all respects except for their DC values, the same matching rules apply as for OUs except that in DC values the left-most DC is the lowest-level (most specific) and the comparison ordering differs accordingly.

Displaying channel authentication records

To display channel authentication records, use the MQSC command **DISPLAY CHLAUTH** or the PCF command **Inquire Channel Authentication Records**. You can choose to return all records that match the supplied channel name, or you can choose an explicit match. The explicit match tells you which channel authentication record would be used if a channel attempted to make a connection from a specific IP address, from a specific queue manager or using a specific user ID, and, optionally, presenting an SSL/TLS personal certificate containing a specified DN.

Related concepts:

 WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*)

Related tasks:

“Securing remote connectivity to the queue manager” on page 732

Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

 DISPLAY CHLAUTH (*WebSphere MQ V7.1 Reference*)

 Set Channel Authentication Record (*WebSphere MQ V7.1 Reference*)

 Inquire Channel Authentication Records (*WebSphere MQ V7.1 Reference*)

Planning for your security requirements

This collection of topics explains what you need to consider when planning security in a WebSphere MQ environment.

You can use WebSphere MQ for a wide variety of applications on a range of platforms. The security requirements are likely to be different for each application. For some, security will be a critical consideration.

WebSphere MQ provides a range of link-level security services, including support for the Secure Sockets Layer (SSL) and Transport Layer Security (TLS).

You must consider certain aspects of security when implementing WebSphere. On IBM i, UNIX, Linux and Windows systems, if you ignore these aspects and do nothing, you cannot use WebSphere MQ. On z/OS, the effect of ignoring these aspects is that your WebSphere MQ resources are unprotected. That is, all users can access and change all WebSphere MQ resources.

Security considerations are described below.

Authority to administer WebSphere MQ

WebSphere MQ administrators need authority to:

- Issue commands to administer WebSphere MQ
- Use the WebSphere MQ Explorer
- Use the operations and control panels on z/OS
- Use IBM i administrative panels and commands.
- Use the WebSphere MQ utility program, CSQUTIL, on z/OS

- Access the queue manager data sets on z/OS

For more information, see:

- “Authority to administer WebSphere MQ on UNIX, Linux and Windows systems” on page 751
- “Authority to administer WebSphere MQ on IBM i” on page 419
- “Authority to administer WebSphere MQ on z/OS” on page 420

Authority to work with WebSphere MQ objects

Applications can access the following WebSphere MQ objects by issuing MQI calls:

- Queue managers
- Queues
- Processes
- Namelists
- Topics

Applications can also use Programmable Command Format (PCF) commands to access these WebSphere MQ objects, and to access channels and authentication information objects as well. These objects can be protected by WebSphere MQ so that the user IDs associated with the applications need authority to access them.

For more information, see “Authorization for applications to use WebSphere MQ” on page 422.

Channel security

The user IDs associated with message channel agents (MCAs) need authority to access various WebSphere MQ resources. For example, an MCA must be able to connect to a queue manager. If it is a sending MCA, it must be able to open the transmission queue for the channel. If it is a receiving MCA, it must be able to open destination queues. The user IDs associated with applications which need to administer channels, channel initiators, and listeners need authority to use the relevant PCF commands. However, most applications do not need such access.

For more information, see “Channel security” on page 447.

Additional considerations

You need to consider the following aspects of security only if you are using certain WebSphere MQ function or base product extensions:

- “Security for queue manager clusters” on page 459
- “Security for WebSphere MQ Publish/Subscribe” on page 460
- “Security for WebSphere MQ internet pass-thru” on page 461

Planning identification and authentication

Decide what user IDs to use, and how and at what levels you want to apply authentication controls.

You must decide how you will identify the users of your IBM WebSphere MQ applications, bearing in mind that different operating systems support user IDs of different lengths. You can use channel authentication records to map from one user ID to another, or to specify a user ID based on some attribute of the connection. WebSphere MQ channels using SSL or TLS use digital certificates as a mechanism for identification and authentication. Each digital certificate has a subject distinguished name which can be mapped onto specific identities using channel authentication records. Additionally, CA certificates in the key repository determine which digital certificates may be used to authenticate to WebSphere MQ. For more information see:

- “Mapping a remote queue manager to an MCAUSER user ID” on page 735
- “Mapping a client asserted user ID to an MCAUSER user ID” on page 736
- “Mapping an SSL or TLS Distinguished Name to an MCAUSER user ID” on page 737
- “Mapping an IP address to an MCAUSER user ID” on page 739

Planning authentication for a client application

You can apply authentication controls at four levels: at the communications level, in security exits, with channel authentication records, and in terms of the identification that is passed to a security exit.

There are four levels of security to consider. The diagram shows a WebSphere MQ MQI client that is connected to a server. Security is applied at four levels, as described in the following text. MCA is a Message Channel Agent.

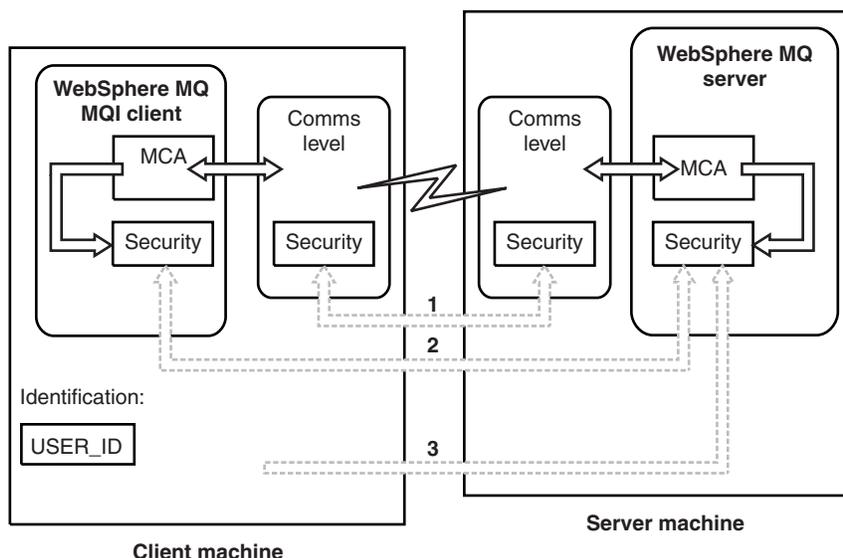


Figure 75. Security in a client/server connection

1. Communications level

See arrow 1. To implement security at the communications level, use SSL or TLS. For more information, see “Cryptographic security protocols: SSL and TLS” on page 376

2. Channel authentication records

See arrows 2 & 3. Authentication can be controlled by using the IP address or SSL/TLS distinguished names at the security level. A user ID can also be blocked or an asserted user ID can be mapped to a valid user ID. A full description is given in “Channel authentication records” on page 408.

3. Channel security exits

See arrow 2. The channel security exits for client to server communication can work in the same way as for server to server communication. A protocol independent pair of exits can be written to provide mutual authentication of both the client and the server. A full description is given in  Channel security exit programs (*WebSphere MQ V7.1 Programming Guide*).

4. Identification that is passed to a channel security exit

See arrow 3. In client to server communication, the channel security exits do not have to operate as a pair. The exit on the IBM WebSphere MQ client side can be omitted. In this case, the user ID is placed in the channel descriptor (MQCD) and the server-side security exit can alter it, if required.

Windows clients also send extra information to assist identification.

- The user ID that is passed to the server is the currently logged-on user ID on the client.

- The security ID of the currently logged-on user.

To assist identification on IBM WebSphere MQ client for HP Integrity NonStop Server, the client passes the OSS Safeguard alias under which the client application is running. This ID is typically of the form <PRIMARYGROUP>.<ALIAS>. If required, you can map this user ID to an alternative user ID on the queue manager by using either channel authentication records or a security exit. For more information about message exits, see “Identity mapping in message exits” on page 693. For more information about defining channel authentication records, see “Mapping a client asserted user ID to an MCAUSER user ID” on page 736.

The values of the user ID and, if available, the security ID, can be used by the server security exit to establish the identity of the IBM WebSphere MQ MQI client.

User IDs:

If the WebSphere MQ MQI client is on Windows and the WebSphere MQ server is also on Windows and has access to the domain on which the client user ID is defined, WebSphere MQ supports user IDs of up to 20 characters. On UNIX and Linux platforms and configurations, the maximum length is 12 characters. On IBM i, it is 10 characters.

A WebSphere MQ for Windows server does not support the connection of a Windows client if the client is running under a user ID that contains the @ character, for example, abc@d. The return code to the **MQCONN** call at the client is **MQRC_NOT_AUTHORIZED**.

However, you can specify the user ID using two @ characters, for example, abc@@d. Using the id@domain format is the preferred practice, to ensure that the user ID is resolved in the correct domain consistently; thus abc@@d@domain.

Note that UNKNOWN is a reserved user ID and the NOBODY user ID also have special meanings to WebSphere MQ. Creating user IDs in the operating system called UNKNOWN or NOBODY could have unintended results.

Although user IDs are used to authenticate, groups are used for authorization, except for Windows.

If you create service accounts, without paying attention to groups, and authorize all the user IDs differently, every user can access the information of every other user.

Planning authorization

Plan the users who will have administrative authority and plan how to authorize users of applications to appropriately use IBM WebSphere MQ objects, including those connecting from a IBM WebSphere MQ MQI client.

Individuals or applications must be granted access in order to use IBM WebSphere MQ. What access they require depend on the roles they undertake and the tasks which they need to perform. Authorization in IBM WebSphere MQ can be subdivided into two main categories:

- Authorization to perform administrative operations
- Authorization for applications to use IBM WebSphere MQ

Both classes of operation are controlled by the same component and an individual can be granted authority to perform both categories of operation.

The following topics give further information about specific areas of authorization that you must consider:

Authority to administer WebSphere MQ

WebSphere MQ administrators need authority to perform various functions. This authority is obtained in different ways on different platforms.

WebSphere MQ administrators need authority to:

- Issue commands to administer WebSphere MQ
- Use the WebSphere MQ Explorer
- Use the operations and control panels on z/OS
- Use the WebSphere MQ utility program, CSQUTIL, on z/OS
- Access the queue manager data sets on z/OS

For more information, see the topic appropriate to your operating system.

Related concepts:

“Authority to administer WebSphere MQ on z/OS” on page 420

“Authority to administer WebSphere MQ on UNIX, Linux and Windows systems” on page 751

Related reference:

“WebSphere MQ authorities on IBM i” on page 494

Authority to administer WebSphere MQ on UNIX and Windows systems:

A IBM WebSphere MQ administrator is a member of the `mqm` group. This group has access to all IBM WebSphere MQ resources and can issue IBM WebSphere MQ control commands. An administrator can grant specific authorities to other users.

To be a WebSphere MQ administrator on UNIX and Windows systems, a user must be a member of the *mqm group*. This group is created automatically when you install WebSphere MQ. To allow users to issue control commands, you must add them to the `mqm` group. This includes the root user on UNIX systems.

Users who are not member of the `mqm` group can be granted administrative privileges, but they are not able to issue IBM WebSphere MQ control commands, and they are authorized to execute only the commands for which they have been granted access.

Additionally, on Windows systems, the `SYSTEM` and `Administrator` accounts have full access to IBM WebSphere MQ resources.

All members of the `mqm` group have access to all WebSphere MQ resources on the system, including being able to administer any queue manager running on the system. This access can be revoked only by removing a user from the `mqm` group. On Windows systems, members of the `Administrators` group also have access to all WebSphere MQ resources.

Administrators can use the control command `runmqsc` to issue WebSphere MQ Script (MQSC) commands. When `runmqsc` is used in indirect mode to send MQSC commands to a remote queue manager, each MQSC command is encapsulated within an Escape PCF command. Administrators must have the required authorities for the MQSC commands to be processed by the remote queue manager.

The WebSphere MQ Explorer issues PCF commands to perform administration tasks. Administrators require no additional authorities to use the WebSphere MQ Explorer to administer a queue manager on the local system. When the WebSphere MQ Explorer is used to administer a queue manager on another system, administrators must have the required authorities for the PCF commands to be processed by the remote queue manager.

For more information about the authority checks carried out when PCF and MQSC commands are processed, see the following topics:

- For commands that operate on queue managers, queues, channels, processes, namelists, and authentication information objects, see “Authorization for applications to use WebSphere MQ” on page 422.
- For commands that operate on channels, channel initiators, listeners, and clusters, see  Channel security.
- For MQSC commands that are processed by the command server on WebSphere MQ for z/OS, see “Command security and command resource security” on page 420.

For more information about the authority you need to administer WebSphere MQ on UNIX and Windows systems, see the related information.

Authority to administer WebSphere MQ on IBM i:

To be a WebSphere MQ administrator on IBM i, you must be a member of the *QMADM* group. This group has properties similar to those of the *mqm* group on UNIX and Windows systems. In particular, the *QMADM* group is created when you install WebSphere MQ for IBM i, and members of the *QMADM* group have access to all WebSphere MQ resources on the system. You also have access to all WebSphere MQ resources if you have **ALLOBJ* authority.

Administrators can use CL commands to administer WebSphere MQ. One of these commands is *GRTMQMAUT*, which is used to grant authorities to other users. Another command, *STRMQMMQSC*, enables an administrator to issue MQSC commands to a local queue manager.

There are two groups of CL command provided by WebSphere MQ for IBM i:

Group 1

To issue a command in this category, a user must be a member of the *QMADM* group or have **ALLOBJ* authority. *GRTMQMAUT* and *STRMQMMQSC* belong to this category, for example.

Group 2

To issue a command in this category, a user does not need to be a member of the *QMADM* group or have **ALLOBJ* authority. Instead, two levels of authority are required:

- The user requires IBM i authority to use the command. This authority is granted by using the *GRTOBJAUT* command.
- The user requires WebSphere MQ authority to access any WebSphere MQ object associated with the command. This authority is granted by using the *GRTMQMAUT* command.

The following are examples of commands in this group:

- *CRTMQMQ*, Create MQM Queue
- *CHGMQMPC*, Change MQM Process
- *DLTMQMNL*, Delete MQM Namelist
- *DSPMQMAUTI*, Display MQM Authentication Information
- *CRTMQMCHL*, Create MQM channel

For more information about this group of commands, see “Authorization for applications to use WebSphere MQ” on page 422.

For a complete list of group 1 and group 2 commands, see “Access authorities for IBM WebSphere MQ objects on IBM i” on page 494

For more information about the authority you need to administer WebSphere MQ on IBM i, see “Administering IBM i” on page 168.

Related reference:

“WebSphere MQ authorities on IBM i” on page 494

Authority to administer WebSphere MQ on z/OS:

This collection of topics describes various aspects of the authority you need to administer WebSphere MQ for z/OS.

Authority checks on z/OS:

WebSphere MQ uses the System Authorization Facility (SAF) to route requests for authority checks to an external security manager (ESM) such as the z/OS Security Server Resource Access Control Facility (RACF). WebSphere MQ does no authority checks of its own.

It is assumed that you are using RACF as your ESM. If you are using a different ESM, you might need to interpret the information provided for RACF in a way that is relevant to your ESM.

You can specify whether you want authority checks turned on or off for each queue manager individually or for every queue manager in a queue-sharing group. This level of control is called *subsystem security*. If you turn subsystem security off for a particular queue manager, no authority checks are carried out for that queue manager.

If you turn subsystem security on for a particular queue manager, authority checks can be performed at two levels:

Queue-sharing group level security

Authority checks use RACF profiles that are shared by all queue managers in the queue-sharing group. This means that there are fewer profiles to define and maintain, making security administration easier.

Queue manager level security

Authority checks use RACF profiles specific to the queue manager.

You can use a combination of queue-sharing group and queue manager level security. For example, you can arrange for profiles specific to a queue manager to override those of the queue-sharing group to which it belongs.

Subsystem security, queue-sharing group level security, and queue manager level security are turned on or off by defining *switch profiles*. A switch profile is a normal RACF profile that has a special meaning to WebSphere MQ.

Command security and command resource security:

Command security relates to the authority to issue a command; command resource authority relates to the authority to perform an operation on a resource. Both are implemented by using RACF classes.

Authority checks are carried out when a WebSphere MQ administrator issues an MQSC command. This is called *command security*.

To implement command security, you must define certain RACF profiles and give the necessary groups and user IDs access to these profiles at the required levels. The name of a profile for command security contains the name of an MQSC command.

Some MQSC commands perform an operation on a WebSphere MQ resource, such as the DEFINE QLOCAL command to create a local queue. When an administrator issues an MQSC command, authority checks are carried out to determine whether the requested operation can be performed on the resource specified in the command. This is called *command resource security*.

To implement command resource security, you must define certain RACF profiles and give the necessary groups and user IDs access to these profiles at the required levels. The name of a profile for command resource security contains the name of a WebSphere MQ resource and its type (QUEUE, PROCESS, NAMELIST, TOPIC, AUTHINFO, or CHANNEL).

Command security and command resource security are independent. For example, when an administrator issues the command:

```
DEFINE QLOCAL(MOON.EUROPA)
```

the following authority checks are performed:

- Command security checks that the administrator is authorized to issue the DEFINE QLOCAL command.
- Command resource security checks that the administrator is authorized to perform an operation on the local queue called MOON.EUROPA.

Command security and command resource security can be turned on or off by defining switch profiles.

MQSC commands and the system command input queue:

Use this topic to understand how the command server processes MQSC commands directed to the system command input queue.

Command security and command resource security are also used when the command server retrieves a message containing an MQSC command from the system command input queue. The user ID that is used for the authority checks is the one found in the *UserIdentifier* field in the message descriptor of the message containing the MQSC command. This user ID must have the required authorities on the queue manager where the command is processed. For more information about the *UserIdentifier* field and how it is set, see Message context.

Messages containing MQSC commands are sent to the system command input queue in the following circumstances:

- The operations and control panels send MQSC commands to the system command input queue of the target queue manager. The MQSC commands correspond to the actions you choose on the panels. The *UserIdentifier* field in each message is set to the TSO user ID of the administrator.
- The COMMAND function of the WebSphere MQ utility program, CSQUTIL, sends the MQSC commands in the input data set to the system command input queue of the target queue manager. The COPY and EMPTY functions send DISPLAY QUEUE and DISPLAY STGCLASS commands. The *UserIdentifier* field in each message is set to the job user ID.
- The MQSC commands in the CSQINPX data sets are sent to the system command input queue of the queue manager to which the channel initiator is connected. The *UserIdentifier* field in each message is set to the channel initiator address space user ID.

No authority checks are performed when MQSC commands are issued from the CSQINP1 and CSQINP2 data sets. You can control who is allowed to update these data sets using RACF data set protection.

- Within a queue-sharing group, a channel initiator might send START CHANNEL commands to the system command input queue of the queue manager to which it is connected. A command is sent when an outbound channel that uses a shared transmission queue is started by triggering. The *UserIdentifier* field in each message is set to the channel initiator address space user ID.
- An application can send MQSC commands to a system command input queue. By default, the *UserIdentifier* field in each message is set to the user ID associated with the application.
- On UNIX, Linux and Windows systems, the **runmqsc** control command can be used in indirect mode to send MQSC commands to the system command input queue of a queue manager on z/OS. The *UserIdentifier* field in each message is set to the user ID of the administrator who issued the **runmqsc** command.

Access to the queue manager data sets:

WebSphere MQ administrators need authority to access the queue manager data sets. Use this topic to understand which data sets need RACF protection.

These data sets include:

- The data sets referred to by CSQINP1, CSQINP2, CSQINPT, and CSQXLIB in the queue manager's started task procedure
- The queue manager's page sets, active log data sets, archive log data sets, and bootstrap data sets (BSDSs)
- The data sets referred to by CSQXLIB and CSQINPX in the channel initiator's started task procedure

You must protect the data sets so that no unauthorized user can start a queue manager or gain access to any queue manager data. To do this, use RACF data set protection.

Authorization for applications to use WebSphere MQ

When applications access objects, the user IDs associated with the applications need appropriate authority.

Applications can access the following WebSphere MQ objects by issuing MQI calls:

- Queue managers
- Queues
- Processes
- Namelists
- Topics

Applications can also use PCF commands to administer WebSphere MQ objects. When the PCF command is processed, it uses the authority context of the user ID that put the PCF message.

Applications, in this context, include those written by users and vendors, and those supplied with WebSphere MQ for z/OS. The applications supplied with WebSphere MQ for z/OS include:

- The operations and control panels
- The WebSphere MQ utility program, CSQUTIL
- The dead letter queue handler utility, CSQUDLQH

Applications that use WebSphere MQ classes for Java, WebSphere MQ classes for JMS, WebSphere MQ classes for .NET, or the Message Service Clients for C/C++ and .NET use the MQI indirectly.

MCAs also issue MQI calls and the user IDs associated with the MCAs need authority to access these WebSphere MQ objects. For more information about these user IDs and the authorities they require, see

 Channel security.

On z/OS, applications can also use MQSC commands to access these WebSphere MQ objects but command security and command resource security provide the authority checks in these circumstances. For more information, see “Command security and command resource security” on page 420 and “MQSC commands and the system command input queue” on page 421.

On IBM i, a user that issues a CL command in Group 2 might require authority to access a WebSphere MQ object associated with the command. For more information, see “When authority checks are performed” on page 423.

When authority checks are performed:

Authority checks are performed when an application attempts to access a queue manager, queue, process, or namelist.

On IBM i, authority checks might also be performed when a user issues a CL command in Group 2 that accesses any of these WebSphere MQ objects. The checks are performed in the following circumstances:

When an application connects to a queue manager using an MQCONN or MQCONNX call

The queue manager asks the operating system for the user ID associated with the application. The queue manager then checks that the user ID is authorized to connect to it and retains the user ID for future checks.

Users do not have to sign on to WebSphere MQ. WebSphere MQ assumes that users are signed on to the underlying operating system and are been authenticated by it.

When an application opens a WebSphere MQ object using an MQOPEN or MQPUT1 call

All authority checks are performed when an object is opened, not when it is accessed later. For example, authority checks are performed when an application opens a queue. They are not performed when the application puts messages on the queue or gets messages from the queue.

When an application opens an object, it specifies the types of operation it needs to perform on the object. For example, an application might open a queue to browse the messages on it, get messages from it, but not to put messages on it. For each type of operation, the queue manager checks that the user ID associated with the application has the authority to perform that operation.

When an application opens a queue, the authority checks are performed against the object named in the `ObjectName` field of the object descriptor. The `ObjectName` field is used on the `MQOPEN` or `MQPUT1` calls. If the object is an alias queue or a remote queue definition, the authority checks are performed against the object itself. They are not performed on the queue to which the alias queue or the remote queue definition resolves. This means that the user does not need permission to access it. Limit the authority to create queues to privileged users. If you do not, users might bypass the normal access control simply by creating an alias.

An application can reference a remote queue explicitly. It sets the `ObjectName` and `ObjectQMgrName` fields in the object descriptor to the names of the remote queue and the remote queue manager. The authority checks are performed against the transmission queue with the same name as the remote queue manager. On z/OS, a check is made on the RACF queue profile that matches the remote queue manager name. On platforms other than z/OS, a check is made against the RQMNAME profile that matches the remote queue manager name, if clustering is being used. An application can reference a cluster queue explicitly by setting the `ObjectName` field in the object descriptor to the name of the cluster queue. The authority checks are performed against the cluster transmission queue, `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

The authority to a dynamic queue is based on the model queue from which it is derived, but is not necessarily the same; see note 1.

The user ID that the queue manager uses for the authority checks is obtained from the operating system. The user ID is obtained when the application connects to the queue manager. A suitably authorized application can issue an `MQOPEN` call specifying an alternative user ID; access control checks are then made on the alternative user ID. Using an alternate user ID does not change the user ID associated with the application, only the one used for access control checks.

When an application subscribes to a topic using an MQSUB call

When an application subscribes to a topic, it specifies the type of operation that it needs to perform. It is either creating a subscription, altering an existing subscription, or resuming an existing subscription without changing it. For each type of operation, the queue manager checks that the user ID that is associated with the application has the authority to perform the operation.

When an application subscribes to a topic, the authority checks are performed against topic objects that are found in the topic tree. The topic objects are at, or above, the point in the topic tree at which the application subscribed. The authority checks might involve checks on more than one topic object. The user ID that the queue manager uses for the authority checks is obtained from the operating system. The user ID is obtained when the application connects to the queue manager.

The queue manager performs authority checks on subscriber queues but not on managed queues.

When an application deletes a permanent dynamic queue using an MQCLOSE call

The object handle specified on the MQCLOSE call is not necessarily the same one returned by the MQOPEN call that created the permanent dynamic queue. If it is different, the queue manager checks the user ID associated with the application that issued the MQCLOSE call. It checks that the user ID is authorized to delete the queue.

When an application that closes a subscription to remove it did not create it, the appropriate authority is required to remove it.

When a PCF command that operates on a WebSphere MQ object is processed by the command server

This rule includes the case where a PCF command operates on an authentication information object.

The user ID that is used for the authority checks is the one found in the `UserIdentifier` field in the message descriptor of the PCF command. This user ID must have the required authorities on the queue manager where the command is processed. The equivalent MQSC command encapsulated within an Escape PCF command is treated in the same way. For more information about the `UserIdentifier` field, and how it is set, see “Message context” on page 425.

On IBM i, when a user issues a CL command in Group 2 that operates on a WebSphere MQ object

This rule includes the case where a CL command in Group 2 operates on an authentication information object.

Checks are performed to determine whether the user has the authority to operate on a WebSphere MQ object associated with the command. The checks are performed unless the user is a member of the QMQMADM group or has *ALL0BJ authority. The authority required depends on the type of operation that the command performs on the object. For example, the command **CHGMQMQ**, Change MQM Queue, requires the authority to change the attributes of the queue specified by the command. In contrast, the command **DSPMQMQ**, Display MQM Queue, requires the authority to display the attributes of the queue specified by the command.

Many commands operate on more than one object. For example, to issue the command **DLTMQMQ**, Delete MQM Queue, the following authorities are required:

- The authority to connect to the queue manager specified by the command
- The authority to delete the queue specified by the command

Some commands operate on no object at all. In this case, the user requires only IBM i authority to issue one of these commands. **STRMQLSR**, Start MQM Listener, is an example of such a command.

Alternate user authority:

When an application opens an object or subscribes to a topic, the application can supply a user ID on the MQOPEN, MQPUT1, or MQSUB call. It can ask the queue manager to use this user ID for authority checks instead of the one associated with the application.

The application succeeds in opening the object only if both the following conditions are met:

- The user ID associated with the application has the authority to supply a different user ID for authority checks. The application is said to have *alternate user authority*.
- The user ID supplied by the application has the authority to open the object for the types of operation requested, or to subscribe to the topic.

Message context:

Message context information allows the application that retrieves a message to find out about the originator of the message. The information is held in fields in the message descriptor and the fields are divided into three logical parts

These parts are as follows:

identity context

These fields contain information about the user of the application that put the message on the queue.

origin context

These fields contain information about the application itself and when the message was put on the queue.

user context

These fields contain message properties that applications can use to select messages that the queue manager should deliver.

When an application puts a message on a queue, the application can ask the queue manager to generate the context information in the message. This is the default action. Alternatively, it can specify that the context fields are to contain no information. The user ID associated with an application requires no special authority to do either of these.

An application can set the identity context fields in a message, allowing the queue manager to generate the origin context, or it can set all the context fields. An application can also pass the identity context fields from a message it has retrieved to a message it is putting on a queue, or it can pass all the context fields. However, the user ID associated with an application requires authority to set or pass context information. An application specifies that it intends to set or pass context information when it opens the queue on which it is about to put messages, and its authority is checked at this time.

Here is a brief description of each of the context fields:

Identity context

UserIdentifier

The user ID associated with the application that put the message. If the queue manager sets this field, it is set to the user ID obtained from the operating system when the application connects to the queue manager.

AccountingToken

Information that can be used to charge for the work done as a result of the message.

ApplIdentityData

If the user ID associated with an application has authority to set the identity context fields, or to set all the context fields, the application can set this field to any value related to identity. If the queue manager sets this field, it is set to blank.

Origin context

PutApplType

The type of the application that put the message; a CICS transaction, for example.

PutApplName

The name of the application that put the message.

PutDate

The date when the message was put.

PutTime

The time when the message was put.

ApplOriginData

If the user ID associated with an application has authority to set all the context fields, the application can set this field to any value related to origin. If the queue manager sets this field, it is set to blank.

User context

The following values are supported for **MQINQMP** or **MQSETMP**:

MQPD_USER_CONTEXT

The property is associated with the user context.

No special authorization is required to be able to set a property associated with the user context using the **MQSETMP** call.

On a V7.0 or subsequent queue manager, a property associated with the user context is saved as described for **MQOO_SAVE_ALL_CONTEXT**. An **MQPUT** with **MQOO_PASS_ALL_CONTEXT** specified causes the property to be copied from the saved context into the new message.

MQPD_NO_CONTEXT

The property is not associated with a message context.

An unrecognized value is rejected with **MQRC_PD_ERROR**. The initial value of this field is **MQPD_NO_CONTEXT**.

For a detailed description of each of the context fields, see  **MQMD – Message descriptor** (*WebSphere MQ V7.1 Reference*). For more information about how to use message context, see  **Message context** (*WebSphere MQ V7.1 Programming Guide*).

Authority to work with IBM WebSphere MQ objects on IBM i, UNIX, Linux and Windows systems:

The authorization service component provided with IBM WebSphere MQ is called the *object authority manager (OAM)*. It provides access control via authentication and authorization checks.

1. Authentication.

The authentication check performed by the OAM provided with IBM WebSphere MQ is basic, and is only performed in specific circumstances. It is not intended to meet the strict requirements expected in a highly secure environment.

The OAM performs its authentication check when an application connects to a queue manager, and the following conditions are true.

If an **MQCSP** structure has been supplied by the connecting application, and the *AuthenticationType* attribute in the **MQCSP** structure is given the value **MQCSP_AUTH_USER_ID_AND_PWD**, then the check is performed by the OAM in its **MQZID_AUTHENTICATE_USER** function. This is the check: the user ID in the **MQCSP** structure is compared against the user ID in the *IdentityContext* (**MQZIC**), to determine whether they match. If they do not match, the check fails.

This basic check is not intended to be a full authentication of the user. For example, there is no check of the authenticity of the user by checking the password supplied in the **MQCSP** structure. Also, if the application omits an **MQCSP** structure, then no check is performed.

If fuller authentication services are required in the queue manager via the authorization service component, then the OAM provided with IBM WebSphere MQ does not offer this. You must write a new authorization service component, or obtain one from a vendor.

2. Authorization.

The authorization checks are comprehensive, and are intended to meet most normal requirements.

Authorization checks are performed when an application issues an MQI call to access a queue manager, queue, process, topic, or namelist. They are also performed at other times, for example, when a command is being performed by the Command Server.

On IBM i, UNIX, Linux and Windows systems, the *authorization service* provides the access control when an application issues an MQI call to access a IBM WebSphere MQ object that is a queue manager, queue, process, topic, or namelist. This includes checks for alternative user authority and the authority to set or pass context information.

As with other versions of Windows, the OAM gives members of the Administrators group the authority to access all IBM WebSphere MQ objects even when UAC is enabled on Windows Vista and Windows Server 2008.

Additionally, on Windows systems, the SYSTEM account has full access to IBM WebSphere MQ resources.

The authorization service also provides authority checks when a PCF command operates on one of these IBM WebSphere MQ objects or an authentication information object. The equivalent MQSC command encapsulated within an Escape PCF command is treated in the same way.

On IBM i, unless the user is a member of the QMQMADM group or has *ALLOBJ authority, the authorization service also provides authority checks when a user issues a CL command in Group 2 that operates on any of these IBM WebSphere MQ objects or an authentication information object.

The authorization service is an *installable service*, which means that it is implemented by one or more *installable service components*. Each component is invoked using a documented interface. This enables users and vendors to provide components to augment or replace those provided by the IBM WebSphere MQ products.

The authorization service component provided with IBM WebSphere MQ is called the *object authority manager (OAM)*. The OAM is automatically enabled for each queue manager you create.

The OAM maintains an access control list (ACL) for each IBM WebSphere MQ object it is controlling access to. On UNIX and Linux systems, only group IDs can appear in an ACL. This means that all members of a group have the same authorities. On IBM i and on Windows systems, both user IDs and group IDs can appear in an ACL. This means that authorities can be granted to individual users and groups.

A 12 character limitation applies to both the group and the user ID. UNIX platforms generally restrict the length of a user ID to 12 characters. AIX and Linux have raised this limit but IBM WebSphere MQ continues to observe a 12 character restriction on all UNIX platforms. If you use a user ID of greater than 12 characters, IBM WebSphere MQ replaces it with the value "UNKNOWN". Do not define a user ID with a value of "UNKNOWN".

The OAM can authenticate a user and change appropriate identity context fields. You enable this by specifying a connection security parameters structure (MQCSP) on an MQCONN call. The structure is passed to the OAM Authenticate User function (MQZ_AUTHENTICATE_USER), which sets appropriate identity context fields. In the case of an MQCONN connection from a IBM WebSphere MQ client, the information in the MQCSP is flowed to the queue manager to which the client is connecting over the client-connection and server-connection channel. If security exits are defined on that channel, the MQCSP is passed into each security exit and can be altered by the exit. Security exits can also create the MQCSP.

For more details of the use of security exits in this context, see  Channel security exit programs (*WebSphere MQ V7.1 Programming Guide*).

On UNIX, Linux and Windows systems, the control command **setmqaut** grants and revokes authorities and is used to maintain the ACLs. For example, the command:

```
setmqaut -m JUPITER -t queue -n MOON.EUROPA -g VOYAGER +browse +get
```

allows the members of the group VOYAGER to browse messages on the queue MOON.EUROPA that is owned by the queue manager JUPITER. It allows the members to get messages from the queue as well. To revoke these authorities later, enter the following command:

```
setmqaut -m JUPITER -t queue -n MOON.EUROPA -g VOYAGER -browse -get
```

The command:

```
setmqaut -m JUPITER -t queue -n MOON.* -g VOYAGER +put
```

allows the members of the group VOYAGER to put messages on any queue with a name that commences with the characters MOON. . MOON.* is the name of a generic profile. A *generic profile* allows you to grant authorities for a set of objects using a single **setmqaut** command.

The control command **dspmqaut** is available to display the current authorities that a user or group has for a specified object. The control command **dmpmqaut** is also available to display the current authorities associated with generic profiles.

On IBM i, an administrator uses the CL command GRMQMAUT to grant authorities and the CL command RVKMQMAUT to revoke authorities. Generic profiles can be used as well. For example, the CL command:

```
GRMQMAUT MQMNAME(JUPITER) OBJTYPE(*Q) OBJ('MOON.*') USER(VOYAGER) AUT(*PUT)
```

provides the same function as the previous example of a **setmqaut** command; it allows the members of the group VOYAGER to put messages on any queue with a name that commences with the characters MOON.

The CL command DSPMQMAUT displays the current authorities that user or group has for a specified object. The CL commands WRKMQMAUT and WRKMQMAUTD are also available to work with the current authorities associated with objects and generic profiles.

If you do not want any authority checks, for example, in a test environment, you can disable the OAM.

Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [dspmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [dmpmqaut \(WebSphere MQ V7.1 Reference\)](#)

Using PCF to access OAM commands:

On IBM i, UNIX, Linux and Windows systems, you can use PCF commands to access OAM administration commands.

The PCF commands and their equivalent OAM commands are as follows:

Table 16. PCF commands and their equivalent OAM commands

PCF command	OAM command
Inquire Authority Records	dmpmqaut
Inquire Entity Authority	dspmqaut
Set Authority Record	setmqaut
Delete Authority Record	setmqaut with -remove option

The **setmqaut** and **dmpmqaut** commands are restricted to members of the mqm group. The equivalent PCF commands can be executed by users in any group who have been granted dsp and chg authorities on the queue manager.

For more information about using these commands, see “Introduction to Programmable Command Formats” on page 6.

Authority to work with WebSphere MQ objects on z/OS:

On z/OS, there are seven categories of authority check associated with calls to the MQI. You must define certain RACF profiles and give appropriate access to these profiles. Use the *RESLEVEL* profile to control how many users IDs are checked.

The seven categories of authority check associated with calls to the MQI:

Connection security

The authority checks that are performed when an application connects to a queue manager

Queue security

The authority checks that are performed when an application opens a queue or deletes a permanent dynamic queue

Process security

The authority checks that are performed when an application opens a process object

Namelist security

The authority checks that are performed when an application opens a namelist object

Alternate user security

The authority checks that are performed when an application requests alternate user authority when opening an object

Context security

The authority checks that are performed when an application opens a queue and specifies that it intends to set or pass the context information in the messages it puts on the queue

Topic security

The authority checks that are performed when an application opens a topic

Each category of authority check is implemented in the same way that command security and command resource security are implemented. You must define certain RACF profiles and give the necessary groups and user IDs access to these profiles at the required levels. For queue security, the level of access determines the types of operation the application can perform on a queue. For context security, the level of access determines whether the application can:

- Pass all the context fields
- Pass all the context fields and set the identity context fields
- Pass and set all the context fields

Each category of authority check can be turned on or off by defining switch profiles.

All the categories, except connection security, are known collectively as *API-resource security*.

By default, when an API-resource security check is performed as a result of an MQI call from an application using a batch connection, only one user ID is checked. When a check is performed as a result of an MQI call from a CICS or IMS application, or from the channel initiator, two user IDs are checked.

By defining a *RESLEVEL profile*, however, you can control whether zero, one, or two users IDs are checked. The number of user IDs that are checked is determined by the user ID associated with the type of connection when an application connects to the queue manager and the access level that user ID has to the RESLEVEL profile. The user ID associated with each type of connection is:

- The user ID of the connecting task for batch connections
- The CICS address space user ID for CICS connections
- The IMS region address space user ID for IMS connections
- The channel initiator address space user ID for channel initiator connections

For more information about the authority to work with WebSphere MQ objects on z/OS, see “Authority to administer WebSphere MQ on z/OS” on page 420.

Security for remote messaging

This section deals with remote messaging aspects of security.

You must provide users with authority to use the IBM WebSphere MQ facilities. This is organized according to actions to be taken with respect to objects and definitions. For example:

- Queue managers can be started and stopped by authorized users
- Applications must connect to the queue manager and have authority to use queues
- Message channels must be created and controlled by authorized users
- Objects are kept in libraries and access to these libraries can be restricted

The message channel agent at a remote site must check that the message being delivered originated from a user with authority to do so at this remote site. In addition, as MCAs can be started remotely, it might be necessary to verify that the remote processes trying to start your MCAs are authorized to do so. There are four possible ways for you to deal with this:

1. Make appropriate use of the PutAuthority attribute of your RCVR, RQSTR, or CLUSRCVR channel definition to control which user is used for authorization checks at the time incoming messages are put to your queues. See the DEFINE CHANNEL command description in the MQSC Command Reference.
2. Implement channel authentication records to reject unwanted connection attempts, or to set an MCAUSER value based on the following: the remote IP address, the remote user ID, the SSL or TLS Subject Distinguished Name (DN) provided, or the remote queue manager name.
3. Implement *user exit* security checking to ensure that the corresponding message channel is authorized. The security of the installation hosting the corresponding channel ensures that all users are properly authorized, so that you do not need to check individual messages.
4. Implement *user exit* message processing to ensure that individual messages are vetted for authorization.

Related concepts:

“Channel authentication records” on page 408

Security of IBM WebSphere MQ for IBM i objects:

This section deals with remote messaging aspects of security.

You must provide users with authority to make use of the IBM WebSphere MQ for IBM i facilities. This authority is organized according to actions to be taken with respect to objects and definitions. For example:

- Queue managers can be started and stopped by authorized users
- Applications need to connect to the queue manager, and have authority to make use of queues
- Message channels need to be created and controlled by authorized users

The message channel agent at a remote site must check that the message being delivered has derived from a user with authority to issue the message at this remote site. In addition, as MCAs can be started remotely, it might be necessary to verify that the remote processes trying to start your MCAs are authorized to do so. There are four possible ways for you to deal with this:

- Decree in the channel definition that messages must contain acceptable *context* authority, otherwise they are discarded.
- Implement channel authentication records to reject unwanted connection attempts, or to set an MCAUSER value based on one of the following: the remote IP address, the remote user ID, the SSL or TLS Distinguished Name (DN) provided, or the remote queue manager name.
- Implement user exit security checking to ensure that the corresponding message channel is authorized. The security of the installation hosting the corresponding channel ensures that all users are properly authorized, so that you do not need to check individual messages.
- Implement user exit message processing to ensure that individual messages are vetted for authorization.

Here are some facts about the way IBM WebSphere MQ for IBM i operates security:

- Users are identified and authenticated by IBM i.
- Queue manager services invoked by applications are run with the authority of the queue manager user profile, but in the user's process.
- Queue manager services invoked by user commands are run with the authority of the queue manager user profile.

Security of objects on UNIX and Linux systems:

Administration users must be part of the mqm group on your system (including root) if this ID is going to use WebSphere MQ administration commands.

You should always run amqcrsta as the “mqm” user ID.

User IDs on UNIX and Linux systems

The queue manager converts all uppercase or mixed case user identifiers into lowercase. The queue manager then inserts the user identifiers into the context part of a message, or checks their authorization. Authorizations are therefore based only on lowercase identifiers.

Security of objects on Windows systems:

Administration users must be part of both the mqm group and the administrators group on Windows systems if this ID is going to use WebSphere MQ administration commands.

User IDs on Windows systems

On Windows systems, *if there is no message exit installed*, the queue manager converts any uppercase or mixed case user identifiers into lowercase. The queue manager then inserts the user identifiers into the context part of a message, or checks their authorization. Authorizations are therefore based only on lowercase identifiers.

User IDs across systems:

Platforms other than Windows, UNIX and Linux systems use uppercase characters for user IDs in messages.

To allow Windows, UNIX and Linux systems to use lowercase user IDs in messages, the following conversions are carried out by the message channel agent (MCA) on these platforms:

At the sending end

The alphabetic characters in all user IDs are converted to uppercase characters, if there is no message exit installed.

At the receiving end

The alphabetic characters in all user IDs are converted to lowercase characters, if there is no message exit installed.

The automatic conversions are not carried out if you provide a message exit on UNIX, Linux and Windows systems for any other reason.

Using a custom authorization service

IBM WebSphere MQ supplies an installable authorization service. You can choose to install an alternative service.

The authorization service component supplied with IBM WebSphere MQ is called the Object Authority Manager (OAM). If the OAM does not supply the authorization facilities you need, you can write your own authorization service component. The installable service functions that must be implemented by an authorization service component are described at  Installable services interface reference information (*WebSphere MQ V7.1 Reference*).

Access control for clients

Access control is based on user IDs. There can be many user IDs to administer, and user IDs can be in different formats. You can set the server-connection channel property MCAUSER to a special user ID value for use by clients.

Access control in WebSphere MQ is based on user IDs. The user ID of the process making MQI calls is normally used. For MQ MQI clients, the server-connection MCA makes MQI calls on behalf of MQ MQI clients. You can select an alternative user ID for the server-connection MCA to use for making MQI calls. The alternative user ID can be associated either with the client workstation, or with anything you choose to organize and control the access of clients. The user ID needs to have the necessary authorities allocated to it on the server to issue MQI calls. Choosing an alternative user ID is preferable to allowing clients to make MQI calls with the authority of the server-connection MCA.

Table 17. The user ID used by a server-connection channel

User ID	When used
The user ID that is set by a security exit	Used unless blocked by a CHLAUTH TYPE(BLOCKUSER) rule. See the following section, "Setting the user ID in a security exit" for more information.
The user ID that is set by a CHLAUTH rule	Used unless over-ridden by a security exit. See "Channel authentication records" on page 408 for more information.
The user ID that is defined in the MCAUSER attribute in the SVRCONN channel definition	Used unless over-ridden by a security exit or a CHLAUTH rule.
The user ID that is flowed from the client machine	Used when no used ID is set by any other means.
The user ID that started the server-connection channel	Used when no user ID is set by any other means and no client user ID is flowed. See the following section, "The user ID that runs the channel program" on page 434 for more information.

Because the server-connection MCA makes MQI calls on behalf of remote users, it is important to consider the security implications of the server-connection MCA issuing MQI calls on behalf of remote clients and how to administer the access of a potentially large number of users.

- One approach is for the server-connection MCA to issue MQI calls on its own authority. But beware, it is normally undesirable for the server-connection MCA, with its powerful access capabilities, to issue MQI calls on behalf of client users.
- Another approach is to use the user ID that flows from the client. The server-connection MCA can issue MQI calls using the access capabilities of the client user ID. This approach presents a number of questions to consider:
 1. There are different formats for the user ID on different platforms. This sometimes causes problems if the format of the user ID on the client differs from the acceptable formats on the server.
 2. There are potentially many clients, with different, and changing user IDs. The IDs need to be defined and managed on the server.
 3. Is the user ID to be trusted? Any user ID can be flowed from a client, not necessarily the ID of the logged on user. For example, the client might flow an ID with full mqm authority that was intentionally only defined on the server for security reasons.
- The preferred approach is to define client identification tokens at the server, and so limit the capabilities of client connected applications. This is typically done by setting the server-connection channel property MCAUSER to a special user ID value to be used by clients, and defining few IDs for use by clients with different level of authorization on the server.

Setting the user ID in a security exit

For IBM WebSphere MQ MQI clients, the process that issues the MQI calls is the server-connection MCA. The user ID used by the server-connection MCA is contained in either the MCAUserIdentifier or LongMCAUserIdentifier fields of the MQCD. The contents of these fields are set by:

- Any values set by security exits
- The user ID from the client
- MCAUSER (in the server-connection channel definition)

The security exit can override the values that are visible to it, when it is invoked.

- If the server-connection channel MCAUSER attribute is set to nonblank, the MCAUSER value is used.
- If the server-connection channel MCAUSER attribute is blank, the user ID received from the client is used.

- If the server-connection channel MCAUSER attribute is blank, and no user ID is received from the client then the user ID that started the server-connection channel is used.

Ensure that the MCAUSER field is restricted to 12 characters on Windows platforms because any extra characters will be truncated which might lead to authorization failures.

The user ID that runs the channel program

When the user ID fields are derived from the user ID that started the server-connection channel, the following value is used:

- For z/OS, the user ID assigned to the channel initiator started task by the z/OS started procedures table.
- For TCP/IP (non-z/OS), the user ID from the `inetd.conf` entry, or the user ID that started the listener.
- For SNA (non-z/OS), the user ID from the SNA Server entry or (if there is none) the incoming attach request, or the user ID that started the listener.
- For NetBIOS or SPX, the user ID that started the listener.

If any server-connection channel definitions exist that have the MCAUSER attribute set to blank, clients can use this channel definition to connect to the queue manager with access authority determined by the user ID supplied by the client. This might be a security exposure if the system on which the queue manager is running allows unauthorized network connections. The IBM WebSphere MQ default server-connection channel (SYSTEM.DEF.SVRCONN) has the MCAUSER attribute set to blank. To prevent unauthorized access, update the MCAUSER attribute of the default definition with a user ID that has no access to IBM WebSphere MQ objects.

Case of user IDs

When you define a channel with `runmqsc`, the MCAUSER attribute is changed to uppercase unless the user ID is contained within single quotation marks.

For servers on UNIX, Linux and Windows systems, the content of the `MCAUserIdentifier` field that is received from the client is changed to lowercase.

For servers on IBM i, the content of the `LongMCAUserIdentifier` field that is received from the client is changed to uppercase.

For servers on UNIX and Linux systems, the content of the `LongMCAUserIdentifier` field that is received from the client is changed to lowercase.

By default, the user ID that is passed when a MQ JMS binding application is used, is the user ID for the JVM the application is running on.

It is also possible to pass a user ID via the `createQueueConnection` method.

Planning confidentiality

Plan how to keep your data confidential.

You can implement confidentiality at the application level or at link level. You might choose to use SSL or TLS, in which case you must plan your usage of digital certificates. You can also use channel exit programs if standard facilities do not satisfy your requirements.

Related concepts:

“Comparing link level security and application level security”

“Channel exit programs” on page 440

“Protecting channels with SSL” on page 449

Comparing link level security and application level security

This topic contains information about various aspects of link level security and application level security, and compares the two levels of security.

Link level and application level security are illustrated in Figure 76.

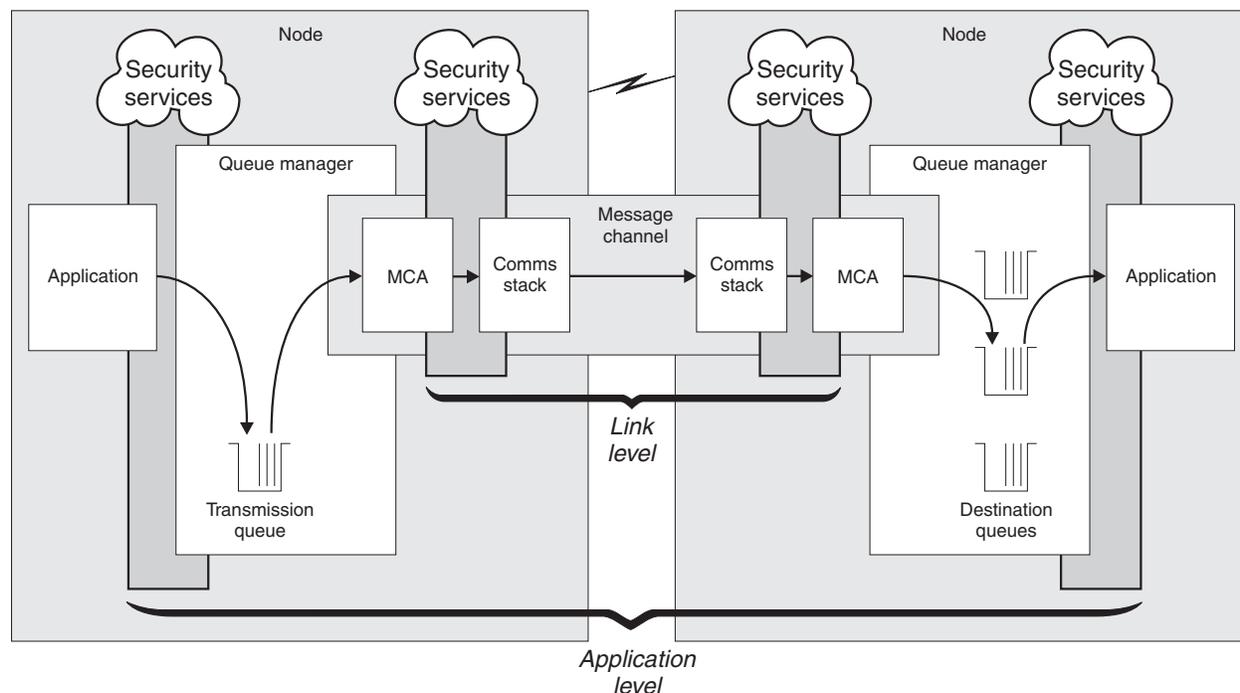


Figure 76. Link level security and application level security

Protecting messages in queues

Link level security can protect messages while they are transferred from one queue manager to another. It is particularly important when messages are transmitted over an insecure network. It cannot, however, protect messages while they are stored in queues at either a source queue manager, a destination queue manager, or an intermediate queue manager.

Application level security, by comparison, can protect messages while they are stored in queues and applies even when distributed queuing is not used. This is the major difference between link level security and application level security and is illustrated in Figure 76.

Queue managers not running in controlled and trusted environments

If a queue manager is running in a controlled and trusted environment, the access control mechanisms provided by WebSphere MQ might be considered sufficient to protect the messages stored on its queues. This is particularly true if only local queuing is involved and messages never leave the queue manager. Application level security in this case might be considered unnecessary.

Application level security might also be considered unnecessary if messages are transferred to another queue manager that is also running in a controlled and trusted environment, or are received from such a

queue manager. The need for application level security becomes greater when messages are transferred to, or received from, a queue manager that is not running in a controlled and trusted environment.

Differences in cost

Application level security might cost more than link level security in terms of administration and performance.

The cost of administration is likely to be greater because there are potentially more constraints to configure and maintain. For example, you might need to ensure that a particular user sends only certain types of message and sends messages only to certain destinations. Conversely, you might need to ensure that a particular user receives only certain types of message and receives messages only from certain sources. Instead of managing the link level security services on a single message channel, you might need to be configuring and maintaining rules for every pair of users who exchange messages across that channel.

There might be an effect on performance if security services are invoked every time an application puts or gets a message.

Organizations tend to consider link level security first because it might be easier to implement. They consider application level security if they discover that link level security does not satisfy all their requirements.

Availability of components

Generally, in a distributed environment, a security service requires a component on at least two systems. For example, a message might be encrypted on one system and decrypted on another. This applies to both link level security and application level security.

In a heterogeneous environment, with different platforms in use, each with different levels of security function, the required components of a security service might not be available for every platform on which they are needed and in a form that is easy to use. This is probably more of an issue for application level security than for link level security, particularly if you intend to provide your own application level security by buying in components from various sources.

Messages in a dead letter queue

If a message is protected by application level security, there might be a problem if, for any reason, the message does not reach its destination and is put on a dead letter queue. If you cannot work out how to process the message from the information in the message descriptor and the dead letter header, you might need to inspect the contents of the application data. You cannot do this if the application data is encrypted and only the intended recipient can decrypt it.

What application level security cannot do

Application level security is not a complete solution. Even if you implement application level security, you might still require some link level security services. For example:

- When a channel starts, the mutual authentication of the two MCAs might still be a requirement. This can be done only by a link level security service.
- Application level security cannot protect the transmission queue header, MQXQH, which includes the embedded message descriptor. Nor can it protect the data in WebSphere MQ channel protocol flows other than message data. Only link level security can provide this protection.

- If application level security services are invoked at the server end of an MQI channel, the services cannot protect the parameters of MQI calls that are sent over the channel. In particular, the application data in an MQPUT, MQPUT1, or MQGET call is unprotected. Only link level security can provide the protection in this case.

Related concepts:

- “Channel exit programs” on page 440
- “The SSPI channel exit program” on page 487
- “SNA LU 6.2 security services” on page 453
- “Channel authentication records” on page 408
- “WebSphere MQ Advanced Message Security” on page 439
- “Providing your own link level security”
- “Providing your own application level security” on page 440

Related reference:

- “WebSphere MQ support for SSL and TLS” on page 386

Link level security:

Link level security refers to those security services that are invoked, directly or indirectly, by an MCA, the communications subsystem, or a combination of the two working together.

Link level security is illustrated in Figure 76 on page 435.

Here are some examples of link level security services:

- The MCA at each end of a message channel can authenticate its partner. This is done when the channel starts and a communications connection has been established, but before any messages start to flow. If authentication fails at either end, the channel is closed and no messages are transferred. This is an example of an identification and authentication service.
- A message can be encrypted at the sending end of a channel and decrypted at the receiving end. This is an example of a confidentiality service.
- A message can be checked at the receiving end of a channel to determine whether its contents have been deliberately modified while it was being transmitted over the network. This is an example of a data integrity service.

Link level security provided by WebSphere MQ

The primary means of provision of confidentiality and data integrity in WebSphere MQ is by the use of SSL or TLS. For more information about the use of SSL and TLS in WebSphere MQ, see “WebSphere MQ support for SSL and TLS” on page 386. For authentication, WebSphere MQ provides the facility to use channel authentication records. Channel authentication records offer precise control over the access granted to connecting systems, at the level of individual channels or groups of channels. For more information, see “Channel authentication records” on page 408.

Providing your own link level security:

This collection of topics describes how you can provide your own link level security services. Writing your own channel exit programs is the main way to provide your own link level security services.

Channel exit programs are introduced in “Channel exit programs” on page 440. The same topic also describes the channel exit program that is supplied with IBM WebSphere MQ for Windows (the SSPI channel exit program). This channel exit program is supplied in source format so that you can modify the source code to suit your requirements. If this channel exit program, or channel exit programs available from other vendors, do not meet your requirements, you can design and write your own. This topic suggests ways in which channel exit programs can provide security services. For information about how

to write a channel exit program, see  Writing channel-exit programs (*WebSphere MQ V7.1 Programming Guide*).

Link level security using a security exit:

Security exits normally work in pairs; one at each end of a channel. They are called immediately after the initial data negotiation has completed on channel startup.

Security exits can be used to provide identification and authentication, access control, and confidentiality.

Link level security using a message exit:

A message exit can be used only on a message channel, not on an MQI channel. It has access to both the transmission queue header, MQXQH, which includes the embedded message descriptor, and the application data in a message. It can modify the contents of the message and change its length.

A message exit can be used for any purpose that requires access to the whole message rather than a portion of it.

Message exits can be used to provide identification and authentication, access control, confidentiality, data integrity, and non-repudiation, and for reasons other than security.

Link level security using send and receive exits:

Send and receive exits can be used on both message and MQI channels. They are called for all types of data that flow on a channel, and for flows in both directions.

Send and receive exits have access to each transmission segment. They can modify its contents and change its length.

On a message channel, if an MCA needs to split a message and send it in more than one transmission segment, a send exit is called for each transmission segment containing a portion of the message and, at the receiving end, a receive exit is called for each transmission segment. The same occurs on an MQI channel if the input or output parameters of an MQI call are too large to be sent in a single transmission segment.

On an MQI channel, byte 10 of a transmission segment identifies the MQI call, and indicates whether the transmission segment contains the input or output parameters of the call. Send and receive exits can examine this byte to determine whether the MQI call contains application data that might need to be protected.

When a send exit is called for the first time, to acquire and initialize any resources it needs, it can ask the MCA to reserve a specified amount of space in the buffer that holds a transmission segment. When it is called later to process a transmission segment, it can use this space to add an encrypted key or a digital signature, for example. The corresponding receive exit at the other end of the channel can remove the data added by the send exit, and use it to process the transmission segment.

Send and receive exits are best suited for purposes in which they do not need to understand the structure of the data they are handling and can therefore treat each transmission segment as a binary object.

Send and receive exits can be used to provide confidentiality and data integrity, and for uses other than security.

Related reference:



Identifying the API call in a send or receive exit program (*WebSphere MQ V7.1 Installing Guide*)

Application level security:

Application level security refers to those security services that are invoked at the interface between an application and a queue manager to which it is connected.

These services are invoked when the application issues MQI calls to the queue manager. The services might be invoked, directly or indirectly, by the application, the queue manager, another product that supports WebSphere MQ, or a combination of any of these working together. Application level security is illustrated in Figure 76 on page 435.

Application level security is also known as *end-to-end security* or *message level security*.

Here are some examples of application level security services:

- When an application puts a message on a queue, the message descriptor contains a user ID associated with the application. However, there is no data present, such as an encrypted password, that can be used to authenticate the user ID. A security service can add this data. When the message is eventually retrieved by the receiving application, another component of the service can authenticate the user ID using the data that has travelled with the message. This is an example of an identification and authentication service.
- A message can be encrypted when it is put on a queue by an application and decrypted when it is retrieved by the receiving application. This is an example of a confidentiality service.
- A message can be checked when it is retrieved by the receiving application. This check determines whether its contents have been deliberately modified since it was first put on a queue by the sending application. This is an example of a data integrity service.

WebSphere MQ Advanced Message Security:

WebSphere MQ Advanced Message Security extends the security features available in WebSphere MQ.

If you are moving highly sensitive or valuable information, especially confidential or payment-related information such as patient records or credit card details, you must pay special attention to information security. Ensuring that information moving around the enterprise retains its integrity and is protected from unauthorized access is an ongoing challenge and responsibility. You are also likely to be required to comply with security regulations, at the risk of penalties for non-compliance.

You can develop your own security extensions to WebSphere MQ. However, such solutions require specialist skills and can be complicated and expensive to maintain. WebSphere MQ Advanced Message Security helps address these challenges when moving information around the enterprise between virtually every type of commercial IT system.

WebSphere MQ Advanced Message Security extends the security features of WebSphere MQ in the following ways:

- It provides application-level, end-to-end data protection for your point to point messaging infrastructure, using either encryption or digital signing of messages.
- It provides comprehensive security without writing complex security code or modifying or recompiling existing applications.
- It uses Public Key Infrastructure (PKI) technology to provide authentication, authorization, confidentiality, and data integrity services for messages.
- It provides administration of security policies for mainframe and distributed servers.
- It supports both WebSphere MQ servers and clients.

- It integrates with WebSphere MQ File Transfer Edition to provide an end-to-end secure messaging solution.

For more information see the WebSphere MQ Advanced Message Security webpage at

➔ <http://www-01.ibm.com/software/integration/wmq/advanced-message-security/> and the

WebSphere MQ Advanced Message Security product documentation at ➔ http://www-01.ibm.com/support/knowledgecenter/SSKSWQ_7.0.1/welcomePage/AMSDisWelcomePage.htm.

Providing your own application level security:

This collection of topics describes how you can provide your own application level security services.

To help you implement application level security, WebSphere MQ provides two exits, the API exit and the API-crossing exit.

These exits can provide identification and authentication, access control, confidentiality, data integrity, and non-repudiation services, and other functions not related to security.

If the API exit or API-crossing exit is not supported in your system environment, you might want to consider other ways of providing your own application level security. One way is to develop a higher level API that encapsulates the MQI. Programmers then use this API, instead of the MQI, to write WebSphere MQ applications.

The most common reasons for using a higher level API are:

- To hide the more advanced features of the MQI from programmers.
- To enforce standards in the use of the MQI.
- To add function to the MQI. This additional function can be security services.

Some vendor products use this technique to provide application level security for WebSphere MQ.

If you are planning to provide security services in this way, note the following regarding data conversion:

- If a security token, such as a digital signature, has been added to the application data in a message, any code performing data conversion must be aware of the presence of this token.
- A security token might have been derived from a binary image of the application data. Therefore, any checking of the token must be done before converting the data.
- If the application data in a message has been encrypted, it must be decrypted before data conversion.

Channel exit programs

Channel exit programs are programs that are called at defined places in the processing sequence of an MCA. Users and vendors can write their own channel exit programs. Some are supplied by IBM.

There are several types of channel exit program, but only four have a role in providing link level security:

- Security exit
- Message exit
- Send exit
- Receive exit

These four types of channel exit program are illustrated in Figure 77 on page 441 and are described in the following topics.

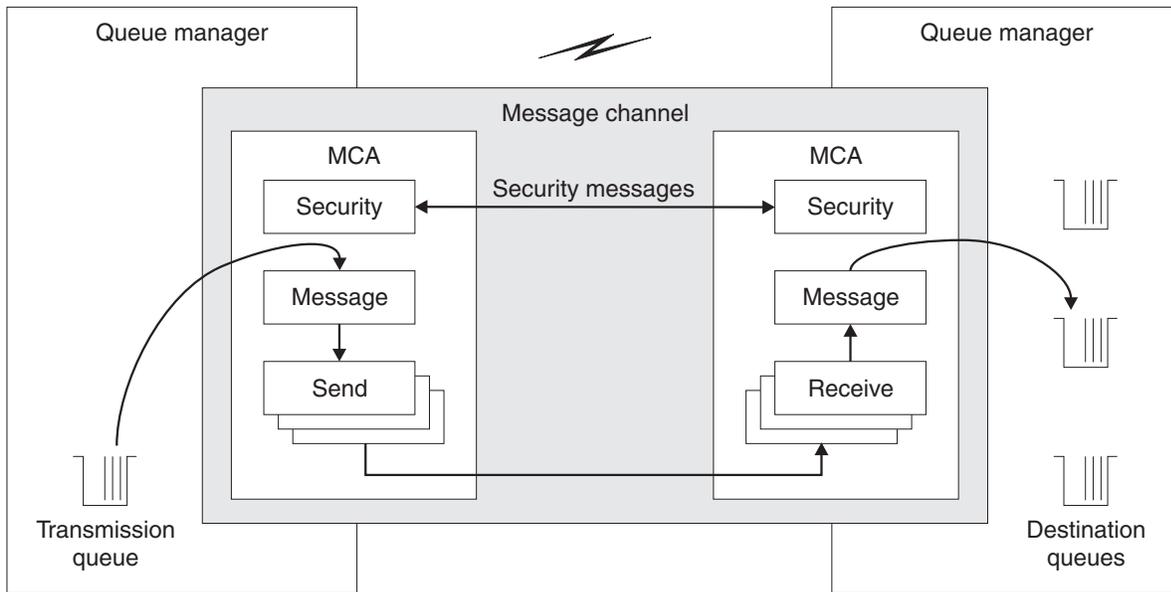


Figure 77. Security, message, send, and receive exits on a message channel

Related concepts:

 Channel-exit programs for messaging channels (*WebSphere MQ V7.1 Programming Guide*)

Security exit overview:

Security exits normally work in pairs. They are called before messages flow and their purpose is to allow an MCA to authenticate its partner.

Security exits normally work in pairs; one at each end of a channel. They are called immediately after the initial data negotiation has completed on channel startup, but before any messages start to flow. The primary purpose of the security exit is to enable the MCA at each end of a channel to authenticate its partner. However, there is nothing to prevent a security exit from performing other function, even function that has nothing to do with security.

Security exits can communicate with each other by sending *security messages*. The format of a security message is not defined and is determined by the user. One possible outcome of the exchange of security messages is that one of the security exits might decide not to proceed any further. In that case, the channel is closed and messages do not flow. If there is a security exit at only one end of a channel, the exit is still called and can elect whether to continue or to close the channel.

Security exits can be called on both message and MQI channels. The name of a security exit is specified as a parameter in the channel definition at each end of a channel.

For more information about security exits, see “Link level security using a security exit” on page 438.

Message exit:

Message exits operate only on message channels and normally work in pairs. A message exit can operate on the whole message and make various changes to it.

Message exits at the sending and receiving ends of a channel normally work in pairs. A message exit at the sending end of a channel is called after the MCA has got a message from the transmission queue. At the receiving end of a channel, a message exit is called before the MCA puts a message on its destination queue.

A message exit has access to both the transmission queue header, MQXQH, which includes the embedded message descriptor, and the application data in a message. A message exit can modify the contents of the message and change its length. A change of length might be the result of compressing, decompressing, encrypting, or decrypting the message. It might also be the result of adding data to the message, or removing data from it.

Message exits can be used for any purpose that requires access to the whole message, rather than a portion of it, and not necessarily for security.

A message exit can determine that the message it is currently processing should not proceed any further towards its destination. The MCA then puts the message on the dead letter queue. A message exit can also close the channel.

Message exits can be called only on message channels, not on MQI channels. This is because the purpose of an MQI channel is to enable the input and output parameters of MQI calls to flow between the WebSphere MQ MQI client application and the queue manager.

The name of a message exit is specified as a parameter in the channel definition at each end of a channel. You can also specify a list of message exits to be run in succession.

For more information about message exits, see “Link level security using a message exit” on page 438.

Send and receive exits:

Send and receive exits typically work in pairs. They operate on transmission segments and are best used where the structure of the data they are processing is not relevant.

A *send exit* at one end of a channel and a *receive exit* at the other end normally work in pairs. A send exit is called just before an MCA issues a communications send to send data over a communications connection. A receive exit is called just after an MCA has regained control following a communications receive and has received data from a communications connection. If sharing conversations is in use, over an MQI channel, a different instance of a send and receive exit is called for each conversation.

The WebSphere MQ channel protocol flows between two MCAs on a message channel contain control information as well as message data. Similarly, on an MQI channel, the flows contain control information as well as the parameters of MQI calls. Send and receive exits are called for all types of data.

Message data flows in only one direction on a message channel but, on an MQI channel, the input parameters of an MQI call flow in one direction and the output parameters flow in the other. On both message and MQI channels, control information flows in both directions. As a result, send and receive exits can be called at both ends of a channel.

The unit of data that is transmitted in a single flow between two MCAs is called a *transmission segment*. Send and receive exits have access to each transmission segment. They can modify its contents and change its length. A send exit, however, must not change the first 8 bytes of a transmission segment. These 8 bytes form part of the WebSphere MQ channel protocol header. There are also restrictions on

how much a send exit can increase the length of a transmission segment. In particular, a send exit cannot increase its length beyond the maximum that was negotiated between the two MCAs at channel startup.

On a message channel, if a message is too large to be sent in a single transmission segment, the sending MCA splits the message and sends it in more than one transmission segment. As a consequence, a send exit is called for each transmission segment containing a portion of the message and, at the receiving end, a receive exit is called for each transmission segment. The receiving MCA reconstitutes the message from the transmission segments after they have been processed by the receive exit.

Similarly, on an MQI channel, the input or output parameters of an MQI call are sent in more than one transmission segment if they are too large. This might occur, for example, on an MQPUT, MQPUT1, or MQGET call if the application data is sufficiently large.

Taking these considerations into account, it is more appropriate to use send and receive exits for purposes in which they do not need to understand the structure of the data they are handling and can therefore treat each transmission segment as a binary object.

A send or a receive exit can close a channel.

The names of a send exit and a receive exit are specified as parameters in the channel definition at each end of a channel. You can also specify a list of send exits to be run in succession. Similarly, you can specify a list of receive exits.

For more information about send and receive exits, see “Link level security using send and receive exits” on page 438.

Planning data integrity

Plan how to preserve the integrity of your data.

You can implement data integrity at the application level or at link level.

At the application level, you might choose to use WebSphere MQ Advanced Message Security (AMS) to digitally sign messages in order to protect against unauthorized modification. You can also use API exit programs if standard facilities do not satisfy your requirements.

At the link level, you might choose to use SSL or TLS, in which case you must plan your usage of digital certificates. You can also use channel exit programs if standard facilities do not satisfy your requirements.

Related concepts:

“Protecting channels with SSL” on page 449

“Data integrity in IBM WebSphere MQ” on page 385

“WebSphere MQ Advanced Message Security” on page 439

Related reference:



API exit reference (*WebSphere MQ V7.1 Reference*)



Channel-exit calls and data structures (*WebSphere MQ V7.1 Reference*)

Planning auditing

Decide what data you need to audit, and how you will capture and process audit information. Consider how to check that your system is correctly configured.

There are several aspects to activity monitoring. The aspects you must consider are often defined by auditor requirements, and these requirements are often driven by regulatory standards such as HIPAA

(Health Insurance Portability and Accountability Act) or SOX (Sarbanes–Oxley). IBM WebSphere MQ provides features intended to help with compliance to such standards.

Consider whether you are interested only in exceptions or whether you are interested in all system behavior.

Some aspects of auditing can also be considered as operational monitoring; one distinction for auditing is that you are often looking at historic data, not just looking at real-time alerts. Monitoring is covered in the section “Monitoring and performance” on page 829.

What data to audit

Consider what types of data or activity you need to audit, as described in the following sections:

Changes made to IBM WebSphere MQ using the IBM WebSphere MQ interfaces

Configure IBM WebSphere MQ to issue instrumentation events, specifically command events and configuration events.

Changes made to IBM WebSphere MQ outside its control

Some changes can affect how IBM WebSphere MQ behaves, but cannot be directly monitored by IBM WebSphere MQ. Examples of such changes include changes to the configuration files `mqs.ini`, `qm.ini`, and `mqclient.ini`, the creation and deletion of queue managers, installation of binary files such as user exit programs, and changes to file permissions. To monitor these activities, you must use tools running at the level of the operating system. Different tools are available and appropriate for different operating systems. You might also have logs created by associated tools such as *sudo*.

Operational control of IBM WebSphere MQ

You might have to use operating system tools to audit activities such as the starting and stopping of queue managers. In some cases, IBM WebSphere MQ can be configured to issue instrumentation events.

Application activity within IBM WebSphere MQ

To audit the actions of applications, for example opening of queues and putting and getting of messages, configure IBM WebSphere MQ to issue appropriate events.

Intruder alerts

To audit attempted breaches of security, configure your system to issue authorization events. Channel events might also be useful to show activity, particularly if a channel ends unexpectedly.

Planning the capture, display, and archiving of audit data

Many of the elements you need are reported as WebSphere MQ event messages. You must choose tools that can read and format these messages. If you are interested in long-term storage and analysis you must move them to an auxiliary storage mechanism such as a database. If you do not process these messages, they remain on the event queue, possibly filling the queue. You might decide to implement a tool that automatically takes action based on some events; for example, to issue an alert when a security failure happens.

Verifying that your system is correctly configured

A set of tests are supplied with the IBM WebSphere MQ Explorer. Use these to check your object definitions for problems.

Also, check periodically that the system configuration is as you expect. Although command and configuration events can report when something is changed, it is also useful to dump the configuration and compare it to a known good copy.

Common Criteria environmental considerations

In order that IBM WebSphere MQ operates in accordance with its Common Criteria certificate, the environmental requirements defined in this topic need to be met.

- There must be one or more competent individuals that are assigned to manage IBM WebSphere MQ and the security of the information that it contains. Such personnel are assumed not to be careless, willfully negligent, or hostile.
- You must verify that your installation image for IBM WebSphere MQ has not been tampered with during the delivery process:
 - Do not use the HTTP download method to download IBM WebSphere MQ, because the HTTP download method does not provide protection against data modification in transit. Use the IBM Download Director instead.
 - If you download IBM WebSphere MQ using the IBM Download Director, the Download Director computes a cryptographic checksum to verify that each file was correctly downloaded.
 - If you obtain IBM WebSphere MQ using a CD, then you should check the CD labels and packing slip against the invoice to check that the order numbers and dates match. You can also contact IBM to obtain a tracking number for your delivery and confirm that the tracking number on the package received is correct.
- The IBM Global Security Kit (GSKit) component is required for TLS support. On Windows, GSKit is always installed if the Server or Client components are selected at installation time. However, on the AIX, HP-UX, Linux, and Solaris platforms, GSKit is an optional component and you must select this option during installation; see  IBM WebSphere MQ components for UNIX and Linux (*WebSphere MQ V7.1 Installing Guide*) for more information.
- The operating system must be configured in accordance with the installation guides provided by the manufacturer and where applicable, in its evaluated configuration. It must be securely configured such that the operating system protects IBM WebSphere MQ from any unauthorized users or processes.
- It is the responsibility of the system administrator to ensure that any connections over the internet are secured by the appropriate protocol (for example, HTTPS or SFTP). The system administrator must also ensure that any external URLs are trusted sites.
- The operating system environment must be configured to allow local login for a single user only.
- Remote logins to the operating system must be disabled. Telnet, rlogin, SSH, and any similar remote login services are included in this restriction.
- The system must be kept in a physically secure environment with appropriate access controls to prevent unauthorized physical access.
- The system administrator must set the system time to the appropriate date and time of the server.
- The operating system environment must provide procedures and mechanisms to ensure that recovery after a system failure, or other discontinuity, does not result in any compromise to security.
- The following operating systems are supported:
 - AIX (V6.1)
 - AIX (V7.1)
 - HP-UX IA64 (V11.31)
 - Microsoft Windows 7 Enterprise (32-bit and 64-bit versions)
 - Microsoft Windows 7 Professional (32-bit and 64-bit versions)
 - Microsoft Windows 7 Ultimate (32-bit and 64-bit versions)
 - Microsoft Windows Server 2008 Enterprise Edition (32-bit and 64-bit versions)
 - Microsoft Windows Server 2008 R2 Enterprise Edition
 - Microsoft Windows Server 2008 R2 Standard Edition
 - Microsoft Windows Server 2008 Standard Edition (32-bit and 64-bit versions)
 - Microsoft Windows Vista Business (32-bit and 64-bit versions)

- Microsoft Windows Vista Enterprise (32-bit and 64-bit versions)
- Microsoft Windows Vista Ultimate (32-bit and 64-bit versions)
- Microsoft Windows Server 2003, Standard Edition (32-bit)
- Microsoft Windows Server 2003, Standard x64 Edition
- Microsoft Windows Server 2003, R2 Standard Edition (32-bit)
- Microsoft Windows Server 2003, R2 Standard x64 Edition
- Microsoft Windows Server 2003, Enterprise Edition (32-bit and 64-bit versions)
- Microsoft Windows Server 2003, Enterprise x64 Edition
- Microsoft Windows Server 2003, R2 Enterprise Edition (32-bit and 64-bit versions)
- Microsoft Windows Server 2003, R2 Enterprise x64 Edition
- Microsoft Windows XP Professional
- Microsoft Windows XP Professional x64 Edition
- Linux for System p: Red Hat Enterprise Linux (RHEL) V5.0
- Linux for System p: Red Hat Enterprise Linux (RHEL) V6.0
- Linux for System p: SUSE Linux Enterprise Server V10
- Linux for System p: SUSE Linux Enterprise Server V11
- Linux for System x (32 bit and 64 bit): Red Hat Enterprise Linux (RHEL) V5.0
- Linux for System x (32 bit and 64 bit): Red Hat Enterprise Linux (RHEL) V6.0
- Linux for System x (32 bit and 64 bit): SUSE Linux Enterprise Server V10
- Linux for System x (32 bit and 64 bit): SUSE Linux Enterprise Server V11
- Linux for System z[®]: Red Hat Enterprise Linux (RHEL) V5.0
- Linux for System z: Red Hat Enterprise Linux (RHEL) V6.0
- Linux for System z: SUSE Linux Enterprise Server V10
- Linux for System z: SUSE Linux Enterprise Server V11
- Oracle Solaris SPARC: Oracle Solaris V10
- Oracle Solaris x86_64: Oracle Solaris V10

IBM WebSphere MQ relies on the operating system to provide user and group IDs and time and date information. In addition, you need an application to read the event logs so that the audit records produced by IBM WebSphere MQ can be read.

Ensure that only IBM WebSphere MQ administrators have the authority to access the event queues. Do not grant any authority to any other users for the event queues.

The evaluation of IBM WebSphere MQ does not include the following aspects:

- The operating system
- AMS security policies
- Java and JMS client applications
- MQTT client and server applications
- Remote administration
- IBM WebSphere MQ Explorer
- Windows Default Configuration application
- Third-party or user-written authorization services not supplied with the IBM WebSphere MQ product.
- Publish/Subscribe configurations involving more than one queue manager.

Related concepts:

“Common Criteria in WebSphere MQ” on page 384

“Configuring IBM WebSphere MQ for Common Criteria” on page 685

Planning security by topology

This collection of topics covers security in specific situations, namely for channels, queue manager clusters, publish/subscribe and multicast applications, and when using a firewall.

Channel security

When you send or receive a message through a channel, you need a user ID that has access to various WebSphere MQ resources.

To receive messages at PUT time for MCAs, you can use either the user ID associated with the MCA, or the user ID associated with the message.

At CONNECT time you can map the asserted user ID to an alternative user, by using **CHLAUTH** channel authentication records.

In WebSphere MQ, channels can be protected by SSL or TLS support.

The user IDs associated with sending and receiving channels, excluding the sender channel where the MCAUSER attribute is unused, require access to the following resources:

- The user ID associated with a sending channel requires access to the queue manager, the transmission queue, the dead-letter queue, and access to any other resources that are required by channel exits.
- The MCAUSER user ID of a receiver channel needs *+setall* authority.

The reason is that the receiver channel has to create the full MQMD, including all context fields, using the data it received from the remote sender channel.

The queue manager therefore requires that the user performing this activity has the *+setall* authority. This *+setall* authority must be granted to the user for:

- All queues that the receiver channel validly puts messages to.
- The queue manager object. See  *Authorizations for context (WebSphere MQ V7.1 Reference)* for further information.
- With the user ID associated with the receiving channel you can open the target queues to put messages onto the queues.

This involves the Message queuing Interface (MQI), so additional access control checks might need to be made if you are not using the WebSphere MQ Object Authority Manager (OAM). You can specify whether the authorization checks are made against the user ID associated with the MCA (as described in this topic), or against the user ID associated with the message (from the MQMD  *UserIdentifier (WebSphere MQ V7.1 Reference)* field).

For the channel types to which it applies, the **PUTAUT** parameter of a channel definition specifies which user ID is used for these checks.

- The channel defaults to using the queue manager's service account, that will have full administrative rights and requires no special authorizations

In the case of server-connection channels, administrative connections are blocked by default by CHLAUTH rules and require explicit provisioning.

Channels of type receiver, requester, and cluster-receiver allow local administration by any adjacent queue manager, unless the administrator takes steps to restrict this access.

- If you use a user ID that lacks WebSphere administrative privileges, then you must grant dsp and ctrlx authority for the channel to that user ID for the channel to work. The MCAUSER attribute is unused for the SDR channel type.

- If you use the user ID associated with the message, it is likely that the user ID is from a remote system.

This remote system user ID must be recognized by the target system. For example, issue the following commands:

```
-
setmqaut -m [QMGR] -t qmgr -g [GROUP] +connect +inq +setall
-
setmqaut -m[QMGR] -t chl -n [CHANNEL] -g [Group] +dsp +ctrlx
-
setmqaut -m [QMGR] -t q -n [DLQ if set] -g [Group] +put +setall
-
setmqaut -m [QMGR] -t q -n[AUTHORIZED QUEUES] -g [Group] +put +setall
```

Attention: Exercise caution when authorizing a user ID to place messages onto the Command Queue or other sensitive system queues.

The user ID associated with the MCA depends on the type of MCA. There are two types of MCA:

Caller MCA

MCAs that initiate a channel. Caller MCAs can be started as individual processes, as threads of the channel initiator, or as threads of a process pool. The user ID used is the user ID associated with the parent process (the channel initiator), or the user ID associated with the process that starts the MCA.

Responder MCA

Responder MCAs are MCAs that are started as a result of a request by a caller MCA. Responder MCAs can be started as individual processes, as threads of the listener, or as threads of a process pool. The user ID can be any one of the following types (in this order of preference):

1. On APPC, the caller MCA can indicate the user ID to be used for the responder MCA. This is called the network user ID and applies only to channels started as individual processes. Set the network user ID by using the **USERID** parameter of the channel definition.
2. If the **USERID** parameter is not used, the channel definition of the responder MCA can specify the user ID that the MCA must use. Set the user ID by using the **MCAUSER** parameter of the channel definition.
3. If the user ID has not been set by either of the previous (two) methods, the user ID of the process that starts the MCA or the user ID of the parent process (the listener) is used.

Related concepts:

“Channel authentication records” on page 408

Protecting channel initiator definitions:

Only members of the mqm group can manipulate channel initiators.

WebSphere MQ channel initiators are not WebSphere MQ objects; access to them is not controlled by the OAM. WebSphere MQ does not allow users or applications to manipulate these objects, unless their user ID is a member of the mqm group. If you have an application that issues the PCF command **StartChannelInitiator**, the user ID specified in the message descriptor of the PCF message must be a member of the mqm group on the target queue manager.

A user ID must also be a member of the mqm group on the target machine to issue the equivalent MQSC commands through the Escape PCF command or using **runmqsc** in indirect mode.

Transmission queues:

Queue managers automatically put remote messages on a transmission queue; no special authority is required for this.

However, if you need to put a message directly on a transmission queue, this requires special authorization; see Table 20 on page 473.

Channel exits:

If channel authentication records are not suitable, you can use channel exits for added security. A security exit forms a secure connection between two security exit programs. One program is for the sending message channel agent (MCA), and one is for the receiving MCA.

See “Channel exit programs” on page 440 for more information about channel exits.

Protecting channels with SSL:

SSL support in WebSphere MQ uses the queue manager authentication information object, and various MQSC commands. You must also consider your use of digital certificates.

Commands and attributes for SSL support

The Secure Sockets Layer (SSL) protocol provides channel security, with protection against eavesdropping, tampering, and impersonation. WebSphere MQ support for SSL enables you to specify, on the channel definition, that a particular channel uses SSL security. You can also specify details of the type of security you want, such as the encryption algorithm you want to use.

The following MQSC commands support SSL:

ALTER AUTHINFO

Modifies the attributes of an authentication information object.

DEFINE AUTHINFO

Creates an authentication information object.

DELETE AUTHINFO

Deletes an authentication information object.

DISPLAY AUTHINFO

Displays the attributes for a specific authentication information object.

The following queue manager parameters support SSL:

SSLCRLNL

The SSLCRLNL attribute specifies a namelist of authentication information objects which are used to provide certificate revocation locations to allow enhanced TLS/SSL certificate checking.

SSLCRYP

On Windows, UNIX and Linux systems, sets the SSLCryptoHardware queue manager attribute. This attribute is the name of the parameter string that you can use to configure the cryptographic hardware you have on your system.

SSLEV

Determines whether an SSL event message is reported if a channel using SSL fails to establish an SSL connection.

SSLFIPS

Specifies whether only FIPS-certified algorithms are to be used if cryptography is carried out in WebSphere MQ, rather than in cryptographic hardware. If cryptographic hardware is configured,

the cryptographic modules provided by the hardware product are used, and these might be FIPS-certified to a particular level. This depends on the hardware product in use.

SSLKEYR

On Windows, UNIX and Linux systems, associates a key repository with a queue manager. The key database is held in a *GSKit* key database. (The IBM Global Security Kit (GSKit) enables you to use SSL security on Windows, UNIX and Linux systems.)

SSLRKEYC

The number of bytes to be sent and received within an SSL conversation before the secret key is renegotiated. The number of bytes includes control information sent by the MCA.

The following channel parameters support SSL:

SSLCAUTH

Defines whether WebSphere MQ requires and validates a certificate from the SSL client.

SSLCIPH

Specifies the encryption strength and function (CipherSpec), for example NULL_MD5 or RC4_MD5_US. The CipherSpec must match at both ends of channel.

SSLPEER

Specifies the distinguished name (unique identifier) of allowed partners.

This section describes the **setmqaut**, **dspmqaut**, **dmpmqaut**, **rcrmqobj**, **rcdmqimg**, and **dspmqfls** commands to support the authentication information object. It also describes the **iKeycmd** command for managing certificates on UNIX and Linux systems, and the **runmqakm** tool for managing certificates on UNIX, Linux and Windows systems. See the following sections:

-  **setmqaut** (*WebSphere MQ V7.1 Reference*)
-  **dspmqaut** (*WebSphere MQ V7.1 Reference*)
-  **dmpmqaut** (*WebSphere MQ V7.1 Reference*)
-  **rcrmqobj** (*WebSphere MQ V7.1 Reference*)
-  **rcdmqimg** (*WebSphere MQ V7.1 Reference*)
-  **dspmqfls** (*WebSphere MQ V7.1 Reference*)
-  **Managing keys and certificates** (*WebSphere MQ V7.1 Reference*)

For an overview of channel security using SSL, see “WebSphere MQ support for SSL and TLS” on page 386.

For details of MQSC commands associated with SSL, see  **ALTER AUTHINFO** (*WebSphere MQ V7.1 Reference*),  **DEFINE AUTHINFO** (*WebSphere MQ V7.1 Reference*),  **DELETE AUTHINFO** (*WebSphere MQ V7.1 Reference*), and  **DISPLAY AUTHINFO** (*WebSphere MQ V7.1 Reference*).

For details of PCF commands associated with SSL, see  **Change, Copy, and Create Authentication Information Object** (*WebSphere MQ V7.1 Reference*),  **Delete Authentication Information Object** (*WebSphere MQ V7.1 Reference*), and  **Inquire Authentication Information Object** (*WebSphere MQ V7.1 Reference*).

Self-signed and CA-signed certificates

It is important to plan your use of digital certificates, both when you are developing and testing your application, and for its use in production. You can use CA-signed certificates or self-signed certificates, depending on the usage of your queue managers and client applications.

CA-signed certificates

For production systems, obtain your certificates from a trusted certificate authority (CA). When you obtain a certificate from an external CA, you pay for the service.

Self-signed certificates

While you are developing your application you can use self-signed certificates or certificates issued by a local CA, depending on platform:

On UNIX, Linux, and Windows systems, you can use self-signed certificates. See “Creating a self-signed personal certificate on UNIX, Linux, and Windows systems” on page 639 for instructions.

On IBM i systems, you can use certificates signed by the local CA. See “Requesting a server certificate on IBM i” on page 624 for instructions.

On z/OS, you can use either self-signed or local CA-signed certificates. See “Creating a self-signed personal certificate on z/OS” on page 677 or “Requesting a personal certificate on z/OS” on page 678 for instructions.

Self-signed certificates are not suitable for production use, for the following reasons:

- Self-signed certificates cannot be revoked, which might allow an attacker to spoof an identity after a private key has been compromised. CAs can revoke a compromised certificate, which prevents its further use. CA-signed certificates are therefore safer to use in a production environment, though self-signed certificates are more convenient for a test system.
- Self-signed certificates never expire. This is both convenient and safe in a test environment, but in a production environment it leaves them open to eventual security breaches. The risk is compounded by the fact that self-signed certificates cannot be revoked.
- A self-signed certificate is used both as a personal certificate and as a root (or trust anchor) CA certificate. A user with a self-signed personal certificate might be able to use it to sign other personal certificates. In general, this is not true of personal certificates issued by a CA, and represents a significant exposure.

CipherSpecs and digital certificates

Only a subset of the supported CipherSpecs can be used with all of the supported types of digital certificate. It is therefore necessary to choose an appropriate CipherSpec for your digital certificate. Similarly, if your organization's security policy requires that a particular CipherSpec be used, then you must obtain a suitable digital certificate.

For more information on the relationship between CipherSpecs and digital certificates, refer to “Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403

Certificate validation policies

The IETF RFC 5280 standard specifies a series of certificate validation rules which compliant application software must implement in order to prevent impersonation attacks. A set of certificate validation rules is known as a certificate validation policy. For more information about certificate validation policies in WebSphere MQ, see “Certificate validation policies in WebSphere MQ” on page 402.

z/OS MQ SVRCONN channel:

The z/OS MQ SVRCONN channel is not secure without implementing channel authentication, or adding a security exit and implementing SSL. SVRCONN channels do not have a security exit defined by default.

Security concerns

SVRCONN channels are not secure as initially defined, SYSTEM.DEF.SVRCONN for example. To secure a SVRCONN channel you must set up channel authentication using the  SET CHLAUTH (*WebSphere MQ V7.1 Reference*) command, or install a security exit and implement SSL, or TLS.

You must use a publicly available sample security exit, write a security exit yourself, or purchase a security exit.

There are several samples available that you can use as a good starting point for writing your own SVRCONN channel security exit.

In WebSphere MQ for z/OS, the member CSQ4BCX3 in your hlq.SCSQC37S library is a security exit sample written in the C language. Sample CSQ4BCX3 is also shipped pre-compiled in your hlq.SCSQAUTH library.

You can implement the CSQ4BCX3 sample exit by copying the compiled member hlq.SCSQAUTH(CSQ4BCX3) into a load library that is allocated to the CSQXLIB DD in your CHIN Proc. Note that the CHIN requires the load library to be set as "Program Controlled".

Alter your SVRCONN channel to set CSQ4BCX3 as the security exit.

When a client connects using that SVRCONN channel, CSQ4BCX3 will authenticate using the RemoteUserIdentifier and RemotePassword pair from MQCD. If authentication is successful it will copy RemoteUserIdentifier into MCAUserIdentifier, changing the identity context of the thread.

If you are writing an MQ Java client you can use pop-ups to query the user and set MQEnvironment.userID and MQEnvironment.password. These values will be passed when the connection is made.

Now that you have a functional security exit, there is the additional concern that the userid and password are being transmitted in plain text across the network when the connection is made, as are the contents of any subsequent MQ messages. You can use SSL to encrypt this initial connection information as well as the contents of any MQ messages.

Example

To secure the MQ Explorer SVRCONN channel SYSTEM.ADMIN.SVRCONN complete the following steps:

- Copy hlq.SCSQAUTH(CSQ4BCX3) into a load library that is allocated to the CSQXLIB DD in the CHINIT Proc
- Verify that load library is Program Controlled
- Alter the SYSTEM ADMIN.SVRCONN to use security exit CSQ4BCX3
- On MQ Explorer V7
 - Right click the z/OS Queue Manager name
 - Connection Details
 - Properties

- Userid - Enter your z/OS Userid
- Connect to the z/OS Queue Manager by entering a password

Additional information

For exit CSQ4BCX3 to run in a Program Controlled environment, everything loaded into the CHIN address space must be loaded from a Program Controlled library, for example, all libraries in STEPLIB and any libraries named on CSQXLIB DD. To set a load library as Program Controlled issue RACF commands. In the following example the load library name is MY.TEST.LOADLIB.

```
RALTER PROGRAM * ADDMEM('MY.TEST.LOADLIB'//NOPADCHK)
SETROPTS WHEN(PROGRAM)REFRESH
```

To alter the SVRCONN channel to implement CSQ4BCX3, issue the following MQ command:

```
ALTER CHANNEL( SYSTEM ADMIN.SVRCONN) CHLTYPE(SVRCONN) SECYEXIT(CSQ4BCX3)
```

In the example above, the SVRCONN channel name being used is SYSTEM ADMIN.SVRCONN.

See “Channel exit programs” on page 440 for more information about channel exits.

Related concepts:



Writing channel exit programs on z/OS

SNA LU 6.2 security services:

SNA LU 6.2 offers session level cryptography, session level authentication, and conversation level authentication.

Note: This collection of topics assumes that you have a basic understanding of Systems Network Architecture (SNA). The other documentation referred to in this section contains a brief introduction to the relevant concepts and terminology. If you require a more comprehensive technical introduction to SNA, see *Systems Network Architecture Technical Overview*, GC30-3073.

SNA LU 6.2 provides three security services:

- Session level cryptography
- Session level authentication
- Conversation level authentication

For session level cryptography and session level authentication, SNA uses the *Data Encryption Standard (DES)* algorithm. The DES algorithm is a block cipher algorithm, which uses a symmetric key for encrypting and decrypting data. Both the block and the key are 8 bytes in length.

Session level cryptography:

Session level cryptography encrypts and decrypts session data using the DES algorithm. It can therefore be used to provide a link level confidentiality service on SNA LU 6.2 channels.

Logical units (LUs) can provide mandatory (or required) data cryptography, selective data cryptography, or no data cryptography.

On a *mandatory cryptographic session*, an LU encrypts all outbound data request units and decrypts all inbound data request units.

On a *selective cryptographic session*, an LU encrypts only the data request units specified by the sending transaction program (TP). The sending LU signals that the data is encrypted by setting an indicator in the request header. By checking this indicator, the receiving LU can tell which request units to decrypt before passing them on to the receiving TP.

In an SNA network, WebSphere MQ MCAs are transaction programs. MCAs do not request encryption for any data that they send. Selective data cryptography is not an option therefore; only mandatory data cryptography or no data cryptography is possible on a session.

For information about how to implement mandatory data cryptography, see the documentation for your SNA subsystem. Refer to the same documentation for information about stronger forms of encryption that might be available for use on your platform, such as Triple DES 24-byte encryption on z/OS.

For more general information about session level cryptography, see *Systems Network Architecture LU 6.2 Reference: Peer Protocols, SC31-6808*.

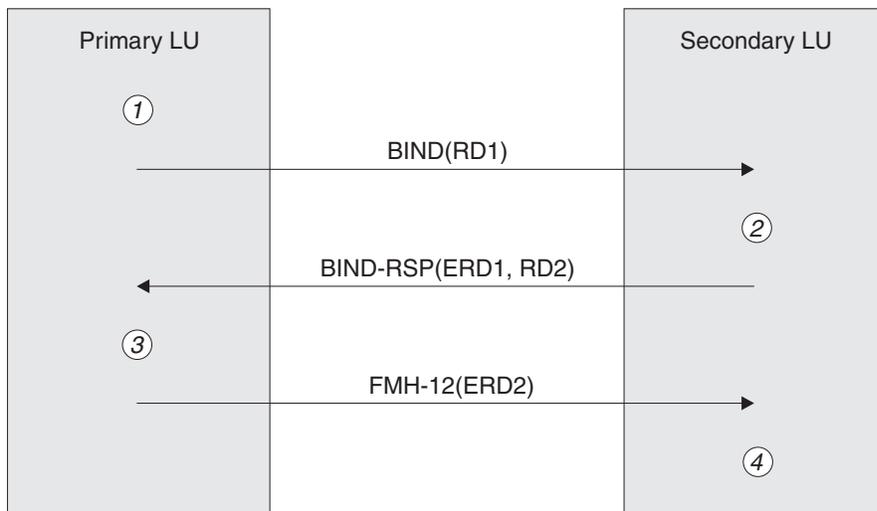
Session level authentication:

Session level authentication is a session level security protocol that enables two LUs to authenticate each other while they are activating a session. It is also known as *LU-LU verification*.

Because an LU is effectively the “gateway” into a system from the network, you might consider this level of authentication to be sufficient in certain circumstances. For example, if your queue manager needs to exchange messages with a remote queue manager that is running in a controlled and trusted environment, you might be prepared to trust the identities of the remaining components of the remote system after the LU has been authenticated.

Session level authentication is achieved by each LU verifying its partner's password. The password is called an *LU-LU password* because one password is established between each pair of LUs. The way that an LU-LU password is established is implementation dependent and outside the scope of SNA.

Figure 78 on page 455 illustrates the flows for session level authentication.



Legend:

BIND = BIND request unit
 BIND-RSP = BIND response unit
 ERD = Encrypted random data
 FMH-12 = Function Management Header 12
 RD = Random data

Figure 78. Flows for session level authentication

The protocol for session level authentication is as follows. The numbers in the procedure correspond to the numbers in Figure 78.

1. The primary LU generates a random data value (RD1) and sends it to the secondary LU in the BIND request.
2. When the secondary LU receives the BIND request with the random data, it encrypts the data using the DES algorithm with its copy of the LU-LU password as the key. The secondary LU then generates a second random data value (RD2) and sends it, with the encrypted data (ERD1), to the primary LU in the BIND response.
3. When the primary LU receives the BIND response, it computes its own version of the encrypted data from the random data it generated originally. It does this by using the DES algorithm with its copy of the LU-LU password as the key. It then compares its version with the encrypted data that it received in the BIND response. If the two values are the same, the primary LU knows that the secondary LU has the same password as it does and the secondary LU is authenticated. If the two values do not match, the primary LU terminates the session.

The primary LU then encrypts the random data that it received in the BIND response and sends the encrypted data (ERD2) to the secondary LU in a Function Management Header 12 (FMH-12).

4. When the secondary LU receives the FMH-12, it computes its own version of the encrypted data from the random data it generated. It then compares its version with the encrypted data that it received in the FMH-12. If the two values are the same, the primary LU is authenticated. If the two values do not match, the secondary LU terminates the session.

In an enhanced version of the protocol, which provides better protection against man in the middle attacks, the secondary LU computes a DES Message Authentication Code (MAC) from RD1, RD2, and the fully qualified name of the secondary LU, using its copy of the LU-LU password as the key. The secondary LU sends the MAC to the primary LU in the BIND response instead of ERD1.

The primary LU authenticates the secondary LU by computing its own version of the MAC, which it compares with the MAC received in the BIND response. The primary LU then computes a second MAC from RD1 and RD2, and sends the MAC to the secondary LU in the FMH-12 instead of ERD2.

The secondary LU authenticates the primary LU by computing its own version of the second MAC, which it compares with the MAC received in the FMH-12.

For information about how to configure session level authentication, see the documentation for your SNA subsystem. For more general information about session level authentication, see *Systems Network Architecture LU 6.2 Reference: Peer Protocols*, SC31-6808.

Conversation level authentication:

When a local TP attempts to allocate a conversation with a partner TP, the local LU sends an attach request to the partner LU, asking it to attach the partner TP. Under certain circumstances, the attach request can contain security information, which the partner LU can use to authenticate the local TP. This is known as *conversation level authentication*, or *end user verification*.

The following topics describe how WebSphere MQ provides support for conversation level authentication.

For more information about conversation level authentication, see *Systems Network Architecture LU 6.2 Reference: Peer Protocols*, SC31-6808. For information specific to z/OS, see *z/OS MVS Planning: APPC/MVS Management*, SA22-7599.

For more information about CPI-C, see *Common Programming Interface Communications CPI-C Specification*, SC31-6180. For more information about APPC/MVS TP Conversation Callable Services, see *z/OS MVS Programming: Writing Transaction Programs for APPC/MVS*, SA22-7621.

Support for conversation level authentication in WebSphere MQ on IBM i, UNIX systems, and Windows systems:

Use this topic to gain an overview of how conversation level authentication works, on platforms other than z/OS.

The support for conversation level authentication in WebSphere MQ for IBM i, WebSphere MQ on UNIX systems, and WebSphere MQ for Windows is illustrated in Figure 79 on page 457. The numbers in the diagram correspond to the numbers in the description that follows.

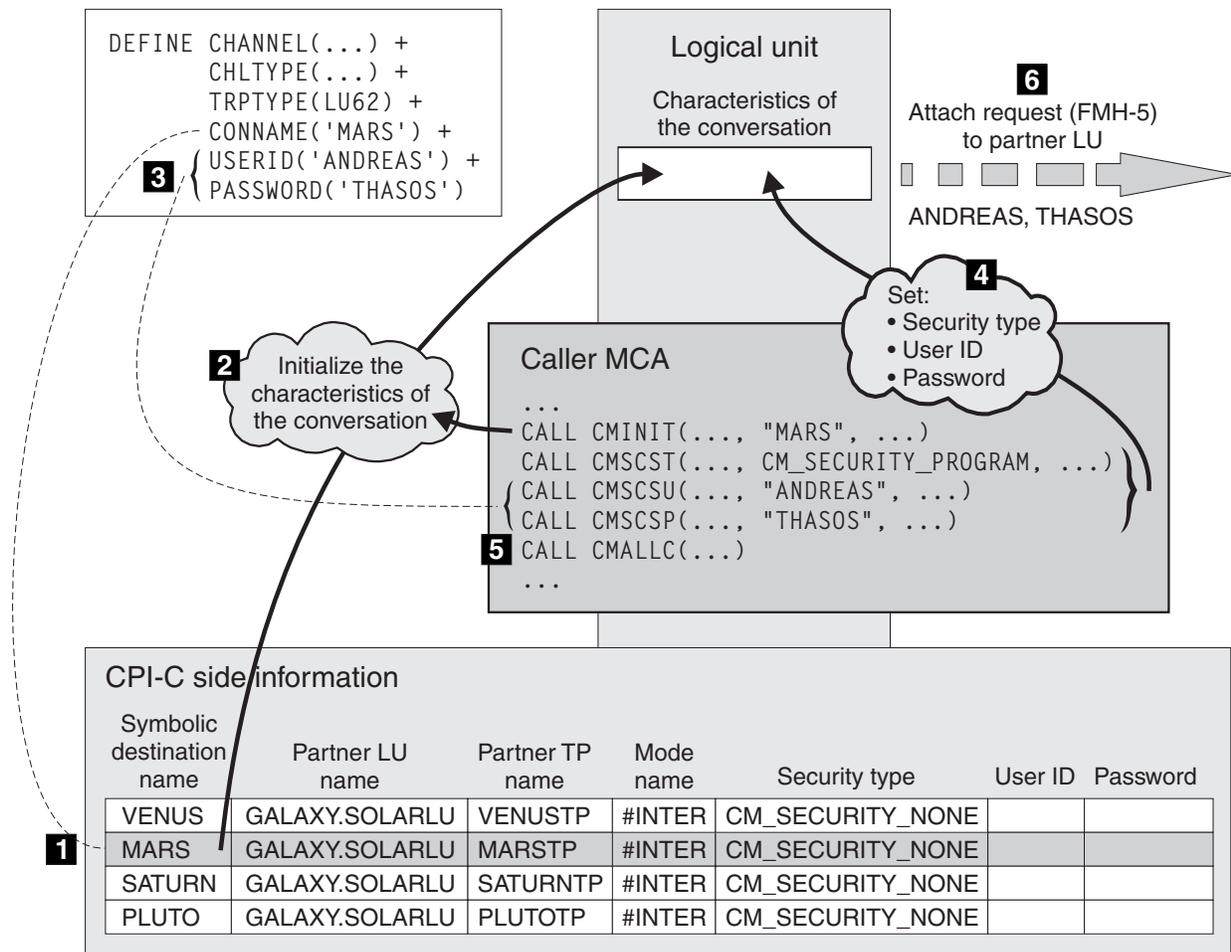


Figure 79. WebSphere MQ support for conversation level authentication

On IBM i, UNIX systems, and Windows systems, an MCA uses Common Programming Interface Communications (CPI-C) calls to communicate with a partner MCA across an SNA network. In the channel definition at the caller end of a channel, the value of the CONNAME parameter is a symbolic destination name, which identifies a CPI-C side information entry (1). This entry specifies:

- The name of the partner LU
- The name of the partner TP, which is a responder MCA
- The name of the mode to be used for the conversation

A side information entry can also specify the following security information:

- A security type.

The commonly implemented security types are CM_SECURITY_NONE, CM_SECURITY_PROGRAM, and CM_SECURITY_SAME, but others are defined in the CPI-C specification.

- A user ID.
- A password.

A caller MCA prepares to allocate a conversation with a responder MCA by issuing the CPI-C call CMINIT, using the value of CONNAME as one of the parameters on the call. The CMINIT call identifies, for the benefit of the local LU, the side information entry that the MCA intends to use for the conversation. The local LU uses the values in this entry to initialize the characteristics of the conversation (2).

The caller MCA then checks the values of the USERID and PASSWORD parameters in the channel definition (3). If USERID is set, the caller MCA issues the following CPI-C calls (4):

- CMSCST, to set the security type for the conversation to CM_SECURITY_PROGRAM.
- CMSCSU, to set the user ID for the conversation to the value of USERID.
- CMSCSP, to set the password for the conversation to the value of PASSWORD. CMSCSP is not called unless PASSWORD is set.

The security type, user ID, and password set by these calls override any values acquired previously from the side information entry.

The caller MCA then issues the CPI-C call CMALLC to allocate the conversation (5). In response to this call, the local LU sends an attach request (Function Management Header 5, or FMH-5) to the partner LU (6).

If the partner LU will accept a user ID and a password, the values of USERID and PASSWORD are included in the attach request. If the partner LU will not accept a user ID and a password, the values are not included in the attach request. The local LU discovers whether the partner LU will accept a user ID and a password as part of an exchange of information when the LUs bind to form a session.

In a later version of the attach request, a password substitute can flow between the LUs instead of a clear password. A password substitute is a DES Message Authentication Code (MAC), or an SHA-1 message digest, formed from the password. Password substitutes can be used only if both LUs support them.

When the partner LU receives an incoming attach request containing a user ID and a password, it might use the user ID and password for the purposes of identification and authentication. By referring to access control lists, the partner LU might also determine whether the user ID has the authority to allocate a conversation and attach the responder MCA.

In addition, the responder MCA might run under the user ID included in the attach request. In this case, the user ID becomes the default user ID for the responder MCA and is used for authority checks when the MCA attempts to connect to the queue manager. It might also be used for authority checks subsequently when the MCA attempts to access the queue manager's resources.

The way in which a user ID and a password in an attach request can be used for identification, authentication, and access control is implementation dependent. For information specific to your SNA subsystem, refer to the appropriate documentation.

If USERID is not set, the caller MCA does not call CMSCST, CMSCSU, and CMSCSP. In this case, the security information that flows in an attach request is determined solely by what is specified in the side information entry and what the partner LU will accept.

Conversation level authentication and WebSphere MQ for z/OS:

Use this topic to gain an overview of how conversation level authentication works, on z/OS.

On WebSphere MQ for z/OS, MCAs do not use CPI-C. Instead, they use APPC/MVS TP Conversation Callable Services, an implementation of Advanced Program-to-Program Communication (APPC), which has some CPI-C features. When a caller MCA allocates a conversation, a security type of SAME is specified on the call. Therefore, because an APPC/MVS LU supports persistent verification only for inbound conversations, not for outbound conversations, there are two possibilities:

- If the partner LU trusts the APPC/MVS LU and will accept an already verified user ID, the APPC/MVS LU sends an attach request containing:
 - The channel initiator address space user ID
 - A security profile name, which, if RACF is used, is the name of the current connect group of the channel initiator address space user ID

- An already verified indicator
- If the partner LU does not trust the APPC/MVS LU and will not accept an already verified user ID, the APPC/MVS LU sends an attach request containing no security information.

On WebSphere MQ for z/OS, the USERID and PASSWORD parameters on the DEFINE CHANNEL command cannot be used for a message channel and are valid only at the client connection end of an MQI channel. Therefore, an attach request from an APPC/MVS LU never contains values specified by these parameters.

Security for queue manager clusters

Though queue manager clusters can be convenient to use, you must pay special attention to their security.

A *queue manager cluster* is a network of queue managers that are logically associated in some way. A queue manager that is a member of a cluster is called a *cluster queue manager*.

A queue that belongs to a cluster queue manager can be made known to other queue managers in the cluster. Such a queue is called a *cluster queue*. Any queue manager in a cluster can send messages to cluster queues without needing any of the following:

- An explicit remote queue definition for each cluster queue
- Explicitly defined channels to and from each remote queue manager
- A separate transmission queue for each outbound channel

You can create a cluster in which two or more queue managers are clones. This means that they have instances of the same local queues, including any local queues declared as cluster queues, and can support instances of the same server applications.

When an application connected to a cluster queue manager sends a message to a cluster queue that has an instance on each of the cloned queue managers, WebSphere MQ decides which queue manager to send it to. When many applications send messages to the cluster queue, WebSphere MQ balances the workload across each of the queue managers that have an instance of the queue. If one of the systems hosting a cloned queue manager fails, WebSphere MQ continues to balance the workload across the remaining queue managers until the system that failed is restarted.

If you are using queue manager clusters, you need to consider the following security issues:

- Allowing only selected queue managers to send messages to your queue manager
- Allowing only selected users of a remote queue manager to send messages to a queue on your queue manager
- Allowing applications connected to your queue manager to send messages only to selected remote queues

These considerations are relevant even if you are not using clusters, but they become more important if you are using clusters.

If an application can send messages to one cluster queue, it can send messages to any other cluster queue without needing additional remote queue definitions, transmission queues, or channels. It therefore becomes more important to consider whether you need to restrict access to the cluster queues on your queue manager, and to restrict the cluster queues to which your applications can send messages.

There are some additional security considerations, which are relevant only if you are using queue manager clusters:

- Allowing only selected queue managers to join a cluster
- Forcing unwanted queue managers to leave a cluster

For more information about all these considerations, see *Keeping clusters secure*. For considerations specific to WebSphere MQ for z/OS, see “Security in queue manager clusters on z/OS” on page 592.

Security for WebSphere MQ Publish/Subscribe

There are additional security considerations if you are using WebSphere MQ Publish/Subscribe.

In a publish/subscribe system, there are two types of application: publisher and subscriber. *Publishers* supply information in the form of WebSphere MQ messages. When a publisher publishes a message, it specifies a *topic*, which identifies the subject of the information inside the message.

Subscribers are the consumers of the information that is published. A subscriber specifies the topics it is interested in by subscribing to them.

The *queue manager* is an application supplied with WebSphere MQ Publish/Subscribe. It receives published messages from publishers and subscription requests from subscribers, and routes the published messages to the subscribers. A subscriber is sent messages only on those topics to which it has subscribed.

For more information, see *Publish/subscribe security*.

Multicast security

Use this information to understand why security processes might be needed with WebSphere MQ Multicast.

WebSphere MQ Multicast does not have in-built security. Security checks are handled in the queue manager at MQOPEN time and the MQMD field setting is handled by the client. Some applications in the network might not be WebSphere MQ applications (For example, LLM applications, see “Multicast interoperability with WebSphere MQ Low Latency Messaging” on page 167 for more information), therefore you might need to implement your own security procedures because receiving applications cannot be certain of the validity of context fields.

There are three security processes to consider:

Access control

Access control in WebSphere MQ is based on user IDs. For more information on this subject, see “Access control for clients” on page 432.

Network security

An isolated network might be a viable security option to prevent fake messages. It is possible for an application on the multicast group address to publish malicious messages using native communication functions, which are indistinguishable from MQ messages because they come from an application on the same multicast group address.

It is also possible for a client on the multicast group address to receive messages that were intended for other clients on the same multicast group address.

Isolating the multicast network ensures that only valid clients and applications have access. This security precaution can prevent malicious messages from coming in, and confidential information from going out.

For information about multicast group network addresses, see: “Setting the appropriate network for multicast traffic” on page 1223

Digital signatures

A digital signature is formed by encrypting a representation of a message. The encryption uses the private key of the signatory and, for efficiency, usually operates on a message digest rather than the message itself. Digitally signing a message before an MQPUT is a good security precaution, but this process might have a detrimental effect on performance if there is a large volume of messages.

Digital signatures vary with the data being signed. If two different messages are signed digitally by the same entity, the two signatures differ, but both signatures can be verified with the same public key, that is, the public key of the entity that signed the messages.

As mentioned previously in this section, it might be possible for an application on the multicast group address to publish malicious messages using native communication functions, which are indistinguishable from MQ messages. Digital signatures provide proof of origin, and only the sender knows the private key, which provides strong evidence that the sender is the originator of the message.

For more information on this subject, see “Cryptographic concepts” on page 369.

Firewalls and Internet pass-thru

You would normally use a firewall to prevent access from hostile IP addresses, for example in a Denial of Service attack. However, you might need to temporarily block IP addresses within IBM WebSphere MQ, perhaps while you wait for a security administrator to update the firewall rules.

To block one or more IP addresses, create a channel authentication record of type BLOCKADDR or ADDRESSMAP. For more information, see “Blocking specific IP addresses” on page 733.

Security for WebSphere MQ internet pass-thru

Internet pass-thru can simplify communication through a firewall, but this has security implications.

WebSphere MQ internet pass-thru is a WebSphere MQ base product extension that is supplied in SupportPac MS81.

WebSphere MQ internet pass-thru enables two queue managers to exchange messages, or a WebSphere MQ client application to connect to a queue manager, over the Internet without requiring a direct TCP/IP connection. This is useful if a firewall prohibits a direct TCP/IP connection between two systems. It makes the passage of WebSphere MQ channel protocol flows into and out of a firewall simpler and more manageable by tunnelling the flows inside HTTP or by acting as a proxy. Using the Secure Sockets Layer (SSL), it can also be used to encrypt and decrypt messages that are sent over the Internet.

When your WebSphere MQ system communicates with IPT, unless you are using SSLProxyMode in IPT, ensure that the CipherSpec used by WebSphere MQ matches the CipherSuite used by IPT:

- When IPT is acting as the SSL or TLS server and WebSphere MQ is connecting as the SSL or TLS client, the CipherSpec used by WebSphere MQ must correspond to a CipherSuite that is enabled in the relevant IPT key ring.
- When IPT is acting as the SSL or TLS client and is connecting to a WebSphere MQ SSL or TLS server, the IPT CipherSuite must match the CipherSpec defined on the receiving WebSphere MQ channel.

If you migrate from IPT to the integrated WebSphere MQ SSL and TLS support, transfer the digital certificates from IPT Using iKeyman.

For more information about WebSphere MQ internet pass-thru, see *MS81: WebSphere MQ internet pass-thru*, available from the following address:  <http://www.ibm.com/software/integration/support/supportpacs/>

IBM WebSphere MQ for z/OS security implementation checklist

This topic gives a step-by-step procedure you can use to work out and define the security implementation for each of your IBM WebSphere MQ queue managers.

RACF provides definitions for the WebSphere MQ security classes in its supplied static Class Descriptor Table (CDT). As you work through the checklist, you can determine which of these classes your setup requires. You must ensure that they are activated as described in “RACF security classes” on page 521.

Refer to other sections for details, in particular “Profiles used to control access to WebSphere MQ resources” on page 530.

If you require security checking, follow this checklist to implement it:

1. Activate the RACF MQADMIN (uppercase profiles) and MXADMIN (mixed case profiles) classes
 - Do you want security at queue-sharing group level, queue-manager level, or a combination of both?
See, “Profiles to control queue-sharing group or queue manager level security” on page 525.
2. Do you need connection security?
 - **Yes:** Activate the MQCONN class. Define appropriate connection profiles at either queue manager level or queue-sharing group level in the MQCONN class. Then permit the appropriate users or groups access to these profiles.

Note: Only users of the MQCONN API request or CICS or IMS address space user IDs need to have access to the corresponding connection profile.
 - **No:** Define an hlq.NO.CONNECT.CHECKS profile at either queue manager level or queue-sharing group level in the MQADMIN or MXADMIN class.
3. Do you need security checking on commands?
 - **Yes:** Activate the MQCMD class. Define appropriate command profiles at either queue manager level or queue-sharing group level in the MQCMD class. Then permit the appropriate users or groups access to these profiles.
If you are using a queue-sharing group, you might need to include the user IDs used by the queue manager itself and the channel initiator. See “Setting up WebSphere MQ for z/OS resource security” on page 583.
 - **No:** Define an hlq.NO.CMD.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
4. Do you need security on the resources used in commands?
 - **Yes:** Ensure the MQADMIN or MXADMIN class is active. Define appropriate profiles for protecting resources on commands at either queue manager level or queue-sharing group level in the MQADMIN or MXADMIN class. Then permit the appropriate users or groups access to these profiles. Set the CMDUSER parameter in CSQ6SYSP to the default user ID to be used for command security checks.
If you are using a queue-sharing group, you might need to include the user IDs used by the queue manager itself and the channel initiator. See “Setting up WebSphere MQ for z/OS resource security” on page 583.
 - **No:** Define an hlq.NO.CMD.RESC.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
5. Do you need queue security?
 - **Yes:** Activate the MQQUEUE or MXQUEUE class. Define appropriate queue profiles for the required queue manager or queue-sharing group in the MQQUEUE or MXQUEUE class. Then permit the appropriate users or groups access to these profiles.
 - **No:** Define an hlq.NO.QUEUE.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.

6. Do you need process security?
 - **Yes:** Activate the MQPROC or MXPROC class. Define appropriate process profiles at either queue manager or queue-sharing group level and permit the appropriate users or groups access to these profiles.
 - **No:** Define an hlq.NO.PROCESS.CHECKS profile for the appropriate queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
7. Do you need namelist security?
 - **Yes:** Activate the MQNLIST or MXNLIST class. Define appropriate namelist profiles at either queue manager level or queue-sharing group level in the MQNLIST or MXNLIST class. Then permit the appropriate users or groups access to these profiles.
 - **No:** Define an hlq.NO.NLIST.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
8. Do you need topic security?
 - **Yes:** Activate the MXTOPIC class. Define appropriate topic profiles at either queue manager level or queue-sharing group level in the MXTOPIC class. Then permit the appropriate users or groups access to these profiles.
 - **No:** Define an hlq.NO.TOPIC.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
9. Do any users need to protect the use of the MQOPEN or MQPUT1 options relating to the use of context?
 - **Yes:** Ensure the MQADMIN or MXADMIN class is active. Define hlq.CONTEXT.queue name profiles at the queue, queue manager, or queue-sharing group level in the MQADMIN or MXADMIN class. Then permit the appropriate users or groups access to these profiles.
 - **No:** Define an hlq.NO.CONTEXT.CHECKS profile for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
10. Do you need to protect the use of alternative user IDs?
 - **Yes:** Ensure the MQADMIN or MXADMIN class is active. Define the appropriate hlq.ALTERNATE.USER.*alternateuserid* profiles for the required queue manager or queue-sharing group and permit the required users or groups access to these profiles.
 - **No:** Define the profile hlq.NO.ALTERNATE.USER.CHECKS for the required queue manager or queue-sharing group in the MQADMIN or MXADMIN class.
11. Do you need to tailor which user IDs are to be used for resource security checks through RESLEVEL?
 - **Yes:** Ensure the MQADMIN or MXADMIN class is active. Define an hlq.RESLEVEL profile at either queue manager level or queue-sharing group level in the MQADMIN or MXADMIN class. Then permit the required users or groups access to the profile.
 - **No:** Ensure that no generic profiles exist in the MQADMIN or MXADMIN class that can apply to hlq.RESLEVEL. Define an hlq.RESLEVEL profile for the required queue manager or queue-sharing group and ensure that no users or groups have access to it.
12. Do you need to 'timeout' unused user IDs from IBM WebSphere MQ?
 - **Yes:** Determine what timeout values you would like to use and issue the MQSC ALTER SECURITY command to change the TIMEOUT and INTERVAL parameters.
 - **No:** Issue the MQSC ALTER SECURITY command to set the INTERVAL value to zero.

Note: Update the CSQINP1 initialization input data set used by your subsystem so that the MQSC ALTER SECURITY command is issued automatically when the queue manager is started.
13. Do you use distributed queuing?
 - **Yes:** Use channel authentication records. For more information, see "Channel authentication records" on page 408.

- You can also determine the appropriate MCAUSER attribute value for each channel, or provide suitable channel security exits.
14. Do you want to use the Secure Sockets Layer (SSL)?
- **Yes:** To specify that any user presenting an SSL/TLS personal certificate containing a specified DN is to use a specific MCAUSER, set a channel authentication record of type SSLPEERMAP. You can specify a single distinguished name or a pattern including wildcards.
 - Plan your SSL infrastructure. Install the System SSL feature of z/OS. In RACF, set up your certificate name filters (CNFs), if you are using them, and your digital certificates. Set up your SSL key ring. Ensure that the SSLKEYR queue manager attribute is nonblank and points to your SSL key ring. Also ensure that the value of SSLTASKS is at least 2.
 - **No:** Ensure that SSLKEYR is blank, and SSLTASKS is zero.

For further details about SSL, see WebSphere MQ support for SSL and TLS.

15. Do you use clients?
- **Yes:** Use channel authentication records.
 - You can also determine the appropriate MCAUSER attribute value for each server-connection channel, or provide suitable channel security exits if required.

16. Check your switch settings.

IBM WebSphere MQ issues messages when the queue manager is started that display your security settings. Use these messages to determine whether your switches are set correctly. For an example of these messages, see User messages issued when the queue manager starts.

Setting up security

This collection of topics contains information specific to different operating systems, and to the use of clients.

Setting up security on Windows, UNIX and Linux systems

Security considerations specific to Windows, UNIX and Linux systems.

WebSphere MQ queue managers transfer information that is potentially valuable, so you need to use an authority system to ensure that unauthorized users cannot access your queue managers. Consider the following types of security controls:

Who can administer WebSphere MQ

You can define the set of users who can issue commands to administer WebSphere MQ.

Who can use WebSphere MQ objects

You can define which users (usually applications) can use MQI calls and PCF commands to do the following:

- Who can connect to a queue manager.
- Who can access objects (queues, process definitions, namelists, channels, client connection channels, listeners, services, and authentication information objects), and what type of access they have to those objects.
- Who can access WebSphere MQ messages.
- Who can access the context information associated with a message.

Channel security

You need to ensure that channels used to send messages to remote systems can access the required resources.

You can use standard operating facilities to grant access to program libraries, MQI link libraries, and commands. However, the directory containing queues and other queue manager data is private to WebSphere MQ; do not use standard operating system commands to grant or revoke authorizations to

MQI resources.

Connecting to WebSphere MQ using Terminal Services

The **Create global objects** user right can cause problems if you are using Terminal Services.

If you are connecting to a Windows system by using Terminal Services and you have problems creating or starting a queue manager, this might be because of the user right, **Create global objects**, in recent versions of Windows.

The **Create global objects** user right limits the users authorized to create objects in the global namespace. In order for an application to create a global object, it must either be running in the global namespace, or the user under which the application is running must have the **Create global objects** user right applied to it.

Administrators have the **Create global objects** user right applied by default, so an administrator can create and start queue managers when connected by using Terminal Services without altering the user rights.

If the various methods of administering WebSphere MQ do not work when you use terminal services, try setting the **Create global objects** user right:

1. Open the Administrative Tools panel:

Windows 2003 and Windows XP

Access this panel using **Control Panel > Administrative Tools**.

Windows Vista and Windows Server 2008

Access this panel using **Control Panel > System and Maintenance > Administrative Tools**.

2. Double-click **Local Security Policy**.
3. Expand **Local Policies**.
4. Click **User Rights Assignment**.
5. Add the new user or group to the **Create global objects** policy.

Creating and managing groups on Windows

These instructions lead you through the process of administering groups on a workstation or member server machine.

For domain controllers, users and groups are administered through Active Directory. For more details on using Active Directory refer to the appropriate operating system instructions.

Any changes you make to a principal's group membership are not recognized until the queue manager is restarted, or you issue the MQSC command REFRESH SECURITY (or the PCF equivalent).

Use the Computer Management panel to work with user and groups. Any changes made to the current logged on user might not be effective until the user logs in again.

Windows 2003 and Windows XP

Access this panel using **Control Panel > Administrative Tools > Computer Management**.

Windows Vista and Windows Server 2008

Access this panel using **Control Panel > System and Maintenance > Administrative Tools > Computer Management**.

Windows 7

Access this panel using **Administrative Tools > Computer Management**

Creating a group on Windows:

Create a group by using the control panel.

Procedure

1. Open the control panel
2. Double-click **Administrative Tools**. The Administrative Tools panel opens.
3. Double-click **Computer Management**. The Computer Management panel opens.
4. Expand **Local Users and Groups**.
5. Right-click **Groups**, and select **New Group...** The New Group panel is displayed.
6. Type an appropriate name in the Group name field, then click **Create**.
7. Click **Close**.

Adding a user to a group on Windows:

Add a user to a group by using the control panel.

Procedure

1. Open the control panel
2. Double-click **Administrative Tools**. The Administrative Tools panel opens.
3. Double-click **Computer Management**. The Computer Management panel opens.
4. From the Computer Management panel, expand **Local Users and Groups**.
5. Select **Users**
6. Double-click the user that you want to add to a group. The user properties panel is displayed.
7. Select the **Member Of** tab.
8. Select the group that you want to add the user to. If the group you want is not visible:
 - a. Click **Add...** The Select Groups panel is displayed.
 - b. Click **Locations...** The Locations panel is displayed.
 - c. Select the location of the group you want to add the user to from the list and click **OK**.
 - d. Type the group name in the field provided.
Alternatively, click **Advanced...** and then **Find Now** to list the groups available in the currently selected location. From here, select the group you want to add the user to and click **OK**.
 - e. Click **OK**. The user properties panel is displayed, showing the group you added.
 - f. Select the group.
9. Click **OK**. The Computer Management panel is displayed.

Displaying who is in a group on Windows:

Display the members of a group by using the control panel.

Procedure

1. Open the control panel
2. Double-click **Administrative Tools**. The Administrative Tools panel opens.
3. Double-click **Computer Management**. The Computer Management panel opens.
4. From the Computer Management panel, expand **Local Users and Groups**.
5. Select **Groups**.
6. Double-click a group. The group properties panel is displayed. The group properties panel is displayed.

Results

The group members are displayed.

Removing a user from a group on Windows:

Remove a user from a group by using the control panel.

Procedure

1. Open the control panel
2. Double-click **Administrative Tools**. The Administrative Tools panel opens.
3. Double-click **Computer Management**. The Computer Management panel opens.
4. From the Computer Management panel, expand **Local Users and Groups**.
5. Select **Users**.
6. Double-click the user that you want to add to a group. The user properties panel is displayed.
7. Select the **Member Of** tab.
8. Select the group that you want to remove the user from, then click **Remove**.
9. Click **OK**. The Computer Management panel is displayed.

Results

You have now removed the user from the group.

Creating and managing groups on HP-UX

On HP-UX, providing you are not using NIS or NIS+, use the System Administration Manager (SAM) to work with groups.

Creating a group on HP-UX:

Add a user to a group by using the System Administration Manager

Procedure

1. From the System Administration Manager (SAM), double click Accounts for Users and Groups.
2. Double click Groups.
3. Select Add from the Actions pull down to display the Add a New Group panel.
4. Enter the name of the group and select the users that you want to add to the group.
5. Click Apply to create the group.

Results

You have now created a group.

Adding a user to a group on HP-UX:

Add a user to a group by using the System Administration Manager.

Procedure

1. From the System Administration Manager (SAM), double click Accounts for Users and Groups.
2. Double click Groups.
3. Highlight the name of the group and select Modify from the Actions pull down to display the Modify an Existing Group panel.
4. Select a user that you want to add to the group and click Add.

5. If you want to add other users to the group, repeat step 4 for each user.
6. When you have finished adding names to the list, click OK.

Results

You have now added a user to a group.

Displaying who is in a group on HP-UX:

Display who is in a group by using the System Administration Manager

Procedure

1. From the System Administration Manager (SAM), double click Accounts for Users and Groups.
2. Double click Groups.
3. Highlight the name of the group and select Modify from the Actions pull down to display the Modify an Existing Group panel, showing a list of the users in the group.

Results

The group members are displayed.

Removing a user from a group on HP-UX:

Remove a user from a group by using the System Administration Manager.

Procedure

1. From the System Administration Manager (SAM), double click Accounts for Users and Groups.
2. Double click Groups.
3. Highlight the name of the group and select Modify from the Actions pull down to display the Modify an Existing Group panel.
4. Select a user that you want to remove from the group and click Remove.
5. If you want to remove other users from the group, repeat step 4 for each user.
6. When you have finished removing names from the list, click OK.

Results

You have now removed a user from a group

Creating and managing groups on AIX

On AIX, providing you are not using NIS or NIS+, use SMITTY to work with groups.

Creating a group:

Create a group using SMITTY.

Procedure

1. From SMITTY, select Security and Users and press Enter.
2. Select Groups and press Enter.
3. Select Add a Group and press Enter.
4. Enter the name of the group and the names of any users that you want to add to the group, separated by commas.
5. Press Enter to create the group.

Results

You have now created a group.

Adding a user to a group:

Add a user to a group by using SMITTY.

Procedure

1. From SMITTY, select Security and Users and press Enter.
2. Select Groups and press Enter.
3. Select Change / Show Characteristics of Groups and press Enter.
4. Enter the name of the group to show a list of the members of the group.
5. Add the names of the users that you want to add to the group, separated by commas.
6. Press Enter to add the names to the group.

Displaying who is in a group:

Display who is in a group using SMITTY.

Procedure

1. From SMITTY, select Security and Users and press Enter.
2. Select Groups and press Enter.
3. Select Change / Show Characteristics of Groups and press Enter.
4. Enter the name of the group to show a list of the members of the group.

Results

The group members are displayed.

Removing a user from a group:

Remove a user from a group by using SMITTY.

Procedure

1. From SMITTY, select Security and Users and press Enter.
2. Select Groups and press Enter.
3. Select Change / Show Characteristics of Groups and press Enter.
4. Enter the name of the group to show a list of the members of the group.
5. Delete the names of the users that you want to remove from the group.
6. Press Enter to remove the names from the group.

Results

You have now removed a user from a group.

Creating and managing groups on Solaris

On Solaris, providing you are not using NIS or NIS+, use the `/etc/group` file to work with groups.

Creating a group on Solaris:

Creating a group by using the **groupadd** command.

Procedure

Type the following command: `groupadd group-name` where *group-name* is the name of the group.

Results

The file `/etc/group` file holds group information.

Adding a user to a group on Solaris:

Add a user to a group by using the **usermod** command.

Procedure

To add a member to a supplementary group, execute the **usermod** command and list the supplementary groups that the user is currently a member of, and the supplementary groups that the user is to become a member of. For example, if the user is a member of the group `groupa`, and is to become a member of `groupb` also, use the following command: `usermod -G groupa,groupb user-name`, where *user-name* is the user name.

Displaying who is in a group on Solaris:

To discover who is a member of a group, look at the entry for that group in the `/etc/group` file.

Removing a user from a group on Solaris:

Remove a user from a group by using the **usermod** command.

Procedure

To remove a member from a supplementary group, execute the **usermod** command listing the supplementary groups that you want the user to remain a member of. For example, if the user's primary group is `users` and the user is also a member of the groups `mqm`, `groupa` and `groupb`, to remove the user from the `mqm` group, the following command is used: `usermod -G groupa,groupb user-name`, where *user-name* is the user name.

Creating and managing groups on Linux

On Linux, providing you are not using NIS or NIS+, use the `/etc/group` file to work with groups.

Creating a group on Linux:

Create a group by using the **groupadd** command.

Procedure

To create a new group, type the following command: `groupadd -g group-ID group-name` , where *group-ID* is the numeric identifier of the group, and *group-name* is the name of the group.

Results

The file `/etc/group` file holds group information.

Adding a user to a group on Linux:

Add a user to a group by using the **usermod** command.

Procedure

To add a member to a supplementary group, execute the **usermod** command and list the supplementary groups that the user is currently a member of, and the supplementary groups that the user is to become a member of. For example, if the user is a member of the group `groupa`, and is to become a member of `groupb` also, the following command is used: `usermod -G groupa,groupb user-name` , where *user-name* is the user name.

Displaying who is in a group on Linux:

Display who is in a group by using the **getent** command.

Procedure

To display who is a member of a group, type the following command: `getent group group-name` , where *group-name* is the name of the group.

Removing a user from a group:

Remove a user from a group by using the **usermod** command.

Procedure

To remove a member from a supplementary group, execute the **usermod** command listing the supplementary groups that you want the user to remain a member of. For example, if the user's primary group is `users` and the user is also a member of the groups `mqm`, `groupa` and `groupb`, to remove the user from the `mqm` group, the following command is used: `usermod -G groupa,groupb user-name` , where *user-name* is the user name.

How authorizations work

The authorization specification tables in the topics in this section define precisely how the authorizations work and the restrictions that apply.

The tables apply to these situations:

- Applications that issue MQI calls
- Administration programs that issue MQSC commands as escape PCFs
- Administration programs that issue PCF commands

In this section, the information is presented as a set of tables that specify the following:

Action to be performed

MQI option, MQSC command, or PCF command.

Access control object

Queue, process, queue manager, namelist, authentication information, channel, client connection channel, listener, or service.

Authorization required

Expressed as an MQZAO_ constant.

In the tables, the constants prefixed by MQZAO_ correspond to the keywords in the authorization list for the **setmqaut** command for the particular entity. For example, MQZAO_BROWSE corresponds to the keyword +browse, MQZAO_SET_ALL_CONTEXT corresponds to the keyword +setall, and so on. These constants are defined in the header file cmqzc.h, supplied with the product.

Authorizations for MQI calls:

MQCONN, **MQOPEN**, **MQPUT1**, and **MQCLOSE** might require authorization checks. The tables in this topic summarize the authorizations needed for each call.

An application is allowed to issue specific MQI calls and options only if the user identifier under which it is running (or whose authorizations it is able to assume) has been granted the relevant authorization.

Four MQI calls might require authorization checks: **MQCONN**, **MQOPEN**, **MQPUT1**, and **MQCLOSE**.

For **MQOPEN** and **MQPUT1**, the authority check is made on the name of the object being opened, and not on the name, or names, resulting after a name has been resolved. For example, an application might be granted authority to open an alias queue without having authority to open the base queue to which the alias resolves. The rule is that the check is carried out on the first definition encountered during the process of resolving a name that is not a queue manager alias, unless the queue manager alias definition is opened directly; that is, its name is displayed in the *ObjectName* field of the object descriptor. Authority is always needed for the object being opened. In some cases additional queue-independent authority, obtained through an authorization for the queue manager object, is required.

Table 18 on page 473, Table 19 on page 473, Table 20 on page 473, and Table 21 on page 474 summarize the authorizations needed for each call. In the tables *Not applicable* means that authorization checking is not relevant to this operation; *No check* means that no authorization checking is performed.

Note: You will find no mention of namelists, channels, client connection channels, listeners, services, or authentication information objects in these tables. This is because none of the authorizations apply to these objects, except for MQOO_INQUIRE, for which the same authorizations apply as for the other objects.

The special authorization MQZAO_ALL_MQI includes all the authorizations in the tables that are relevant to the object type, except MQZAO_DELETE and MQZAO_DISPLAY, which are classed as administration authorizations.

In order to modify any of the message context options, you must have the appropriate authorizations to issue the call. For example, in order to use MQOO_SET_IDENTITY_CONTEXT or MQPMO_SET_IDENTITY_CONTEXT, you must have +setid permission.

Table 18. Security authorization needed for MQCONN calls

Authorization required for:	Queue object (1 on page 474)	Process object	Queue manager object
MQCONN	Not applicable	Not applicable	MQZAO_CONNECT

Table 19. Security authorization needed for MQOPEN calls

Authorization required for:	Queue object (1 on page 474)	Process object	Queue manager object
MQOO_INQUIRE	MQZAO_INQUIRE	MQZAO_INQUIRE	MQZAO_INQUIRE
MQOO_BROWSE	MQZAO_BROWSE	Not applicable	No check
MQOO_INPUT_*	MQZAO_INPUT	Not applicable	No check
MQOO_SAVE_ALL_CONTEXT (2 on page 474)	MQZAO_INPUT	Not applicable	Not applicable
MQOO_OUTPUT (Normal queue) (3 on page 474)	MQZAO_OUTPUT	Not applicable	Not applicable
MQOO_PASS_IDENTITY_CONTEXT (4 on page 474)	MQZAO_PASS_IDENTITY_CONTEXT	Not applicable	No check
MQOO_PASS_ALL_CONTEXT (4 on page 474, 5 on page 474)	MQZAO_PASS_ALL_CONTEXT	Not applicable	No check
MQOO_SET_IDENTITY_CONTEXT (4 on page 474, 5 on page 474)	MQZAO_SET_IDENTITY_CONTEXT	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (6 on page 474)
MQOO_SET_ALL_CONTEXT (4 on page 474, 7 on page 474)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (6 on page 474)
MQOO_OUTPUT (Transmission queue) (8 on page 474)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (6 on page 474)
MQOO_SET	MQZAO_SET	Not applicable	No check
MQOO_ALTERNATE_USER_AUTHORITY	(9 on page 474)	(9 on page 474)	MQZAO_ALTERNATE_USER_AUTHORITY (9 on page 474, 10 on page 474)

Table 20. Security authorization needed for MQPUT1 calls

Authorization required for:	Queue object (1 on page 474)	Process object	Queue manager object
MQPMO_PASS_IDENTITY_CONTEXT	MQZAO_PASS_IDENTITY_CONTEXT (11 on page 474)	Not applicable	No check
MQPMO_PASS_ALL_CONTEXT	MQZAO_PASS_ALL_CONTEXT (11 on page 474)	Not applicable	No check
MQPMO_SET_IDENTITY_CONTEXT	MQZAO_SET_IDENTITY_CONTEXT (11 on page 474)	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (6 on page 474)

Table 20. Security authorization needed for MQPUT1 calls (continued)

Authorization required for:	Queue object (1)	Process object	Queue manager object
MQPMO_SET_ ALL_CONTEXT	MQZAO_SET_ ALL_CONTEXT (11)	Not applicable	MQZAO_SET_ ALL_CONTEXT (6)
(Transmission queue) (8)	MQZAO_SET_ ALL_CONTEXT	Not applicable	MQZAO_SET_ ALL_CONTEXT (6)
MQPMO_ALTERNATE_ USER_AUTHORITY	(12)	Not applicable	MQZAO_ALTERNATE_ USER_AUTHORITY (10)

Table 21. Security authorization needed for MQCLOSE calls

Authorization required for:	Queue object (1)	Process object	Queue manager object
MQCO_DELETE	MQZAO_DELETE (13)	Not applicable	Not applicable
MQCO_DELETE_PURGE	MQZAO_DELETE (13)	Not applicable	Not applicable

Notes for the tables:

1. If opening a model queue:
 - MQZAO_DISPLAY authority is needed for the model queue, in addition to the authority to open the model queue for the type of access for which you are opening.
 - MQZAO_CREATE authority is not needed to create the dynamic queue.
 - The user identifier used to open the model queue is automatically granted all the queue-specific authorities (equivalent to MQZAO_ALL) for the dynamic queue created.
2. MQOO_INPUT_* must also be specified. This is valid for a local, model, or alias queue.
3. This check is performed for all output cases, except transmission queues (see note 8).
4. MQOO_OUTPUT must also be specified.
5. MQOO_PASS_IDENTITY_CONTEXT is also implied by this option.
6. This authority is required for both the queue manager object and the particular queue.
7. MQOO_PASS_IDENTITY_CONTEXT, MQOO_PASS_ALL_CONTEXT, and MQOO_SET_IDENTITY_CONTEXT are also implied by this option.
8. This check is performed for a local or model queue that has a *Usage* queue attribute of MQUS_TRANSMISSION, and is being opened directly for output. It does not apply if a remote queue is being opened (either by specifying the names of the remote queue manager and remote queue, or by specifying the name of a local definition of the remote queue).
9. At least one of MQOO_INQUIRE (for any object type), or MQOO_BROWSE, MQOO_INPUT_*, MQOO_OUTPUT, or MQOO_SET (for queues) must also be specified. The check carried out is as for the other options specified, using the supplied alternate-user identifier for the specific-named object authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.
10. This authorization allows any *AlternateUserId* to be specified.
11. An MQZAO_OUTPUT check is also carried out if the queue does not have a *Usage* queue attribute of MQUS_TRANSMISSION.
12. The check carried out is as for the other options specified, using the supplied alternate-user identifier for the specific-named queue authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.
13. The check is carried out only if both of the following are true:
 - A permanent dynamic queue is being closed and deleted.
 - The queue was not created by the **MQOPEN** call that returned the object handle being used.

Otherwise, there is no check.

Authorizations for MQSC commands in escape PCFs:

This information summarizes the authorizations needed for each MQSC command contained in Escape PCF.

Not applicable means that this operation is not relevant to this object type.

The user ID under which the program that submits the command is running must also have the following authorities:

- MQZAO_CONNECT authority to the queue manager
- MQZAO_DISPLAY authority on the queue manager in order to perform PCF commands
- Authority to issue the MQSC command within the text of the Escape PCF command

ALTER *object*

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	MQZAO_CHANGE
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE
Communication information	MQZAO_CHANGE

CLEAR *object*

Object	Authorization required
Queue	MQZAO_CLEAR
Topic	MQZAO_CLEAR
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	Not applicable
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable
Communication information	Not applicable

DEFINE *object* NOREPLACE (1 on page 479)

Object	Authorization required
Queue	MQZAO_CREATE (2 on page 479)
Topic	MQZAO_CREATE (2 on page 479)
Process	MQZAO_CREATE (2 on page 479)
Queue manager	Not applicable
Namelist	MQZAO_CREATE (2 on page 479)
Authentication information	MQZAO_CREATE (2 on page 479)
Channel	MQZAO_CREATE (2 on page 479)
Client connection channel	MQZAO_CREATE (2 on page 479)
Listener	MQZAO_CREATE (2 on page 479)
Service	MQZAO_CREATE (2 on page 479)
Communication information	MQZAO_CREATE (2 on page 479)

DEFINE *object* REPLACE (1 on page 479, 3 on page 479)

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	Not applicable
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE
Communication information	MQZAO_CHANGE

DELETE *object*

Object	Authorization required
Queue	MQZAO_DELETE
Topic	MQZAO_DELETE
Process	MQZAO_DELETE
Queue manager	Not applicable
Namelist	MQZAO_DELETE
Authentication information	MQZAO_DELETE
Channel	MQZAO_DELETE
Client connection channel	MQZAO_DELETE
Listener	MQZAO_DELETE
Service	MQZAO_DELETE
Communication information	MQZAO_DELETE

DISPLAY object

Object	Authorization required
Queue	MQZAO_DISPLAY
Topic	MQZAO_DISPLAY
Process	MQZAO_DISPLAY
Queue manager	MQZAO_DISPLAY
Namelist	MQZAO_DISPLAY
Authentication information	MQZAO_DISPLAY
Channel	MQZAO_DISPLAY
Client connection channel	MQZAO_DISPLAY
Listener	MQZAO_DISPLAY
Service	MQZAO_DISPLAY
Communication information	MQZAO_DISPLAY

START object

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	MQZAO_CONTROL
Service	MQZAO_CONTROL
Communication information	Not applicable

STOP object

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	MQZAO_CONTROL
Service	MQZAO_CONTROL
Communication information	Not applicable

Channel Commands

Command	Object	Authorization required
PING CHANNEL	Channel	MQZAO_CONTROL
RESET CHANNEL	Channel	MQZAO_CONTROL_EXTENDED
RESOLVE CHANNEL	Channel	MQZAO_CONTROL_EXTENDED

Subscription Commands

Command	Object	Authorization required
ALTER SUB	Topic	MQZAO_CONTROL
DEFINE SUB	Topic	MQZAO_CONTROL
DELETE SUB	Topic	MQZAO_CONTROL
DISPLAY SUB	Topic	MQZAO_DISPLAY

Security Commands

Command	Object	Authorization required
SET AUTHREC	Queue manager	MQZAO_CHANGE
DELETE AUTHREC	Queue manager	MQZAO_CHANGE
DISPLAY AUTHREC	Queue manager	MQZAO_DISPLAY
DISPLAY AUTHSERV	Queue manager	MQZAO_DISPLAY
DISPLAY ENTAUTH	Queue manager	MQZAO_DISPLAY
SET CHLAUTH	Queue manager	MQZAO_CHANGE
DISPLAY CHLAUTH	Queue manager	MQZAO_DISPLAY
REFRESH SECURITY	Queue manager	MQZAO_CHANGE

Status Displays

Command	Object	Authorization required
DISPLAY CHSTATUS	Queue manager	MQZAO_DISPLAY
DISPLAY LSSTATUS	Queue manager	MQZAO_DISPLAY
DISPLAY PUBSUB	Queue manager	MQZAO_DISPLAY
DISPLAY SBSTATUS	Queue manager	MQZAO_DISPLAY
DISPLAY SVSTATUS	Queue manager	MQZAO_DISPLAY
DISPLAY TPSTATUS	Queue manager	MQZAO_DISPLAY

Cluster Commands

Command	Object	Authorization required
DISPLAY CLUSQMGR	Queue manager	MQZAO_DISPLAY
REFRESH CLUSTER	'mqm' group membership required	
RESET CLUSTER	'mqm' group membership required	
SUSPEND QMGR	'mqm' group membership required	
RESUME QMGR	'mqm' group membership required	

Other Administrative Commands

Command	Object	Authorization required
PING QMGR	Queue manager	MQZAO_DISPLAY
REFRESH QMGR	Queue manager	MQZAO_CHANGE
RESET QMGR	Queue manager	MQZAO_CHANGE
DISPLAY CONN	Queue manager	MQZAO_DISPLAY
STOP CONN	Queue manager	MQZAO_CHANGE

Note:

1. For DEFINE commands, MQZAO_DISPLAY authority is also needed for the LIKE object if one is specified, or on the appropriate SYSTEM.DEFAULT.xxx object if LIKE is omitted.
2. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects for a specified queue manager, by specifying an object type of QMGR on the **setmqaut** command.
3. This applies if the object to be replaced already exists. If it does not, the check is as for DEFINE *object* NOREPLACE.

Related concepts:

 Clustering: Using REFRESH CLUSTER best practices

Authorizations for PCF commands:

This section summarizes the authorizations needed for each PCF command.

No check means that no authorization checking is carried out; *Not applicable* means that this operation is not relevant to this object type.

The user ID under which the program that submits the command is running must also have the following authorities:

- MQZAO_CONNECT authority to the queue manager
- MQZAO_DISPLAY authority on the queue manager in order to perform PCF commands

The special authorization MQZAO_ALL_ADMIN includes all the authorizations in the following list that are relevant to the object type, except MQZAO_CREATE, which is not specific to a particular object or object type.

Change *object*

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Queue manager (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE

Clear object

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CLEAR
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CLEAR
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	Not applicable
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable
Communication information	Not applicable

Copy object (without replace) (1)

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
Queue manager	Not applicable
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2 on page 487)

Copy object (with replace) (1, 4)

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
Queue manager	Not applicable
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE

Create object (without replace) (3)

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
Queue manager	Not applicable
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CREATE (2)

Create object (with replace) (3, 4)

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
Queue manager	Not applicable
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CHANGE

Delete object

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
Queue manager	Not applicable
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DELETE

Inquire object

Object	Authorization required
 Queue (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Topic (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Process (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Queue manager (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Namelist (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Authentication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Client connection channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY
 Communication information (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_DISPLAY

Inquire *object* names

Object	Authorization required
Queue	No check
Topic	No check
Process	No check
Queue manager	No check
Namelist	No check
Authentication information	No check
Channel	No check
Client connection channel	No check
Listener	No check
Service	No check
Communication information	No check

Start *object*

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
Client connection channel	Not applicable
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
Communication information	Not applicable

Stop *object*

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
 Channel (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
Client connection channel	Not applicable

Object	Authorization required
 Listener (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
 Service (<i>WebSphere MQ V7.1 Reference</i>)	MQZAO_CONTROL
Communication information	Not applicable

Channel Commands

Command	Object	Authorization required
 Ping Channel (<i>WebSphere MQ V7.1 Reference</i>)	Channel	MQZAO_CONTROL
 Reset Channel (<i>WebSphere MQ V7.1 Reference</i>)	Channel	MQZAO_CONTROL_EXTENDED
 Resolve Channel (<i>WebSphere MQ V7.1 Reference</i>)	Channel	MQZAO_CONTROL_EXTENDED

Subscription Commands

Command	Object	Authorization required
 Change Subscription (<i>WebSphere MQ V7.1 Reference</i>)	Topic	MQZAO_CONTROL
 Create Subscription (<i>WebSphere MQ V7.1 Reference</i>)	Topic	MQZAO_CONTROL
 Delete Subscription (<i>WebSphere MQ V7.1 Reference</i>)	Topic	MQZAO_CONTROL
 Inquire Subscription (<i>WebSphere MQ V7.1 Reference</i>)	Topic	MQZAO_DISPLAY

Security Commands

Command	Object	Authorization required
 Set Authority Record (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE
 Delete Authority Record (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE
 Inquire Authority Records (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Authority Service (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Entity Authority (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY

Command	Object	Authorization required
 Set Channel Authentication Record (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE
 Inquire Channel Authentication Records (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Refresh Security (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE

Status Displays

Command	Object	Authorization required
 Inquire Channel Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Channel Listener Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Pub/Sub Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Subscription Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Service Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Inquire Topic Status (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY

Cluster Commands

Command	Object	Authorization required
 Inquire Cluster Queue Manager (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Refresh Cluster (<i>WebSphere MQ V7.1 Reference</i>)	'mqm' group membership required	
 Reset Cluster (<i>WebSphere MQ V7.1 Reference</i>)	'mqm' group membership required	
 Suspend Queue Manager Cluster (<i>WebSphere MQ V7.1 Reference</i>)	'mqm' group membership required	
 Resume Queue Manager Cluster (<i>WebSphere MQ V7.1 Reference</i>)	'mqm' group membership required	

Other Administrative Commands

Command	Object	Authorization required
 Ping Queue Manager (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Refresh Queue Manager (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE
 Reset Queue Manager (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE
 Reset Queue Statistics (<i>WebSphere MQ V7.1 Reference</i>)	Queue	MQZAO_DISPLAY and MQZAO_CHANGE
 Inquire Connection (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_DISPLAY
 Stop Connection (<i>WebSphere MQ V7.1 Reference</i>)	Queue manager	MQZAO_CHANGE

Note:

1. For Copy commands, MQZAO_DISPLAY authority is also needed for the From object.
2. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects for a specified queue manager, by specifying an object type of QMGR on the **setmqaut** command.
3. For Create commands, MQZAO_DISPLAY authority is also needed for the appropriate SYSTEM.DEFAULT.* object.
4. This applies if the object to be replaced already exists. If it does not, the check is as for Copy or Create without replace.

Special considerations for security on Windows

Some security functions behave differently on different version of Windows.

WebSphere MQ for Windows runs on Windows 2003, Windows Vista, Windows Server 2008, and Windows XP, but the operation of WebSphere MQ security can be affected by differences between the platforms.

WebSphere MQ security relies on calls to the operating system API for information about user authorizations and group memberships. Some functions do not behave identically on the Windows systems. This collection of topics includes descriptions of how those differences might affect WebSphere MQ security when you are running WebSphere MQ in a Windows environment.

The SSPI channel exit program:

WebSphere MQ for Windows supplies a security exit program, which can be used on both message and MQI channels. The exit is supplied as source and object code, and provides one-way and two-way authentication.

The security exit uses the Security Support Provider Interface (SSPI), which provides the integrated security facilities of Windows platforms.

The security exit provides the following identification and authentication services:

One way authentication

This uses Windows NT LAN Manager (NTLM) authentication support. NTLM allows servers to

authenticate their clients. It does not allow a client to authenticate a server, or one server to authenticate another. NTLM was designed for a network environment in which servers are assumed to be genuine. NTLM is supported on all Windows platforms that are supported by WebSphere MQ Version 7.0.

This service is typically used on an MQI channel to enable a server queue manager to authenticate a WebSphere MQ MQI client application. A client application is identified by the user ID associated with the process that is running.

To perform the authentication, the security exit at the client end of a channel acquires an authentication token from NTLM and sends the token in a security message to its partner at the other end of the channel. The partner security exit passes the token to NTLM, which checks that the token is authentic. If the partner security exit is not satisfied with the authenticity of the token, it instructs the MCA to close the channel.

Two way, or mutual, authentication

This uses Kerberos authentication services. The Kerberos protocol does not assume that servers in a network environment are genuine. Servers can authenticate clients and other servers, and clients can authenticate servers. Kerberos is supported on all Windows platforms that are supported by WebSphere MQ Version 7.0.

This service can be used on both message and MQI channels. On a message channel, it provides mutual authentication of the two queue managers. On an MQI channel, it enables the server queue manager and the WebSphere MQ MQI client application to authenticate each other. A queue manager is identified by its name prefixed by the string `ibmMQSeries/`. A client application is identified by the user ID associated with the process that is running.

To perform the mutual authentication, the initiating security exit acquires an authentication token from the Kerberos security server and sends the token in a security message to its partner. The partner security exit passes the token to the Kerberos server, which checks that it is authentic. The Kerberos security server generates a second token, which the partner sends in a security message to the initiating security exit. The initiating security exit then asks the Kerberos server to check that the second token is authentic. During this exchange, if either security exit is not satisfied with the authenticity of the token sent by the other, it instructs the MCA to close the channel.

The security exit is supplied in both source and object format. You can use the source code as a starting point for writing your own channel exit programs or you can use the object module as supplied. The object module has two entry points, one for one way authentication using NTLM authentication support and the other for two way authentication using Kerberos authentication services.

For more information about how the SSPI channel exit program works, and for instructions on how to implement it, see  Using the SSPI security exit on Windows systems (*WebSphere MQ V7.1 Programming Guide*).

When you get a 'group not found' error on Windows:

This problem can arise because WebSphere MQ loses access to the local `mqm` group when Windows servers are promoted to, or demoted from, domain controllers. To remedy this problem, re-create the local `mqm` group.

The symptom is an error indicating the lack of a local `mqm` group, for example:

```
>crtmqm qm0  
AMQ8066:Local mqm group not found.
```

Altering the state of a machine between server and domain controller can affect the operation of WebSphere MQ, because WebSphere MQ uses a locally-defined `mqm` group. When a server is promoted to be a domain controller, the scope changes from local to domain local. When the machine is demoted to

server, all domain local groups are removed. This means that changing a machine from server to domain controller and back to server loses access to a local mqm group.

To remedy this problem, re-create the local mqm group using the standard Windows management tools. Because all group membership information is lost, you must reinstate privileged WebSphere MQ users in the newly-created local mqm group. If the machine is a domain member, you must also add the domain mqm group to the local mqm group to grant privileged domain WebSphere MQ user IDs the required level of authority.

When you have problems with WebSphere MQ and domain controllers on Windows:

Certain problems can arise with security settings when Windows servers are promoted to domain controllers.

While promoting Windows 2000, Windows 2003, or Windows Server 2008 servers to domain controllers, you are presented with the option of selecting a default or non-default security setting relating to user and group permissions. This option controls whether arbitrary users are able to retrieve group memberships from the active directory. Because WebSphere MQ relies on group membership information to implement its security policy, it is important that the user ID that is performing WebSphere MQ operations can determine the group memberships of other users.

On Windows 2000, when a domain is created using the default security option, the default user ID created by WebSphere MQ during the installation process can obtain group memberships for other users as required. The product then installs normally, creating default objects, and the queue manager can determine the access authority of local and domain users if required.

On Windows 2000, when a domain is created using the non-default security option, or on Windows 2003 and Windows Server 2008 when a domain is created using the default security option, the user ID created by WebSphere MQ during the installation cannot always determine the required group memberships. In this case, you need to know:

- How Windows 2000 with non-default, or Windows 2003 and Windows Server 2008 with default, security permissions behaves
- How to allow domain mqm group members to read group membership
- How to configure a IBM WebSphere MQ Windows service to run under a domain user

Windows 2000 domain with non-default, or Windows 2003 and Windows Server 2008 domain with default, security permissions:

Installation of WebSphere MQ behaves differently on these operating systems depending on whether a local user or domain user performs the installation.

If a **local** user installs WebSphere MQ, the Prepare WebSphere MQ Wizard detects that the local user created for the IBM WebSphere MQ Windows service can retrieve the group membership information of the installing user. The Prepare WebSphere MQ Wizard asks the user questions about the network configuration to determine whether there are other user accounts defined on domain controllers running on Windows 2000 or later. If so, the IBM WebSphere MQ Windows service needs to run under a domain user account with particular settings and authorities. The Prepare WebSphere MQ Wizard prompts the user for the account details of this user. Its online help provides details of the domain user account required that can be sent to the domain administrator.

If a **domain** user installs WebSphere MQ, the Prepare WebSphere MQ Wizard detects that the local user created for the IBM WebSphere MQ Windows service cannot retrieve the group membership information of the installing user. In this case, the Prepare WebSphere MQ Wizard always prompts the user for the account details of the domain user account for the IBM WebSphere MQ Windows service to use.

When the IBM WebSphere MQ Windows service needs to use a domain user account, WebSphere MQ cannot operate correctly until this has been configured using the Prepare WebSphere MQ Wizard. The Prepare WebSphere MQ Wizard does not allow the user to continue with other tasks, until the Windows service has been configured with a suitable account.

If a Windows 2000 domain has been configured with non-default security permissions, the usual solution to enable WebSphere MQ to work correctly is to configure it with a suitable domain user account, as described above.

See  [Creating and setting up domain accounts for WebSphere MQ \(WebSphere MQ V7.1 Installing Guide\)](#) for more information.

Configuring IBM WebSphere MQ Services to run under a domain user on Windows:

Use the Prepare IBM WebSphere MQ wizard to enter the account details of the domain user account. Alternatively, you can use the Computer Management panel to alter the **Log On** details for the installation specific IBM WebSphere MQ Service.

For more information see “Changing the password of the IBM WebSphere MQ Windows service user account” on page 70

Applying security template files to Windows:

Applying a template might affect the security settings applied to WebSphere MQ files and directories. If you use the highly secure template, apply it before installing WebSphere MQ.

Windows supports text-based security template files that you can use to apply uniform security settings to one or more computers with the Security Configuration and Analysis MMC snap-in. In particular, Windows supplies several templates that include a range of security settings with the aim of providing specific levels of security. These templates include Compatible, Secure, and Highly Secure.

Applying one of these templates might affect the security settings applied to WebSphere MQ files and directories. If you want to use the Highly Secure template, configure your machine before you install WebSphere MQ.

If you apply the highly secure template to a machine on which WebSphere MQ is already installed, all the permissions you have set on the WebSphere MQ files and directories are removed. Because these permissions are removed, you lose *Administrator*, *mqm*, and, when applicable, *Everyone* group access from the error directories.

Nested groups:

There are restrictions on the use of nested groups. These result partly from the domain functional level and partly from WebSphere MQ restrictions.

Active Directory can support different group types within a Domain context depending on the Domain functional level. By default, Windows 2003 domains are in the *Windows 2000 mixed* functional level. (Windows server 2003, Windows XP, Windows Vista, and Windows Server 2008 all follow the Windows 2003 domain model.) The domain functional level determines the supported group types and level of nesting allowed when configuring user IDs in a domain environment. Refer to Active Directory documentation for details on the Group Scope and inclusion criteria.

In addition to Active Directory requirements, further restrictions are imposed on IDs used by WebSphere MQ. The network APIs used by WebSphere MQ do not support all the configurations that are supported by the domain functional level. As a result, WebSphere MQ is not able to query the group memberships of any Domain IDs present in a Domain Local group which is then nested in a local group. Furthermore,

multiple nesting of global and universal groups is not supported. However, immediately nested global or universal groups are supported.

Configuring additional authority for Windows applications connecting to WebSphere MQ:

The account under which WebSphere MQ processes run might need additional authorization before SYNCHRONIZE access to application processes can be granted.

You might experience problems if you have Windows applications, for example ASP pages, connecting to WebSphere MQ that are configured to run at a security level higher than usual.

WebSphere MQ requires SYNCHRONIZE access to application processes in order to coordinate certain actions. APAR IC35116 changed WebSphere MQ so that the appropriate privileges are specified. However, the account under which WebSphere MQ processes run might need additional authorization before the requested access can be granted.

When a server application first attempts to connect to a queue manager WebSphere MQ will modify the process to grant SYNCHRONIZE authority for WebSphere MQ administrators. To configure additional authority to the user ID under which WebSphere MQ processes are running, complete the following steps:

1. Start the Local Security Policy tool, click Security Settings ->Local Policies->User Right Assignments, click "Debug Programs".
2. Double click "Debug Programs", then add your WebSphere MQ user ID to the list

If the system is in a Windows domain and the effective policy setting is still not set, even though the local policy setting is set, the user ID must be authorized in the same way at domain level, using the Domain Security Policy tool.

Setting up security on HP Integrity NonStop Server

Security considerations specific to HP Integrity NonStop Server systems.

The IBM WebSphere MQ client for HP Integrity NonStop Server supports both the Transport Layer Security (TLS) and the Secure Sockets Layer (SSL) protocols to provide link level security when you are connecting to a queue manager. These protocols are supported by using an implementation of OpenSSL. OpenSSL requires a source of random data for providing strong cryptographic operations.

OpenSSL

OpenSSL security overview for IBM WebSphere MQ client for HP Integrity NonStop Server.

The OpenSSL toolkit is an open source implementation of the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols for secure communications over a network.

The toolkit is developed by the OpenSSL Project. For more information about the OpenSSL Project, see

 <http://www.openssl.org>. IBM WebSphere MQ client for HP Integrity NonStop Server contains modified versions of the OpenSSL libraries and the **openssl** command. The libraries and **openssl** command are ported from the OpenSSL toolkit 1.0.1c, and are supplied as object code only. No source code is provided.

The OpenSSL libraries are loaded by IBM WebSphere MQ client application programs dynamically as required. Only the OpenSSL libraries that are provided by IBM WebSphere MQ are supported for use with IBM WebSphere MQ client applications.

The **openssl** command, which can be used for certificate management purposes, is installed in the OSS directory `opt_installation_path/opt/mqm/bin`.

Using the **openssl** command, you can create and manage keys and digital certificates with various common data formats, and carry out simple certificate authority (CA) tasks.

The default format for key and certificate data that is processed by OpenSSL is the Privacy Enhanced Mail (PEM) format. Data in PEM format is base64 encoded ASCII data. The data can therefore be transferred by using text-based systems such as email, and can be cut and pasted by using text editors and web browsers. PEM is an Internet standard for text-based cryptographic exchanges and is specified in Internet RFCs 1421, 1422, 1423, and 1424. IBM WebSphere MQ assumes that a file with extension `.pem` contains data in PEM format. A file in PEM format can contain multiple certificates and other encoded objects, and can include comments.

The IBM WebSphere MQ SSL support on other operating systems might require key and certificate data in files to be encoded by using Distinguished Encoding Rules (DER). DER is a set of encoding rules for using the ASN.1 notation in secure communications. Data that is encoded by using DER is binary data, and the format of key and certificate data that is encoded by using DER is also known as PKCS#12 or PFX. A file that contains this data commonly has an extension of `.p12` or `.pfx`. The **openssl** command can convert between PEM and PKCS#12 format.

Entropy Daemon

OpenSSL requires a source of random data for providing strong cryptographic operations. Random number generation is a capability that is usually provided by the operating system or by a system-wide daemon process. The HP Integrity NonStop Server operating system does not provide this capability within the operating system.

When you are using the SSL and TLS support supplied with IBM WebSphere MQ client for HP Integrity NonStop Server, a process that is called an entropy daemon is needed to provide the source of random data. When you start a client channel that requires SSL or TLS, OpenSSL expects an entropy daemon to be running and providing its services on a socket in the OSS file system at `/etc/egd-pool`.

An entropy daemon is not provided by IBM WebSphere MQ client for HP Integrity NonStop Server. The HP Integrity NonStop Server client is tested with the following entropy daemons:

- `amqjkd0` (as provided by the IBM WebSphere MQ 5.3 server)
- `/usr/local/bin/prngd` (Version 0.9.27, as provided by HP Integrity NonStop Server Open Source Technical Library)

Setting up security on IBM i

Security for WebSphere MQ for IBM i is implemented using the WebSphere MQ Object Authority Manager (OAM) and IBM i object level security.

Security considerations that must be made when determining access authority to WebSphere MQ objects.

You need to consider the following points when setting up authorities to the users in your enterprise:

1. Grant and revoke authorities to the WebSphere MQ for IBM i commands using the IBM i **GRTOBJAUT** and **RVKOBJAUT** commands.

In the QMQM library, certain noncommand (`*cmd`) objects are set to have ***PUBLIC** authority to ***USE**. Do not change the authorities of these objects or use an authorization list to provide authority. Any incorrect authority might compromise WebSphere MQ functionality.

2. During installation of WebSphere MQ for IBM i, the following special user profiles are created:

QMQM

Is used primarily for internal product-only functions. However, it can be used to run trusted applications using `MQCNO_FASTPATH_BINDINGS`; see  Connecting to a queue manager using the `MQCONN` call (*WebSphere MQ V7.1 Programming Guide*) for further information.

QMOMADM

Is used as a group profile for administrators of WebSphere MQ. The group profile gives access to CL commands and WebSphere MQ resources.

When using SBMJOB to submit programs that call WebSphere MQ commands, USER must not be set explicitly to QMOMADM. Instead, set USER to QMQM or another user profile that has QMOMADM specified as a group.

3. If you are sending channel commands to remote queue managers, ensure that your user profile is a member of the group QMOMADM on the target system. For a list of PCF and MQSC channel commands, see  WebSphere MQ for IBM i CL commands (*WebSphere MQ V7.1 Reference*).
4. The group set associated with a user is cached when the group authorizations are computed by the OAM.
Any changes made to a user's group memberships after the group set has been cached are not recognized until you restart the queue manager or execute RFRMQMAUT to refresh security.
5. Limit the number of users who have authority to work with commands that are particularly sensitive. These commands include:
 - Create Message Queue Manager (CRTMQM)
 - Delete Message Queue Manager (DLTMQM)
 - Start Message Queue Manager (STRMQM)
 - End Message Queue Manager (ENDMQM)
 - Start Command Server (STRMQMCSVR)
 - End Command Server (ENDMQMCSVR)
6. Channel definitions contain a security exit program specification. Channel creation and modification requires special considerations. Details of security exits are given in "Security exit overview" on page 441.
7. The channel exit and trigger monitor programs can be substituted. The security of such replacements is the responsibility of the programmer.

The object authority manager on IBM i

The object authority manager (OAM) manages users' authorizations to manipulate WebSphere MQ objects, including queues and process definitions. It also provides a command interface through which you can grant or revoke access authority to an object for a specific group of users. The decision to allow access to a resource is made by the OAM, and the queue manager follows that decision. If the OAM cannot make a decision, the queue manager prevents access to that resource.

Through the OAM you can control:

- Access to WebSphere MQ objects through the MQI. When an application program attempts to access an object, the OAM checks that the user profile making the request has the authorization for the operation requested.

In particular, this means that queues, and the messages on queues, can be protected from unauthorized access.

- Permission to use PCF and MQSC commands.

Different groups of users can have different access authority to the same object. For example, for a specific queue, one group could perform both put and get operations; another group might be allowed only to browse the queue (MQGET with browse option). Similarly, some groups might have get and put authority to a queue, but not be allowed to alter or delete the queue.

WebSphere MQ for IBM i commands and perform operations on WebSphere MQ for IBM i objects

WebSphere MQ authorities on IBM i

To access WebSphere MQ objects, you need authority to issue the command and to access the object referenced. Administrators have access to all WebSphere MQ resources.

Access to WebSphere MQ objects is controlled by authorities to:

1. Issue the WebSphere MQ command
2. Access the WebSphere MQ objects referenced by the command

All WebSphere MQ for IBM i CL commands are shipped with an owner of QMQM, and the administration profile (QMQMADM) has *USE rights with the *PUBLIC access set to *EXCLUDE.

Note: The QSRDUPER program is used by the WebSphere MQ for IBM i licensed program install to duplicate Command (*CMD) objects in QSYS. In IBM i V5R4 and later, the QSRDUPER program was changed so that the default behavior is to create a proxy command rather than a duplicate of the original command. A proxy command redirects command execution to another command and has an attribute of PRX. If a proxy command by the same name as the command being copied exists in library QSYS, private authorities to the proxy command are not granted to the command in the product library. Attempts to prompt or run the proxy command in QSYS check the authority of the target command in the product library. Any changes in authority to *CMD objects therefore need to be done in the product library (QMQM) and those in QSYS do not need to be modified. For example:

```
GRTOBJAUT OBJ(QMQM/DSPMQM) OBJTYPE(*CMD) USER(MQUSER) AUT(*USE)
```

Changes to the authority structure of some of the product's CL commands allows public use of these commands, if you have the required OAM authority to the WebSphere MQ objects to make these changes.

To be a WebSphere MQ administrator on IBM i, you must be a member of the *QMQMADM* group. This group has properties like the properties of the mqm group on UNIX, Linux and Windows systems. In particular, the QMQMADM group is created when you install WebSphere MQ for IBM i, and members of the QMQMADM group have access to all WebSphere MQ resources on the system. You also have access to all WebSphere MQ resources if you have *ALLOBJ authority.

Administrators can use CL commands to administer WebSphere MQ. One of these commands is GRTMQMAUT, which is used to grant authorities to other users. Another command, STRMQMMQSC, enables an administrator to issue MQSC commands to a local queue manager.

Access authorities for IBM WebSphere MQ objects on IBM i:

Access authorities required for running IBM WebSphere MQ CL commands.

IBM WebSphere MQ for IBM i categorizes the product's CL commands into two groups:

Group 1

Users must be in the QMQMADM user group, or have *ALLOBJ authority, to process these commands. Users having either of these authorities can process all commands in all categories without requiring any extra authority.

Note: These authorities override any OAM authority.

These commands can be grouped as follows:

- Command Server Commands
 - ENDMQMCSVR, End IBM WebSphere MQ Command Server
 - STRMQMCSVR, Start IBM WebSphere MQ Command Server
- Dead-Letter Queue Handler Command
 - STRMQMDLQ, Start IBM WebSphere MQ Dead-Letter Queue Handler

- Listener Command
 - ENDMQMLSR, End IBM WebSphere MQ listener
 - STRMQMLSR, Start non-object listener
- Media Recovery Commands
 - RCDMQMIMG, Record IBM WebSphere MQ Object Image
 - RCRMQMOBJ, Re-create IBM WebSphere MQ Object
 - WRKMQMTRN, Work with IBM WebSphere MQ Transactions
- Queue Manager Commands
 - CRTMQM, Create Message Queue Manager
 - DLTMQM, Delete Message Queue Manager
 - ENDMQM, End Message Queue Manager
 - STRMQM, Start Message Queue Manager
- Security Commands
 - GRMQMAUT, Grant IBM WebSphere MQ Object Authority
 - RVKMQMAUT, Revoke IBM WebSphere MQ Object Authority
- Trace Command
 - TRCMQM, Trace IBM WebSphere MQ Job
- Transaction Commands
 - RSVMQMTRN, Resolve IBM WebSphere MQ Transaction
- Trigger Monitor Commands
 - STRMQMTRM, Start Trigger Monitor
- WebSphere MQSC Commands
 - RUNMQSC, Run WebSphere MQSC Commands
 - STRMQMMQSC, Start WebSphere MQSC Commands

Group 2

The rest of the commands, for which two levels of authority are required:

1. IBM i authority to run the command. A IBM WebSphere MQ administrator sets this using the **GRTOBJAUT** command to override the *PUBLIC(*EXCLUDE) restriction for a user or group of users.

For example:

```
GRTOBJAUT OBJ(DSPMQMQ) OBJTYPE(*CMD) USER(MQUSER) AUT(*USE)
```

2. IBM WebSphere MQ authority to manipulate the IBM WebSphere MQ objects associated with the command, or commands, given the correct IBM i authority in Step 1.

This authority is controlled by the user having the appropriate OAM authority for the required action, set by a IBM WebSphere MQ administrator using the **GRMQMAUT** command

For example:

```
GRMQMAUT *connect authority to the queue manager + *admchg authority to
the queue
```

The commands can be grouped as follows:

- Channel Commands
 - CHGMQMCHL, Change WebSphere MQ Channel
 - This requires *connect authority to the queue manager and *admchg authority to the channel.
 - CPYMQMCHL, Copy WebSphere MQ Channel

This requires *connect and *admcrtr authority to the queue manager, *admdsp authority to the default channel type to be copied, and *admcrtr authority to the channel object class. For example, copying a Sender channel, needs *admdsp authority to SYSTEM.DEF.SENDER channel

- CRTMQMCHL, Create WebSphere MQ Channel

This requires *connect and *admcrtr authority to the queue manager, *admdsp authority to the default channel type to be created and *admcrtr authority to the channel object class.

For example, creating a Sender channel, needs *admdsp authority to SYSTEM.DEF.SENDER channel

- DLTMQMCHL, Delete WebSphere MQ Channel

This requires *connect authority to the queue manager and *admdlt authority to the channel.

- RSVMQMCHL, Resolve WebSphere MQ Channel

This requires *connect authority to the queue manager and *ctrlx authority to the channel.

- Display commands

To process the DSP commands you must grant the user *connect and *admdsp authority to the queue manager, together with any specific option listed:

- DSPMQM, Display Message Queue Manager
- DSPMQMAUT, Display IBM WebSphere MQ Object Authority
- DSPMQMAUTI, Display IBM WebSphere MQ Authentication Information – *admdsp to the authentication information object
- DSPMQMCHL, Display IBM WebSphere MQ Channel – *admdsp to the channel
- DSPMQMCSVR, Display IBM WebSphere MQ Command Server
- DSPMQMNL, Display IBM WebSphere MQ Namelist – *admdsp to the namelist
- DSPMQMOBJN, Display IBM WebSphere MQ Object Names
- DSPMQMPRC, Display IBM WebSphere MQ Process – *admdsp to the process
- DSPMQMQ, Display IBM WebSphere MQ Queue – *admdsp to the queue
- DSPMQMTOP, Display IBM WebSphere MQ Topic – *admdsp to the topic

- Work with commands

To process the WRK commands and display the options panel you must grant the user *connect and *admdsp authority to the queue manager, together with any specific option listed:

- WRKMQM, Work with Message Queue Managers
- WRKMQMAUT, Work with IBM WebSphere MQ Object Authority
- WRKMQMAUTD, Work with IBM WebSphere MQ Object Authority Data
- WRKMQMAUTI, Work with IBM WebSphere MQ Authentication Information
 - *admchg for the Change IBM WebSphere MQ Authentication Information Object command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Authentication Information Object command.
 - *admdlt for the Delete IBM WebSphere MQ Authentication Information Object command.
 - *admdsp for the Display IBM WebSphere MQ Authentication Information Object command.
- WRKMQMCHL, Work with IBM WebSphere MQ Channel

This requires the following authorities:

 - *admchg for the Change IBM WebSphere MQ Channel command.
 - *admclr for the Clear IBM WebSphere MQ Channel command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Channel command.
 - *admdlt for the Delete IBM WebSphere MQ Channel command.
 - *admdsp for the Display IBM WebSphere MQ Channel command.

- *ctrl for the Start IBM WebSphere MQ Channel command.
- *ctrl for the End IBM WebSphere MQ Channel command.
- *ctrl for the Ping IBM WebSphere MQ Channel command.
- *ctrlx for the Reset IBM WebSphere MQ Channel command.
- *ctrlx for the Resolve IBM WebSphere MQ Channel command.
- WRKMQMCHST, Work with IBM WebSphere MQ Channel Status
This requires *admdsp authority to the channel.
- WRKMQMCL, Work with IBM WebSphere MQ Clusters
- WRKMQMCLQ, Work with IBM WebSphere MQ Cluster Queues
- WRKMQMCLQM, Work with IBM WebSphere MQ Cluster Queue Manager
- WRKMQMLSR, Work with IBM WebSphere MQ Listener
- WRKMQMMSG, Work with IBM WebSphere MQ Messages
This requires *browse authority to the queue
- WRKMQMNL, Work with IBM WebSphere MQ Namelists
This requires the following authorities:
 - *admchg for the Change IBM WebSphere MQ Namelist command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Namelist command.
 - *admdlt for the Delete IBM WebSphere MQ Namelist command.
 - *admdsp for the Display IBM WebSphere MQ Namelist command.
- WRKMQMPRC, Work with IBM WebSphere MQ Processes
This requires the following authorities:
 - *admchg for the Change IBM WebSphere MQ Process command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Process command.
 - *admdlt for the Delete IBM WebSphere MQ Process command.
 - *admdsp for the Display IBM WebSphere MQ Process command.
- WRKMQMQ, Work with IBM WebSphere MQ queues
This requires the following authorities:
 - *admchg for the Change IBM WebSphere MQ Queue command.
 - *admcldr for the Clear IBM WebSphere MQ Queue command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Queue command.
 - *admdlt for the Delete IBM WebSphere MQ Queue command.
 - *admdsp for the Display IBM WebSphere MQ Queue command.
- WRKMQMQSTS, Work with IBM WebSphere MQ Queue Status
- WRKMQMTOP, Work with IBM WebSphere MQ Topics
This requires the following authorities
 - *admchg for the Change IBM WebSphere MQ Topic command.
 - *admcrtr for the Create and Copy IBM WebSphere MQ Topic command.
 - *admdlt for the Delete IBM WebSphere MQ Topic command.
 - *admdsp for the Display IBM WebSphere MQ Topic command.
- WRKMQMSUB, Work with IBM WebSphere MQ Subscriptions
- Other Channel commands
To process the channel commands you must grant the user the specific authorities listed:
 - ENDMQMCHL, End IBM WebSphere MQ Channel
This requires *connect authority to the queue manager and *allmqi authority to the transmission queue associated with the channel.

- ENDMQMLSR, End IBM WebSphere MQ Listener
This requires *connect authority to the queue manager and *ctrl authority to the named listener object.
- PNGMQMCHL, Ping IBM WebSphere MQ Channel
This requires *connect and *inq authority to the queue manager and *ctrl authority to the channel object.
- RSTMQMCHL, Reset IBM WebSphere MQ Channel
This requires *connect authority to the queue manager.
- STRMQMCHL, Start IBM WebSphere MQ Channel
This requires *connect authority to the queue manager and *ctrl authority to the channel object.
- STRMQMCHLI, Start IBM WebSphere MQ Channel Initiator
This requires *connect and *inq authority to the queue manager, and *allmqi authority to the initiation queue associated with the transmission queue of the channel.
- STRMQMLSR, Start IBM WebSphere MQ Listener
This requires *connect authority to the queue manager and *ctrl authority to the named listener object.
- Other commands:
To process the following commands you must grant the user the specific authorities listed:
 - CCTMQM, Connect to Message Queue Manager
This requires no IBM WebSphere MQ object authority.
 - CHGMQM, Change Message Queue Manager
This requires *connect and *admchg authority to the queue manager.
 - CHGMQMAUTI, Change IBM WebSphere MQ Authentication Information
This requires *connect authority to the queue manager and *admchg and *admdsp authority to the authentication information object.
 - CHGMQMNL, Change IBM WebSphere MQ Namelist
This requires *connect authority to the queue manager and *admchg authority to the namelist.
 - CHGMQMPRC, Change IBM WebSphere MQ Process
This requires *connect authority to the queue manager and *admchg authority to the process.
 - CHGMQMQ, Change IBM WebSphere MQ Queue
This requires *connect authority to the queue manager and *admchg authority to the queue.
 - CLRMQMQ, Clear IBM WebSphere MQ Queue
This requires *connect authority to the queue manager and *admclr authority to the queue.
 - CPYMQMAUTI, Copy IBM WebSphere MQ Authentication Information
This requires *connect authority to the queue manager and *admdsp authority to the authentication information object and *admcrta authority to the authentication information object class.
 - CPYMQMNL, Copy IBM WebSphere MQ Namelist
This requires *connect and *admcrta authority to the queue manager.
 - CPYMQMPRC, Copy IBM WebSphere MQ Process
This requires *connect and *admcrta authority to the queue manager.
 - CPYMQMQ, Copy IBM WebSphere MQ Queue
This requires *connect and *admcrta authority to the queue manager.
 - CRTMQMAUTI, Create IBM WebSphere MQ Authentication Information

- This requires *connect authority to the queue manager and *admdsp authority to the authentication information object and *admcrt authority to the authentication information object class.
- CRTMQMNL, Create IBM WebSphere MQ Namelist
This requires *connect and *admcrt authority to the queue manager and *admdsp authority to the default namelist.
 - CRTMQMPRC, Create IBM WebSphere MQ Process
This requires *connect and *admcrt authority to the queue manager and *admdsp authority to the default process.
 - CRTMQMQ, Create IBM WebSphere MQ Queue
This requires *connect and *admcrt authority to the queue manager and *admdsp authority to the default queue.
 - CVTMQMDDTA, Convert IBM WebSphere MQ Data Type Command
This requires no IBM WebSphere MQ object authority.
 - DLTMQMAUTI, Delete IBM WebSphere MQ Authentication Information
This requires *connect authority to the queue manager and *ctrlx authority to the authentication information object.
 - DLTMQMNL, Delete IBM WebSphere MQ Namelist
This requires *connect authority to the queue manager and *admdlt authority to the namelist.
 - DLTMQMPRC, Delete IBM WebSphere MQ Process
This requires *connect authority to the queue manager and *admdlt authority to the process.
 - DLTMQMQ, Delete IBM WebSphere MQ Queue
This requires *connect authority to the queue manager and *admdlt authority to the queue.
 - DSCMQM, Disconnect from Message Queue Manager
This requires no IBM WebSphere MQ object authority.
 - RFRMQMAUT, Refresh Security
This requires *connect authority to the queue manager.
 - RFRMQMCL, Refresh Cluster
This requires *connect authority to the queue manager.
 - RSMMQMCLQM, Resume Cluster Queue Manager
This requires *connect authority to the queue manager.
 - RSTMQMCL, Reset Cluster
This requires *connect authority to the queue manager.
 - SPDMQMCLQM, Suspend Cluster Queue Manager
This requires *connect authority to the queue manager.

Access authorizations on IBM i:

Use this information to understand the access authorization commands.

Authorizations defined by the AUT keyword on the **GRTMQMAUT** and **RVKMQMAUT** commands can be categorized as follows:

- Authorizations related to MQI calls
- Authorization-related administration commands
- Context authorizations
- General authorizations, that is, for MQI calls, for commands, or both

The following tables list the different authorities, using the AUT parameter for MQI calls, Context calls, MQSC and PCF commands, and generic operations.

Table 22. Authorizations for MQI calls

AUT	Description
*ALTUSR	Allow another user's authority to be used for MQOPEN and MQPUT1 calls.
*BROWSE	Retrieve a message from a queue by issuing an MQGET call with the BROWSE option.
*CONNECT	Connect the application to the specified queue manager by issuing an MQCONN call.
*GET	Retrieve a message from a queue by issuing an MQGET call.
*INQ	Make an inquiry on a specific queue by issuing an MQINQ call.
*PUB	Open a topic to publish a message using an MQPUT call.
*PUT	Put a message on a specific queue by issuing an MQPUT call.
*RESUME	Resume a subscription using an MQSUB call.
*SET	Set attributes on a queue from the MQI by issuing an MQSET call. If you open a queue for multiple options, you must be authorized for each of them.
*SUB	Create, Alter or Resume a subscription to a topic using an MQSUB call.

Table 23. Authorizations for context calls

AUT	Description
*PASSALL	Pass all context on the specified queue. All the context fields are copied from the original request.
*PASSID	Pass identity context on the specified queue. The identity context is the same as that of the request.
*SETALL	Set all context on the specified queue. This is used by special system utilities.
*SETID	Set identity context on the specified queue. This is used by special system utilities.

Table 24. Authorizations for MQSC and PCF calls

AUT	Description
*ADMCHG	Change the attributes of the specified object.
*ADMCLR	Clear the specified object (PCF Clear object command only).
*ADMCRRT	Create objects of the specified type.
*ADMDLT	Delete the specified object.
*ADMDSPP	Display the attributes of the specified object.

Table 25. Authorizations for generic operations

AUT	Description
*ALL	Use all operations applicable to the object. all authority is equivalent to the union of the authorities alladm, allmqi, and system appropriate to the object type.
*ALLADM	Perform all administration operations applicable to the object.
*ALLMQI	Use all MQI calls applicable to the object.
*CTRL	Control startup and shutdown of channels, listeners, and services.
*CTRLX	Reset sequence number and resolve indoubt channels.

Using the access authorization commands on IBM i:

Use this information to learn about the access authorization commands, and use the command examples.

Using the GRTMQMAUT command

If you have the required authorization, you can use the **GRTMQMAUT** command to grant authorization of a user profile or user group to access a particular object. The following examples illustrate how the **GRTMQMAUT** command is used:

1.

```
GRTMQMAUT OBJ(RED.LOCAL.QUEUE) OBJTYPE(*LCLQ) USER(GROUPA) +  
AUT(*BROWSE *PUT) MQMNAME('saturn.queue.manager')
```

In this example:

- RED.LOCAL.QUEUE is the object name.
 - *LCLQ (local queue) is the object type.
 - GROUPA is the name of a user profile on the system for which authorizations are to change. This profile can be used as a group profile for other users.
 - *BROWSE and *PUT are the authorizations being granted to the specified queue.
 - *BROWSE adds authorization to browse messages on the queue (to issue MQGET with the browse option).
 - *PUT adds authorization to put (MQPUT) messages on the queue.
 - saturn.queue.manager is the queue manager name.
2. The following command grants to users JACK and JILL all applicable authorizations, to all process definitions, for the default queue manager.

```
GRTMQMAUT OBJ(*ALL) OBJTYPE(*PRC) USER(JACK JILL) AUT(*ALL)
```

3. The following command grants user GEORGE authority to put a message on the queue ORDERS, on the queue manager TRENT.

```
GRTMQMAUT OBJ(TRENT) OBJTYPE(*MQM) USER(GEORGE) AUT(*CONNECT) MQMNAME (TRENT)  
GRTMQMAUT OBJ(ORDERS) OBJTYPE(*Q) USER(GEORGE) AUT(*PUT) MQMNAME (TRENT)
```

Using the RVKMQMAUT command

If you have the required authorization, you can use the **RVKMQMAUT** command to remove previously granted authorization of a user profile or user group to access a particular object. The following examples illustrate how the **RVKMQMAUT** command is used:

1.

```
RVKMQMAUT OBJ(RED.LOCAL.QUEUE) OBJTYPE(*LCLQ) USER(GROUPA) +  
AUT(*PUT) MQMNAME('saturn.queue.manager')
```

The authority to put messages to the specified queue, that was granted in the previous example, is removed for GROUPA.

2.

```
RVKMQMAUT OBJ(PAY*) OBJTYPE(*Q) USER(*PUBLIC) AUT(*GET) +  
MQMNAME (PAYROLLQM)
```

Authority to get messages from any queue with a name starting with the characters PAY, owned by queue manager PAYROLLQM, is removed from all users of the system unless they, or a group to which they belong, have been separately authorized.

Using the DSPMQMAUT command

The display MQM authority (**DSPMQMAUT**) command shows, for the specified object and user, the list of authorizations that the user has for the object. The following example illustrates how the command is used:

```
DSPMQMAUT OBJ(ADMINNL) OBJTYPE(*NMLIST) USER(JOE) OUTPUT(*PRINT) +
MQMNAME(ADMINQM)
```

Using the RFRMQMAUT command

The refresh MQM security (**RFRMQMAUT**) command enables you to update the OAM's authorization group information immediately, reflecting changes made at the operating system level, without needing to stop and restart the queue manager. The following example illustrates how the command is used:

```
RFRMQMAUT MQMNAME(ADMINQM)
```

Authorization specification tables for IBM i

Use this information to determine what authorization is required to use particular API calls, and particular options of those calls, on queue objects, process objects, and queue manager objects.

The authorization specification tables starting in Table 26 on page 503 define precisely how the authorizations work and the restrictions that apply. The tables apply to these situations:

- Applications that issue MQI calls
- Administration programs that issue MQSC commands as escape PCFs
- Administration programs that issue PCF commands

In this section, the information is presented as a set of tables that specify the following data:

Action to be performed

MQI option, MQSC command, or PCF command.

Access control object

Queue, process definition, queue manager, namelist, channel, client connection channel, listener, service, or authentication information object.

Authorization required

Expressed as an MQZAO_ constant.

In the tables, the constants prefixed by MQZAO_ correspond to the keywords in the authorization list for the **GRTMQMAUT** and **RVKMQMAUT** commands for the particular entity. For example, MQZAO_BROWSE corresponds to the keyword *BROWSE; similarly, the keyword MQZAO_SET_ALL_CONTEXT corresponds to the keyword *SETALL, and so on. These constants are defined in the header file cmqzc.h, which is supplied with the product.

MQI authorizations

An application is allowed to issue specific MQI calls and options only if the user identifier under which it is running (or whose authorizations it is able to assume) has been granted the relevant authorization.

Four MQI calls require authorization checks: MQCONN, MQOPEN, MQPUT1, and MQCLOSE.

For MQOPEN and MQPUT1, the authority check is made on the name of the object being opened, and not on the name, or names, resulting after a name has been resolved. For example, an application can be granted authority to open an alias queue without having authority to open the base queue to which the alias resolves. The rule is that the check is carried out on the first definition encountered during the process of name resolution that is not a queue-manager alias, unless the queue-manager alias definition is opened directly; that is, its name appears in the *ObjectName* field of the object descriptor. Authority is

always needed for the particular object being opened; in some cases additional queue-independent authority, obtained through an authorization for the queue-manager object, is required.

Table 26, Table 27, Table 28 on page 504, and Table 29 on page 504 summarize the authorizations needed for each call.

Note: These tables do not mention namelists, channels, client connection channels, listeners, services, or authentication information objects. This is because none of the authorizations apply to these objects, except for MQOO_INQUIRE, for which the same authorizations apply as for the other objects.

Table 26. Security authorization needed for MQCONN calls

Authorization required for:	Queue object (1 on page 504)	Process object	Queue manager object
MQCONN option	Not applicable	Not applicable	MQZAO_CONNECT

Table 27. Security authorization needed for MQOPEN calls

Authorization required for:	Queue object (1 on page 504)	Process object	Queue manager object
MQOO_INQUIRE	MQZAO_INQUIRE (2 on page 504)	MQZAO_INQUIRE (2 on page 504)	MQZAO_INQUIRE (2 on page 504)
MQOO_BROWSE	MQZAO_BROWSE	Not applicable	No check
MQOO_INPUT_*	MQZAO_INPUT	Not applicable	No check
MQOO_SAVE_ALL_CONTEXT (3 on page 504)	MQZAO_INPUT	Not applicable	Not applicable
MQOO_OUTPUT (Normal queue) (4 on page 504)	MQZAO_OUTPUT	Not applicable	Not applicable
MQOO_PASS_IDENTITY_CONTEXT (5 on page 504)	MQZAO_PASS_IDENTITY_CONTEXT	Not applicable	No check
MQOO_PASS_ALL_CONTEXT (5 on page 504, 6 on page 504)	MQZAO_PASS_ALL_CONTEXT	Not applicable	No check
MQOO_SET_IDENTITY_CONTEXT (5 on page 504, 6 on page 504)	MQZAO_SET_IDENTITY_CONTEXT	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (7 on page 504)
MQOO_SET_ALL_CONTEXT (5 on page 504, 8 on page 504)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7 on page 504)
MQOO_OUTPUT (Transmission queue) (9 on page 504)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7 on page 504)
MQOO_SET	MQZAO_SET	Not applicable	No check
MQOO_ALTERNATE_USER_AUTHORITY	(10 on page 504)	(10 on page 504)	MQZAO_ALTERNATE_USER_AUTHORITY (10 on page 504, 11 on page 505)

Table 28. Security authorization needed for MQPUT1 calls

Authorization required for:	Queue object (1)	Process object	Queue manager object
MQPMO_PASS_IDENTITY_CONTEXT	MQZAO_PASS_IDENTITY_CONTEXT (12 on page 505)	Not applicable	No check
MQPMO_PASS_ALL_CONTEXT	MQZAO_PASS_ALL_CONTEXT (12 on page 505)	Not applicable	No check
MQPMO_SET_IDENTITY_CONTEXT	MQZAO_SET_IDENTITY_CONTEXT (12 on page 505)	Not applicable	MQZAO_SET_IDENTITY_CONTEXT (7)
MQPMO_SET_ALL_CONTEXT	MQZAO_SET_ALL_CONTEXT (12 on page 505)	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
(Transmission queue) (9)	MQZAO_SET_ALL_CONTEXT	Not applicable	MQZAO_SET_ALL_CONTEXT (7)
MQPMO_ALTERNATE_USER_AUTHORITY	(13 on page 505)	Not applicable	MQZAO_ALTERNATE_USER_AUTHORITY (11 on page 505)

Table 29. Security authorization needed for MQCLOSE calls

Authorization required for:	Queue object (1)	Process object	Queue manager object
MQCO_DELETE	MQZAO_DELETE (14 on page 505)	Not applicable	Not applicable
MQCO_DELETE_PURGE	MQZAO_DELETE (14 on page 505)	Not applicable	Not applicable

Notes for the tables:

1. If a model queue is being opened:
 - MQZAO_DISPLAY authority is needed for the model queue, in addition to the authority to open the model queue for the type of access for which you are opening.
 - MQZAO_CREATE authority is not needed to create the dynamic queue.
 - The user identifier used to open the model queue is automatically granted all the queue-specific authorities (equivalent to MQZAO_ALL) for the dynamic queue created.
2. Either the queue, process, namelist, or queue manager object is checked, depending on the type of object being opened.
3. MQOO_INPUT_* must also be specified. This option is valid for a local, model, or alias queue.
4. This check is performed for all output cases, except the case specified in note 9.
5. MQOO_OUTPUT must also be specified.
6. MQOO_PASS_IDENTITY_CONTEXT is also implied by this option.
7. This authority is required for both the queue manager object and the particular queue.
8. MQOO_PASS_IDENTITY_CONTEXT, MQOO_PASS_ALL_CONTEXT, and MQOO_SET_IDENTITY_CONTEXT are also implied by this option.
9. This check is performed for a local or model queue that has a *Usage* queue attribute of MQUS_TRANSMISSION, and is being opened directly for output. It does not apply if a remote queue is being opened (either by specifying the names of the remote queue manager and remote queue, or by specifying the name of a local definition of the remote queue).
10. At least one of MQOO_INQUIRE (for any object type), or (for queues) MQOO_BROWSE, MQOO_INPUT_*, MQOO_OUTPUT, or MQOO_SET must also be specified. The check carried out is

as for the other options specified, using the supplied alternate-user identifier for the specific-named object authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.

11. This authorization allows any *AlternateUserId* to be specified.
12. An MQZAO_OUTPUT check is also carried out if the queue does not have a *Usage* queue attribute of MQUS_TRANSMISSION.
13. The check carried out is as for the other options specified, using the supplied alternate-user identifier for the named queue authority, and the current application authority for the MQZAO_ALTERNATE_USER_IDENTIFIER check.
14. The check is carried out only if both of the following are true:
 - A permanent dynamic queue is being closed and deleted.
 - The queue was not created by the MQOPEN that returned the object handle being used.Otherwise, there is no check.

General notes:

1. The special authorization MQZAO_ALL_MQI includes all the following authorizations that are relevant to the object type:
 - MQZAO_CONNECT
 - MQZAO_INQUIRE
 - MQZAO_SET
 - MQZAO_BROWSE
 - MQZAO_INPUT
 - MQZAO_OUTPUT
 - MQZAO_PASS_IDENTITY_CONTEXT
 - MQZAO_PASS_ALL_CONTEXT
 - MQZAO_SET_IDENTITY_CONTEXT
 - MQZAO_SET_ALL_CONTEXT
 - MQZAO_ALTERNATE_USER_AUTHORITY
2. MQZAO_DELETE (see note 14) and MQZAO_DISPLAY are classed as administration authorizations. They are not therefore included in MQZAO_ALL_MQI.
3. *No check* means that no authorization checking is carried out.
4. *Not applicable* means that authorization checking is not relevant to this operation. For example, you cannot issue an MQPUT call to a process object.

Authorizations for MQSC commands in escape PCFs on IBM i:

These authorizations allow a user to issue administration commands as an escape PCF message. These methods allow a program to send an administration command as a message to a queue manager, for execution on behalf of that user.

This section summarizes the authorizations needed for each MQSC command contained in Escape PCF.

Not applicable means that authorization checking is not relevant to this operation.

The user ID under which the program that submits the command is running must also have the following authorities:

- MQZAO_CONNECT authority to the queue manager
- DISPLAY authority on the queue manager in order to perform PCF commands
- Authority to issue the MQSC commands within the text of the Escape PCF command

ALTER object

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	MQZAO_CHANGE
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

CLEAR object

Object	Authorization required
Queue	MQZAO_CLEAR
Topic	MQZAO_CLEAR
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	Not applicable
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

DEFINE object NOREPLACE (1 on page 509)

Object	Authorization required
Queue	MQZAO_CREATE (2 on page 509)
Topic	MQZAO_CREATE (2 on page 509)
Process	MQZAO_CREATE (2 on page 509)
Queue manager	Not applicable
Namelist	MQZAO_CREATE (2 on page 509)
Authentication information	MQZAO_CREATE (2 on page 509)
Channel	MQZAO_CREATE (2 on page 509)
Client connection channel	MQZAO_CREATE (2 on page 509)
Listener	MQZAO_CREATE (2 on page 509)
Service	MQZAO_CREATE (2 on page 509)

DEFINE object REPLACE (1 on page 509, 3 on page 509)

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	Not applicable
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

DELETE *object*

Object	Authorization required
Queue	MQZAO_DELETE
Topic	MQZAO_DELETE
Process	MQZAO_DELETE
Queue manager	Not applicable
Namelist	MQZAO_DELETE
Authentication information	MQZAO_DELETE
Channel	MQZAO_DELETE
Client connection channel	MQZAO_DELETE
Listener	MQZAO_DELETE
Service	MQZAO_DELETE

DISPLAY *object*

Object	Authorization required
Queue	MQZAO_DISPLAY
Topic	MQZAO_DISPLAY
Process	MQZAO_DISPLAY
Queue manager	MQZAO_DISPLAY
Namelist	MQZAO_DISPLAY
Authentication information	MQZAO_DISPLAY
Channel	MQZAO_DISPLAY
Client connection channel	MQZAO_DISPLAY
Listener	
Service	

PING CHANNEL

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

RESET CHANNEL

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL_EXTENDED
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

RESOLVE CHANNEL

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL_EXTENDED
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

START *object*

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	MQZAO_CONTROL
Service	MQZAO_CONTROL

STOP *object*

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	MQZAO_CONTROL
Service	MQZAO_CONTROL

Note:

1. For DEFINE commands, MQZAO_DISPLAY authority is also needed for the LIKE object if one is specified, or on the appropriate SYSTEM.DEFAULT.xxx object if LIKE is omitted.
2. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects for a specified queue manager, by specifying an object type of QMGR on the **GRTMQMAUT** command.
3. This option applies if the object to be replaced already exists. If it does not, the check is as for DEFINE *object* NOREPLACE.

Authorizations for PCF commands on IBM i:

These authorizations allow a user to issue administration commands as PCF commands. These methods allow a program to send an administration command as a message to a queue manager, for execution on behalf of that user.

This section summarizes the authorizations needed for each PCF command.

No check means that no authorization checking is carried out; *Not applicable* means that authorization checking is not relevant to this operation.

The user ID under which the program that submits the command is running must also have the following authorities:

- MQZAO_CONNECT authority to the queue manager
- DISPLAY authority on the queue manager in order to perform PCF commands

The special authorization MQZAO_ALL_ADMIN includes the following authorizations:

- MQZAO_CHANGE
- MQZAO_CLEAR
- MQZAO_DELETE
- MQZAO_DISPLAY
- MQZAO_CONTROL
- MQZAO_CONTROL_EXTENDED

MQZAO_CREATE is not included as it is not specific to a particular object or object type

Change object

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	MQZAO_CHANGE
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

Clear object

Object	Authorization required
Queue	MQZAO_CLEAR
Topic	MQZAO_CLEAR
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	Not applicable
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Copy object (without replace) (1 on page 515)

Object	Authorization required
Queue	MQZAO_CREATE (2 on page 515)
Topic	MQZAO_CREATE (2 on page 515)
Process	MQZAO_CREATE (2 on page 515)
Queue manager	Not applicable
NamelistMQZAO_CREATE	MQZAO_CREATE (2 on page 515)
Authentication information	MQZAO_CREATE (2 on page 515)
Channel	MQZAO_CREATE (2 on page 515)
Client connection channel	MQZAO_CREATE (2 on page 515)
Listener	MQZAO_CREATE (2 on page 515)
Service	MQZAO_CREATE (2 on page 515)

Copy object (with replace) (1 on page 515, 4 on page 515)

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	Not applicable
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

Create object (without replace) (3 on page 515)

Object	Authorization required
Queue	MQZAO_CREATE (2 on page 515)
Topic	MQZAO_CREATE (2 on page 515)
Process	MQZAO_CREATE (2 on page 515)
Queue manager	Not applicable
Namelist	MQZAO_CREATE (2 on page 515)
Authentication information	MQZAO_CREATE (2 on page 515)
Channel	MQZAO_CREATE (2 on page 515)
Client connection channel	MQZAO_CREATE (2 on page 515)
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

Create object (with replace) (3 on page 515, 4 on page 515)

Object	Authorization required
Queue	MQZAO_CHANGE
Topic	MQZAO_CHANGE
Process	MQZAO_CHANGE
Queue manager	Not applicable
Namelist	MQZAO_CHANGE
Authentication information	MQZAO_CHANGE
Channel	MQZAO_CHANGE
Client connection channel	MQZAO_CHANGE
Listener	MQZAO_CHANGE
Service	MQZAO_CHANGE

Delete object

Object	Authorization required
Queue	MQZAO_DELETE
Topic	MQZAO_DELETE
Process	MQZAO_DELETE
Queue manager	MQZAO_DELETE
Namelist	MQZAO_DELETE
Authentication information	MQZAO_DELETE
Channel	MQZAO_DELETE
Client connection channel	MQZAO_DELETE
Listener	MQZAO_DELETE
Service	MQZAO_DELETE

Inquire object

Object	Authorization required
Queue	MQZAO_DISPLAY
Topic	MQZAO_DISPLAY
Process	MQZAO_DISPLAY
Queue manager	MQZAO_DISPLAY
Namelist	MQZAO_DISPLAY
Authentication information	MQZAO_DISPLAY
Channel	MQZAO_DISPLAY
Client connection channel	MQZAO_DISPLAY
Listener	MQZAO_DISPLAY
Service	MQZAO_DISPLAY

Inquire object names

Object	Authorization required
Queue	No check
Topic	No check
Process	No check
Queue manager	No check
Namelist	No check
Authentication information	No check
Channel	No check
Client connection channel	No check
Listener	No check
Service	No check

Ping Channel

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Reset Channel

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL_EXTENDED
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Reset Queue Statistics

Object	Authorization required
Queue	MQZAO_DISPLAY and MQZAO_CHANGE
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	Not applicable
Client connection channel	Not applicable
Listener	
Service	

Resolve Channel

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL_EXTENDED
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Start Channel

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Stop Channel

Object	Authorization required
Queue	Not applicable
Topic	Not applicable
Process	Not applicable
Queue manager	Not applicable
Namelist	Not applicable
Authentication information	Not applicable
Channel	MQZAO_CONTROL
Client connection channel	Not applicable
Listener	Not applicable
Service	Not applicable

Note:

1. For Copy commands, MQZAO_DISPLAY authority is also needed for the From object.
2. The MQZAO_CREATE authority is not specific to a particular object or object type. Create authority is granted for all objects for a specified queue manager, by specifying an object type of QMGR on the **GRTMQMAUT** command.
3. For Create commands, MQZAO_DISPLAY authority is also needed for the appropriate SYSTEM.DEFAULT.* object.
4. This option applies if the object to be replaced already exists. If it does not, the check is as for Copy or Create without replace.

Generic OAM profiles on IBM i

Object authority manager (OAM) generic profiles enable you to set the authority a user has to many objects at once, rather than having to issue separate **GRTMQMAUT** commands against each individual object when it is created. Using generic profiles in the **GRTMQMAUT** command enables you to set a generic authority for all future objects created that fit that profile.

The rest of this section describes the use of generic profiles in more detail:

- “Using wildcard characters”
- “Profile priorities” on page 516

Using wildcard characters

What makes a profile generic is the use of special characters (wildcard characters) in the profile name. For example, the question mark (?) wildcard character matches any single character in a name. So, if you specify ABC.?EF, the authorization you give to that profile applies to any objects created with the names ABC.DEF, ABC.CEF, ABC.BEF, and so on.

The wildcard characters available are:

- ? Use the question mark (?) instead of any single character. For example, AB.?D would apply to the objects AB.CD, AB.ED, and AB.FD.
- * Use the asterisk (*) as:
 - A *qualifier* in a profile name to match any one qualifier in an object name. A qualifier is the part of an object name delimited by a period. For example, in ABC.DEF.GHI, the qualifiers are ABC, DEF, and GHI.
For example, ABC.*.JKL would apply to the objects ABC.DEF.JKL, and ABC.GHI.JKL. (Note that it would **not** apply to ABC.JKL; * used in this context always indicates one qualifier.)

- A character within a qualifier in a profile name to match zero or more characters within the qualifier in an object name.

For example, ABC.DE*.JKL would apply to the objects ABC.DE.JKL, ABC.DEF.JKL, and ABC.DEGH.JKL.

** Use the double asterisk (**) *once* in a profile name as:

- The entire profile name to match all object names. For example, if you use the keyword OBJTYPE (*PRC) to identify processes, then use ** as the profile name, you change the authorizations for all processes.
- As either the beginning, middle, or ending qualifier in a profile name to match zero or more qualifiers in an object name. For example, **.ABC identifies all objects with the final qualifier ABC.

Profile priorities

An important point to understand when using generic profiles is the priority that profiles are given when deciding what authorities to apply to an object being created. For example, suppose that you have issued the commands:

```
GRTMQMAUT OBJ(AB.*) OBJTYPE(*Q) USER(FRED) AUT(*PUT) MQMNAME(MYQMGR)
GRTMQMAUT OBJ(AB.C*) OBJTYPE(*Q) USER(FRED) AUT(*GET) MQMNAME(MYQMGR)
```

The first gives put authority to all queues for the principal FRED with names that match the profile AB.*; the second gives get authority to the same types of queue that match the profile AB.C*.

Suppose that you now create a queue called AB.CD. According to the rules for wildcard matching, either GRTMQMAUT could apply to that queue. So, does it have put or get authority?

To find the answer, you apply the rule that, whenever multiple profiles can apply to an object, **only the most specific applies**. The way that you apply this rule is by comparing the profile names from left to right. Wherever they differ, a non-generic character is more specific than a generic character. So, in the example above, the queue AB.CD has **get** authority (AB.C* is more specific than AB.*).

When you are comparing generic characters, the order of *specificity* is:

1. ?
2. *
3. **

Specifying the installed authorization service on IBM i

You can specify which authorization service component to use.

The parameter **Service Component name** on **GRTMQMAUT** and **RVKMQMAUT** allows you to specify the name of the installed authorization service component.

Selecting **F24** on the initial panel, followed by **F9=All parameters** on the next panel of either command, allows you to specify either the installed authorization component (*DFT) or the name of the required authorization service component specified in the Service stanza of the queue manager's qm.ini file.

DSPMQMAUT also has this extra parameter. This parameter allows you to search all the installed authorization components (*DFT), or the specified authorization-service component name, for the specified object name, object type, and user

Working with and without authority profiles on IBM i

Use this information to learn how to work with authority profiles and how to work without authority profiles.

You can work with authority profiles, as explained in “Working with authority profiles,” or without them, as explained here:

To work without authority profiles, use *NONE as an Authority parameter on **GRTMQMAUT** to create profiles without authority. This leaves any existing profiles unchanged.

On **RVKMQMAUT**, use *REMOVE as an Authority parameter to remove an existing authority profile.

Working with authority profiles

There are two commands associated with authority profiling:

- **WRKMQMAUT**
- **WRKMQMAUTD**

You can access these commands directly from the command line, or from the **WRKMQM** panel by:

1. Typing in the queue manager name and pressing the Enter key to access the **WRKMQM** results panel.
2. Selecting F23=More options on this panel.

Option 24 selects the results panel for the **WRKMQMAUT** command and option 25 selects the **WRKMQMAUTI** command, which is used with the SSL bindings layer.

WRKMQMAUT

This command allows you to work with the authority data held in the authority queue.

Note: To run this command you must have *connect and *admdsp authority to the queue manager. However, to create or delete a profile, you need QMQMADM authority.

If you output the information to the screen, a list of authority profile names, together with their types, is displayed. If you print the output, you receive a detailed list of all the authority data, the registered users, and their authorities.

Entering an object or profile name on this panel, and pressing ENTER takes you to the results panel for **WRKMQMAUT**.

If you select 4=Delete, you go to a new panel from which you can confirm that you want to delete all the user names registered to the generic authority profile name you specify. This option runs **RVKMQMAUT** with the option *REMOVE for all the users, and applies **only** to generic profile names.

If you select 12=Work with profile you go to the **WRKMQMAUTD** command results panel, as explained in “WRKMQMAUTD.”

WRKMQMAUTD

This command allows you to display all the users registered with a particular authority profile name and object type. To run this command you must have *connect and *admdsp authority to the queue manager. However, to grant, run, create, or delete a profile you need QMQMADM authority.

Selecting F24=More keys from the initial input panel, followed by option F9=All Parameters displays the Service Component Name as for **GRTMQMAUT** and **RVKMQMAUT**.

Note: The F11=Display Object Authorizations key toggles between the following types of authorities:

- Object authorizations
- Context authorizations
- MQI authorizations

The options on the screen are:

2=Grant

Takes you to the **GRTMQMAUT** panel to add to the current authorities.

3=Revoke

Takes you to the **RVKMQMAUT** panel to remove some of the current definitions

4=Delete

Takes you to a panel that allows you to delete the authority data for specified users. This runs **RVKMQMAUT** with the option *REMOVE.

5=Display

Takes you to the existing **DSPMQMAUT** command

F6=Create

Takes you to the **GRTMQMAUT** panel that allows you to create a profile authority record.

Object authority manager guidelines for IBM i

Additional hints and tips for using the object authority manager (OAM)

Limit access to sensitive operations

Some operations are sensitive; limit them to privileged users. For example,

- Accessing some special queues, such as transmission queues or the command queue
SYSTEM.ADMIN.COMMAND.QUEUE
- Running programs that use full MQI context options
- Creating and copying application queues

Queue manager directories

The directories and libraries containing queues and other queue manager data are private to the product. Do not use standard operating system commands to grant or revoke authorizations to MQI resources.

Queues

The authority to a dynamic queue is based on, but is not necessarily the same as, that of the model queue from which it is derived.

For alias queues and remote queues, the authorization is that of the object itself, not the queue to which the alias or remote queue resolves. It is possible to authorize a user profile to access an alias queue that resolves to a local queue to which the user profile has no access permissions.

Limit the authority to create queues to privileged users. If you do not, users can bypass the normal access control by creating an alias.

Alternate-user authority

Alternate-user authority controls whether one user profile can use the authority of another user profile when accessing a IBM WebSphere MQ object. This technique is essential where a server receives requests

from a program and the server wants to ensure that the program has the required authority for the request. The server might have the required authority, but it needs to know whether the program has the authority for the actions it has requested.

For example:

- A server program running under user profile PAYSERV retrieves a request message from a queue that was put on the queue by user profile USER1.
- When the server program gets the request message, it processes the request and puts the reply back into the reply-to queue specified with the request message.
- Instead of using its own user profile (PAYSERV) to authorize opening the reply-to queue, the server can specify some other user profile, in this case, USER1. In this example, you can use alternate-user authority to control whether PAYSERV is allowed to specify USER1 as an alternate-user profile when it opens the reply-to queue.

The alternate-user profile is specified on the *AlternateUserId* field of the object descriptor.

Note: You can use alternate-user profiles on any IBM WebSphere MQ object. Use of an alternate-user profile does not affect the user profile used by any other resource managers.

Context authority

Context is information that applies to a particular message and is contained in the message descriptor, MQMD, which is part of the message.

For descriptions of the message descriptor fields relating to context, see  [Overview for MQMD \(WebSphere MQ V7.1 Reference\)](#).

For information about the context options, see  [Message context \(WebSphere MQ V7.1 Programming Guide\)](#).

Remote security considerations

For remote security, consider:

Put authority

For security across queue managers, you can specify the put authority that is used when a channel receives a message sent from another queue manager.

This parameter is valid only for RCVR, RQSTR, or CLUSRCVR channel types. Specify the channel attribute PUTAUT as follows:

DEF Default user profile. This is the QMQM user profile under which the message channel agent is running.

CTX The user profile in the message context.

Transmission queues

Queue managers automatically put remote messages on a transmission queue; no special authority is required. However, putting a message directly on a transmission queue requires special authorization.

Channel exits

Channel exits can be used for added security.

Channel authentication records

Use to exercise more precise control over the access granted to connecting systems at a channel level.

For more information about remote security, see  Channel security.

Protecting channels with SSL or TLS

The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols provide channel security, with protection against eavesdropping, tampering, and impersonation. IBM WebSphere MQ support for SSL and TLS enables you to specify, on the channel definition, that a particular channel uses SSL or TLS security. You can also specify details of the security you want, such as the encryption algorithm you want to use.

SSL and TLS support in IBM WebSphere MQ uses the queue manager *authentication information object* and various CL and MQSC commands and queue manager and channel parameters that define the SSL or TLS support required in detail.

The following CL commands support SSL or TLS:

WRKMQMAUTI

Work with the attributes of an authentication information object.

CHGMQMAUTI

Modify the attributes of an authentication information object.

CRTMQMAUTI

Create an authentication information object.

CPYMQMAUTI

Create an authentication information object by copying an existing one.

DLTMQMAUTI

Delete an authentication information object.

DSPMQMAUTI

Displays the attributes for a specific authentication information object.

For an overview of channel security using SSL or TLS, see “Protecting channels with SSL” on page 449.

For details of PCF commands associated with SSL or TLS, see  Change, Copy, and Create Authentication Information Object (*WebSphere MQ V7.1 Reference*),  Delete Authentication Information Object (*WebSphere MQ V7.1 Reference*), and  Inquire Authentication Information Object (*WebSphere MQ V7.1 Reference*).

Setting up security on z/OS

Security considerations specific to z/OS.

Security in WebSphere MQ for z/OS is controlled using RACF or an equivalent external security manager (ESM). The instructions given here assume that you are using RACF.

RACF security classes

RACF classes are used to hold the profiles required for WebSphere MQ security checking. Many of the member classes have equivalent group classes. You must activate the classes and enable them to accept generic profiles

Each RACF class holds one or more profiles used at some point in the checking sequence, as shown in Table 30.

Table 30. RACF classes used by WebSphere MQ

Member class	Group class	Contents
MQADMIN	GMQADMIN	Profiles: Used mainly for holding profiles for administration-type functions. For example: <ul style="list-style-type: none"> • Profiles for WebSphere MQ security switches • The RESLEVEL security profile • Profiles for alternate user security • The context security profile • Profiles for command resource security
MXADMIN	GMXADMIN	Profiles: Used mainly for holding profiles for administration-type functions. For example: <ul style="list-style-type: none"> • Profiles for WebSphere MQ security switches • The RESLEVEL security profile • Profiles for alternate user security • The context security profile • Profiles for command resource security <p>This class can hold both uppercase and mixed case RACF profiles.</p>
MQCONN		Profiles used for connection security
MQCMD5		Profiles used for command security
MQQUEUE	GMQQUEUE	Profiles used in queue resource security
MXQUEUE	GMXQUEUE	Mixed case and uppercase profiles used in queue resource security
MQPROC	GMQPROC	Profiles used in process resource security
MXPROC	GMXPROC	Mixed case and uppercase profiles used in process resource security
MQNLIST	GMQNLIST	Profiles used in namelist resource security
MXNLIST	GMXNLIST	Mixed case and uppercase profiles used in namelist resource security
MXTOPIC	GMXTOPIC	Mixed case and uppercase profiles used in topic security

Some classes have a related *group class* that enables you to put together groups of resources that have similar access requirements. For details about the difference between the member and group classes and when to use a member or group class, see the *z/OS SecureWay Security Server RACF Security Administrator's Guide*.

The classes must be activated before security checks can be made. To activate all the WebSphere MQ classes, you use can use this RACF command:

```
SETROPTS CLASSACT(MQADMIN,MXADMIN,MQQUEUE,MXQUEUE,MQPROC,MXPROC,
MQNLIST,MXNLIST,MXTOPIC,MQCONN,MQCMD5)
```

You should also ensure that you set up the classes so that they can accept generic profiles. You also do this with the RACF command SETROPTS, for example:

```
SETROPTS GENERIC(MQADMIN,MXADMIN,MQQUEUE,MXQUEUE,MQPROC,MXPROC,
MQNLIST,MXNLIST,MXTOPIC,MQCONN,MQCMD5)
```

RACF profiles

All RACF profiles used by WebSphere MQ contain a prefix, which is either the queue manager name or the queue-sharing group name. Be careful when you use the percent sign as a wildcard.

All RACF profiles used by WebSphere MQ contain a prefix. For queue-sharing group level security, this is the queue-sharing group name. For queue manager level security, the prefix is the queue manager name. If you are using a mixture of queue manager and queue-sharing group level security, you will use profiles with both types of prefix. (Queue-sharing group and queue manager level security are described in  WebSphere MQ for z/OS concepts: security (*WebSphere MQ V7.1 Product Overview Guide*)).

For example, if you want to protect a queue called QUEUE_FOR_SUBSCRIBER_LIST in queue-sharing group QSG1 at queue-sharing group level, the appropriate profile would be defined to RACF as:

```
RDEFINE MQQUEUE QSG1.QUEUE_FOR_SUBSCRIBER_LIST
```

If you want to protect a queue called QUEUE_FOR_LOST_CARD_LIST, that belongs to queue manager STCD at queue manager level, the appropriate profile would be defined to RACF as:

```
RDEFINE MQQUEUE STCD.QUEUE_FOR_LOST_CARD_LIST
```

This means that different queue managers and queue-sharing groups can share the same RACF database and yet have different security options.

Do not use generic queue manager names in profiles to avoid unanticipated user access.

WebSphere MQ allows the use of the percent sign (%) in object names. However, RACF uses the % character as a single-character wildcard. This means that when you define an object name with a % character in its name, you must consider this when you define the corresponding profile.

For example, for the queue CREDIT_CARD_%_RATE_INQUIRY, on queue manager CRDP, the profile would be defined to RACF as follows:

```
RDEFINE MQQUEUE CRDP.CREDIT_CARD_%_RATE_INQUIRY
```

This queue cannot be protected by a generic profile, such as, CRDP.**.

WebSphere MQ allows the use of mixed case characters in object names. You can protect these objects by defining:

1. Mixed case profiles in the appropriate mixed case RACF classes, or
2. Generic profiles in the appropriate uppercase RACF classes.

There are some profiles, or parts of profiles, that remain uppercase only as the values are provided by WebSphere MQ. These are:

- Switch profiles.
- All high-level qualifiers (HLQ) including subsystem and Queue-Sharing Group identifiers.
- Profiles for SYSTEM objects.
- Profiles for Default objects.
- The **MQCMDS** class, so all command profiles are uppercase only.
- The **MQCONN** class, so all connection profiles are uppercase only.
- **RESLEVEL** profiles.
- The 'object' qualification in command resource profiles; for example, hlq.QUEUE.queueName. The resource name only is mixed case.
- Dynamic queue profiles hlq.CSQOREXX.* , hlq.CSQUTIL.* , and CSQXCMD.*.
- The 'CONTEXT' part of hlq.CONTEXT.resourcename.
- The 'ALTERNATE.USER' part of hlq.ALTERNATE.USER.userid.

For example, if you have a queue called PAYROLL.Dept1 on Queue Manager QM01 and you are using:

- Mixed case classes; you can define a profile in the WebSphere MQ RACF class MXQUEUE
RDEFINE MXQUEUE QM01.PAYROLL.Dept1
- Uppercase classes; you can define a profile in the WebSphere MQ RACF class MQQUEUE
RDEFINE MQQUEUE QM01.PAYROLL.*

The first example, using mixed case classes, gives you more granular control over granting authority to access the resource.

Switch profiles

To control the security checking performed by WebSphere MQ, you use *switch profiles*. A switch profile is a normal RACF profile that has a special meaning to WebSphere MQ. The access list in switch profiles is not used by WebSphere MQ.

WebSphere MQ maintains an internal switch for each switch type shown in tables Switch profiles for subsystem level security, Switch profiles for queue-sharing group or queue manager level security, and Switch profiles for resource checking. Switch profiles can be maintained at queue-sharing group level, or at queue manager level, or at a combination of both. Using a single set of queue-sharing group security switch profiles, you can control security on all the queue managers within a queue-sharing group.

When a security switch is set on, the security checks associated with the switch are performed. When a security switch is set off, the security checks associated with the switch are bypassed. The default is that all security switches are set on.

Switches and classes:

When you start a queue manager or refresh security, WebSphere MQ sets switches according to the state of various RACF classes.

When a queue manager is started (or when the MQADMIN or MXADMIN class is refreshed by the WebSphere MQ REFRESH SECURITY command), WebSphere MQ first checks the status of RACF and the appropriate class:

- The MQADMIN class if you are using uppercase profiles
- The MXADMIN class if you are using mixed case profile.

It sets the subsystem security switch off if any of these conditions is true:

- RACF is inactive or not installed.

- The MQADMIN or MXADMIN class is not defined (these classes are always defined for RACF because they are included in the class descriptor table (CDT)).
- The MQADMIN or MXADMIN class has not been activated.

If both RACF and the MQADMIN or MXADMIN class are active, WebSphere MQ checks the MQADMIN or MXADMIN class to see whether any of the switch profiles have been defined. It first checks the profiles described in “Profiles to control subsystem security” on page 525. If subsystem security is not required, WebSphere MQ sets the internal subsystem security switch off, and performs no further checks.

The profiles determine whether the corresponding WebSphere MQ switch is set on or off.

- If the switch is off, that type of security is deactivated.
- If any WebSphere MQ switch is set on, WebSphere MQ checks the status of the RACF class associated with the type of security corresponding to the WebSphere MQ switch. If the class is not installed or not active, the WebSphere MQ switch is set off. For example, process security checks are not carried out if the MQPROC or MXPROC class has not been activated. The class not being active is equivalent to defining NO.PROCESS.CHECKS profile for every queue manager and queue-sharing group that uses this RACF database.

How switches work:

To set off a security switch, define a NO.* switch profile for it. You can override a NO.* profile set at the queue-sharing group level by defining a YES.* profile for a queue manager.

To set off a security switch, you need to define a NO.* switch profile for it. The existence of a NO.* profile means that security checks are **not** performed for that type of resource, unless you choose to override a queue-sharing group level setting on a particular queue manager. This is described in “Overriding queue-sharing group level settings.”

If your queue manager is not a member of a queue-sharing group, you do not need to define any queue-sharing group level profiles or any override profiles. However, you must remember to define these profiles if the queue manager joins a queue-sharing group at a later date.

Each NO.* switch profile that WebSphere MQ detects turns off the checking for that type of resource. Switch profiles are activated during startup of the queue manager. If you change the switch profiles while any affected queue managers are running, you can get WebSphere MQ to recognize the changes by issuing the WebSphere MQ REFRESH SECURITY command.

The switch profiles must always be defined in the MQADMIN or MXADMIN class. Do not define them in the GMQADMIN or GMXADMIN class. Tables Switch profiles for subsystem level security and Switch profiles for resource checking show the valid switch profiles and the security type they control.

Overriding queue-sharing group level settings

You can override queue-sharing group level security settings for a particular queue manager that is a member of that group. If you want to perform queue manager checks on an individual queue manager that are not performed on other queue managers in the group, use the (qmgr-name.YES.*) switch profiles.

Conversely, if you do not want to perform a certain check on one particular queue manager within a queue-sharing group, define a (qmgr-name.NO.*) profile for that particular resource type on the queue manager, and do not define a profile for the queue-sharing group. (WebSphere MQ only checks for a queue-sharing group level profile if it does not find a queue manager level profile.)

Profiles to control subsystem security:

WebSphere MQ checks whether subsystem security checks are required for the subsystem, for the queue manager, and for the queue-sharing group.

The first security check made by WebSphere MQ is used to determine whether security checks are required for the whole WebSphere MQ subsystem. If you specify that you do not want subsystem security, no further checks are made.

The following switch profiles are checked to determine whether subsystem security is required. Figure 80 shows the order in which they are checked.

Table 31. Switch profiles for subsystem level security

Switch profile name	Type of resource or checking that is controlled
qmgr-name.NO.SUBSYS.SECURITY	Subsystem security for this queue manager
qsg-name.NO.SUBSYS.SECURITY	Subsystem security for this queue-sharing group
qmgr-name.YES.SUBSYS.SECURITY	Subsystem security override for this queue manager

If your queue manager is not a member of a queue-sharing group, WebSphere MQ checks for the qmgr-name.NO.SUBSYS.SECURITY switch profile only.

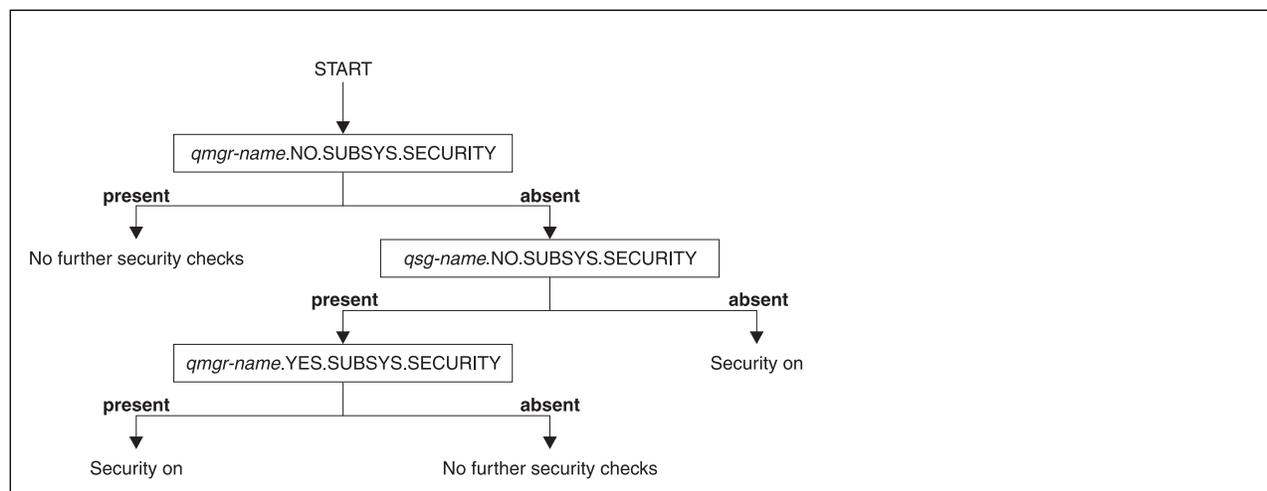


Figure 80. Checking for subsystem security

Profiles to control queue-sharing group or queue manager level security:

If subsystem security checking is required, WebSphere MQ checks whether security checking is required at queue-sharing group or queue manager level.

When WebSphere MQ has determined that security checking is required, it then determines whether checking is required at queue-sharing group or queue manager level, or both. These checks are not performed if your queue manager is not a member of a queue sharing group.

The following switch profiles are checked to determine the level required. Figure 81 on page 526 and Figure 82 on page 526 show the order in which they are checked.

Table 32. Switch profiles for queue-sharing group or queue manager level security

Switch profile name	Type of resource or checking that is controlled
qmgr-name.NO.QMGR.CHECKS	No queue manager level checks for this queue manager
qsg-name.NO.QMGR.CHECKS	No queue manager level checks for this queue-sharing group
qmgr-name.YES.QMGR.CHECKS	Queue manager level checks override for this queue manager
qmgr-name.NO.QSG.CHECKS	No queue-sharing group level checks for this queue manager
qsg-name.NO.QSG.CHECKS	No queue-sharing group level checks for this queue-sharing group
qmgr-name.YES.QSG.CHECKS	Queue-sharing group level checks override for this queue manager

If subsystem security is active, you cannot switch off both queue-sharing group and queue manager level security. If you try to do so, WebSphere MQ sets security checking on at both levels.

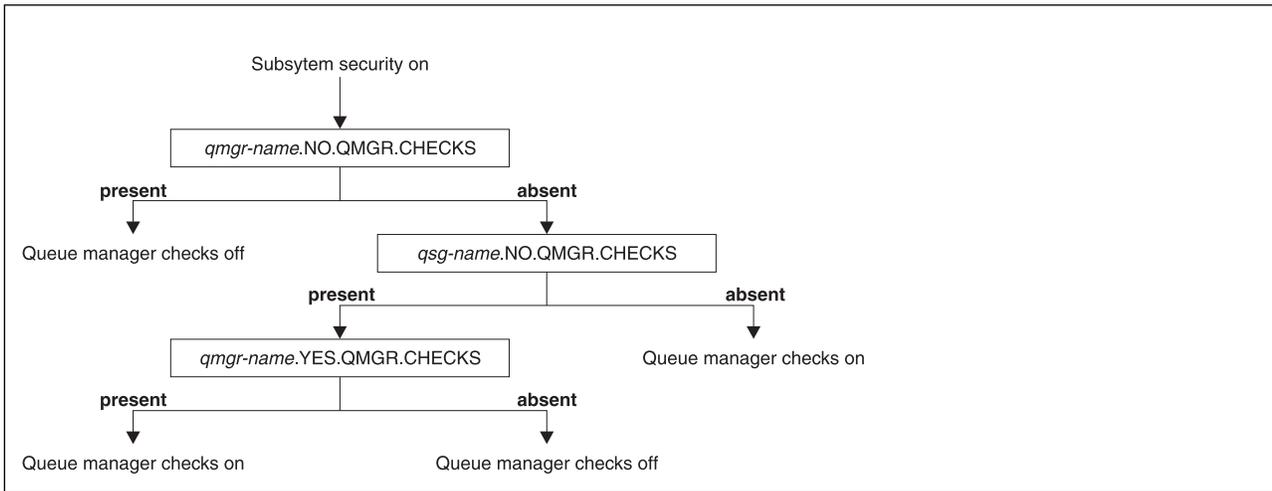


Figure 81. Checking for queue manager level security

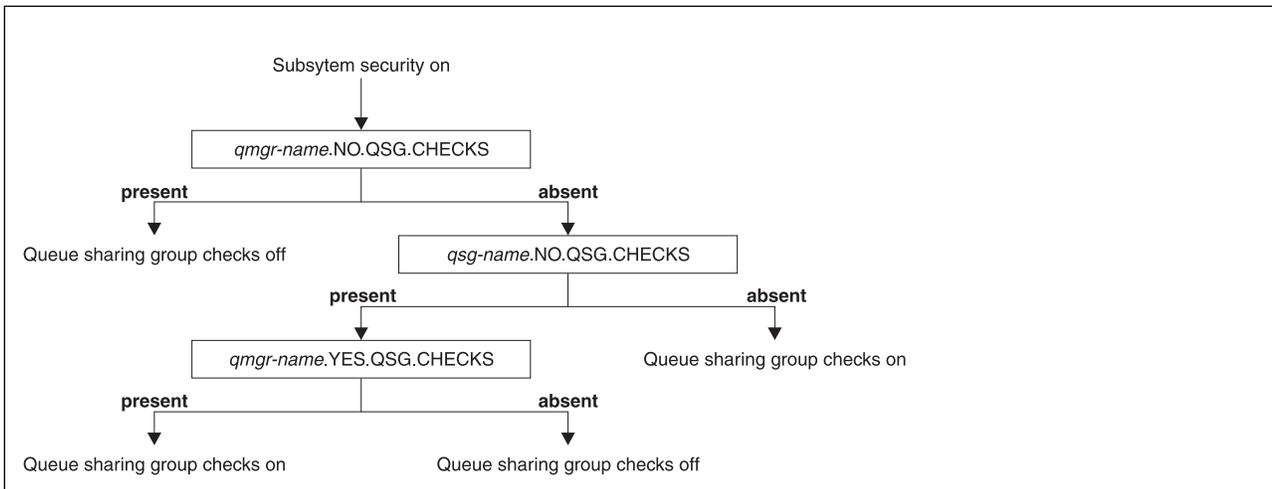


Figure 82. Checking for queue-sharing group level security

Valid combinations of security switches:

Only certain combinations of switches are valid. If you use a combination of switch settings that is not valid, message CSQH026I is issued and security checking is set on at both queue-sharing group and queue manager level.

Table 33, Table 34, Table 35, and Table 36 on page 528 show the sets of combinations of switch settings that are valid for each type of security level.

Table 33. Valid security switch combinations for queue manager level security

Combinations
qmgr-name.NO.QSG.CHECKS
qsg-name.NO.QSG.CHECKS
qmgr-name.NO.QSG.CHECKS qsg-name.NO.QMGR.CHECKS qmgr-name.YES.QMGR.CHECKS
qsg-name.NO.QSG.CHECKS qsg-name.NO.QMGR.CHECKS qmgr-name.YES.QMGR.CHECKS

Table 34. Valid security switch combinations for queue-sharing group level security

Combinations
qmgr-name.NO.QMGR.CHECKS
qsg-name.NO.QMGR.CHECKS
qmgr-name.NO.QMGR.CHECKS qsg-name.NO.QSG.CHECKS qmgr-name.YES.QSG.CHECKS
qsg-name.NO.QMGR.CHECKS qsg-name.NO.QSG.CHECKS qmgr-name.YES.QSG.CHECKS

Table 35. Valid security switch combinations for queue manager and queue-sharing group level security

Combinations
qsg-name.NO.QMGR.CHECKS qmgr-name.YES.QMGR.CHECKS No QSG.* profiles defined
No QMGR.* profiles defined qsg-name.NO.QSG.CHECKS qmgr-name.YES.QSG.CHECKS
qsg-name.NO.QMGR.CHECKS qmgr-name.YES.QMGR.CHECKS qsg-name.NO.QSG.CHECKS qmgr-name.YES.QSG.CHECKS
No profiles for either switch defined

Table 36. Other valid security switch combinations that switch both levels of checking on.

Combinations
qmgr-name.NO.QMGR.CHECKS qmgr-name.NO.QSG.CHECKS
qsg-name.NO.QMGR.CHECKS qsg-name.NO.QSG.CHECKS
qmgr-name.NO.QMGR.CHECKS qsg-name.NO.QSG.CHECKS
qsg-name.NO.QMGR.CHECKS qmgr-name.NO.QSG.CHECKS

Resource level checks:

A number of switch profiles are used to control access to resources. Some stop checking being performed on either a queue manager or a queue-sharing group. These can be overridden by profiles that enable checking for specific queue managers.

Table 37 shows the switch profiles used to control access to WebSphere MQ resources.

If your queue manager is part of a queue sharing group and you have both queue manager and queue-sharing group security active, you can use a YES.* switch profile to override queue-sharing group level profiles and specifically turn on security for a particular queue manager.

Some profiles apply to both queue managers and queue-sharing groups. These are prefixed by the string *hlq* and you should substitute the name of your queue-sharing group or queue manager, as applicable. Profile names shown prefixed by *qmgr-name* are queue-manager override profiles; you should substitute the name of your queue manager.

Table 37. Switch profiles for resource checking

Type of resource checking that is controlled	Switch profile name	Override profile for a particular queue manager
Connection security	hlq.NO.CONNECT.CHECKS	qmgr-name.YES.CONNECT.CHECKS
Queue security	hlq.NO.QUEUE.CHECKS	qmgr-name.YES.QUEUE.CHECKS
Process security	hlq.NO.PROCESS.CHECKS	qmgr-name.YES.PROCESS.CHECKS
Namelist security	hlq.NO.NLIST.CHECKS	qmgr-name.YES.NLIST.CHECKS
Context security	hlq.NO.CONTEXT.CHECKS	qmgr-name.YES.CONTEXT.CHECKS
Alternate user security	hlq.NO.ALTERNATE.USER.CHECKS	qmgr-name.YES.ALTERNATE.USER.CHECKS
Command security	hlq.NO.CMD.CHECKS	qmgr-name.YES.CMD.CHECKS
Command resource security	hlq.NO.CMD.RESC.CHECKS	qmgr-name.YES.CMD.RESC.CHECKS
Topic security	hlq.NO.TOPIC.CHECKS	qmgr-name.YES.TOPIC.CHECKS
Note: Generic switch profiles such as hlq.NO.** are ignored by WebSphere MQ		

For example, if you want to perform process security checks on queue manager QM01, which is a member of queue-sharing group QSG3 but you do not want to perform process security checks on any of the other queue managers in the group, define the following switch profiles:

```
QSG3.NO.PROCESS.CHECKS
QM01.YES.PROCESS.CHECKS
```

If you want to have queue security checks performed on all the queue managers in the queue-sharing group, except QM02, define the following switch profile:

```
QM02.NO.QUEUE.CHECKS
```

(There is no need to define a profile for the queue sharing group because the checks are automatically enabled if there is no profile defined.)

An example of defining switches:

Different WebSphere MQ subsystems have different security requirements, which can be implemented using different switch profiles.

Four WebSphere MQ subsystems have been defined:

- MQP1 (a production system)
- MQP2 (a production system)
- MQD1 (a development system)
- MQT1 (a test system)

All four queue managers are members of queue-sharing group QS01. All WebSphere MQ RACF classes have been defined and activated.

These subsystems have different security requirements:

- The production systems require full WebSphere MQ security checking to be active at queue-sharing group level on both systems.

This is done by specifying the following profile:

```
RDEFINE MQADMIN QS01.NO.QMGR.CHECKS
```

This sets queue-sharing group level checking for all the queue managers in the queue-sharing group. You do not need to define any other switch profiles for the production queue managers because you want to check everything for these systems.

- Test queue manager MQT1 also requires full security checking. However, because you might want to change this later, security can be defined at queue-manager level so that you can change the security settings for this queue manager without affecting the other members of the queue-sharing group.

This is done by defining the NO.QSG.CHECKS profile for MQT1 as follows:

```
RDEFINE MQADMIN MQT1.NO.QSG.CHECKS
```

- Development queue manager MQD1 has different security requirements from the rest of the queue-sharing group. It requires only connection and queue security to be active.

This is done by defining a MQD1.YES.QMGR.CHECKS profile for this queue manager, and then defining the following profiles to switch off security checking for the resources that do not need to be checked:

```
RDEFINE MQADMIN MQD1.NO.CMD.CHECKS
RDEFINE MQADMIN MQD1.NO.CMD.RESC.CHECKS
RDEFINE MQADMIN MQD1.NO.PROCESS.CHECKS
RDEFINE MQADMIN MQD1.NO.NLIST.CHECKS
RDEFINE MQADMIN MQD1.NO.CONTEXT.CHECKS
RDEFINE MQADMIN MQD1.NO.ALTERNATE.USER.CHECKS
```

When the queue manager is active, you can display the current security settings by issuing the DISPLAY SECURITY MQSC command.

You can also change the switch settings when the queue manager is running by defining or deleting the appropriate switch profile in the MQADMIN class. To make the changes to the switch settings active, you must issue the REFRESH SECURITY command for the MQADMIN class.

See “Refreshing queue manager security on z/OS” on page 578 for more details about using the DISPLAY SECURITY and REFRESH SECURITY commands.

Profiles used to control access to WebSphere MQ resources

You must define RACF profiles to control access to WebSphere MQ resources, in addition to the switch profiles that might have been defined. This collection of topics contains information about the RACF profiles for the different types of WebSphere MQ resource.

If you do not have a resource profile defined for a particular security check, and a user issues a request that would involve making that check, WebSphere MQ denies access. You do not have to define profiles for security types relating to any security switches that you have deactivated.

Profiles for connection security:

If connection security is active, you must define profiles in the MQCONN class and permit the necessary groups or user IDs access to those profiles, so that they can connect to WebSphere MQ.

To enable a connection to be made, you must grant users RACF READ access to the appropriate profile. (If no queue manager level profile exists, and your queue manager is a member of a queue-sharing group, checks might be made against queue-sharing group level profiles, if the security is set up to do this.)

A connection profile qualified with a queue manager name controls access to a specific queue manager and users given access to this profile can connect to that queue manager. A connection profile qualified with queue-sharing group name controls access to all queue managers within the queue-sharing group for that connection type. For example, a user with access to QS01.BATCH can use a batch connection to any queue manager in queue-sharing group QS01 that has not got a queue manager level profile defined.

Note:

1. For information about the user IDs checked for different security requests, see “User IDs for security checking” on page 568.
2. Resource level security (RESLEVEL) checks are also made at connection time. For details, see “The RESLEVEL security profile” on page 561.

WebSphere MQ security recognizes the following different types of connection:

- Batch (and batch-type) connections, these include:
 - z/OS batch jobs
 - TSO applications
 - USS sign-ons
 - Db2 stored procedures
- CICS connections
- IMS connections from control and application processing regions
- The WebSphere MQ channel initiator

Connection security profiles for batch connections:

Profiles for checking batch-type connections are composed of the queue manager or queue-sharing group name followed by the word *BATCH*. Give the user ID associated with the connecting address space READ access to the connection profile.

Profiles for checking batch and batch-type connections take the form:

```
hlq.BATCH
```

where *hlq* can be either the *qmgr-name* (queue manager name) or *qsg-name* (queue-sharing group name). If you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name. If it does not find one, it looks for a profile prefixed by the queue-sharing group name. If it fails to find either profile, the connection request fails.

For batch or batch-type connection requests, you must permit the user ID associated with the connecting address space to access the connection profile. For example, the following RACF command allows users in the CONNTQM1 group to connect to the queue manager TQM1; these user IDs will be permitted to use any batch or batch-type connection.

```
RDEFINE MQCONN TQM1.BATCH UACC(NONE)
PERMIT TQM1.BATCH CLASS(MQCONN) ID(CONNTQM1) ACCESS(READ)
```

Connection security profiles for CICS connections:

Profiles for checking CICS connections are composed of the queue manager or queue-sharing group name followed by the word *CICS*. Give the user ID associated with the CICS address space READ access to the connection profile.

Profiles for checking connections from CICS take the form:

```
hlq.CICS
```

where *hlq* can be either *qmgr-name* (queue manager name) or *qsg-name* (queue-sharing group name). If you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name. If it does not find one, it looks for a profile prefixed by the queue-sharing group name. If it fails to find either profile, the connection request fails.

For connection requests by CICS, you need only permit the CICS address space user ID access to the connection profile.

For example, the following RACF commands allow the CICS address space user ID KCBCICS to connect to the queue manager TQM1:

```
RDEFINE MQCONN TQM1.CICS UACC(NONE)
PERMIT TQM1.CICS CLASS(MQCONN) ID(KCBCICS) ACCESS(READ)
```

Connection security profiles for IMS connections:

Profiles for checking IMS connections are composed of the queue manager or queue-sharing group name followed by the word *IMS*. Give the IMS control and dependent region user IDs READ access to the connection profile.

Profiles for checking connections from IMS take the form:

```
hlq.IMS
```

where *hlq* can be either *qmgr-name* (queue manager name) or *qsg-name* (queue-sharing group name). If you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name. If it does not find one, it looks for a profile prefixed by the queue-sharing group name. If it fails to find either profile, the connection request fails

For connection requests by IMS, permit access to the connection profile for the IMS control and dependent region user IDs.

For example, the following RACF commands allow:

- The IMS region user ID, *IMSREG*, to connect to the queue manager *TQM1*.
- Users in group *BMPGRP* to submit BMP jobs.

```
RDEFINE MQCONN TQM1.IMS UACC(NONE)  
PERMIT TQM1.IMS CLASS(MQCONN) ID(IMSREG,BMPGRP) ACCESS(READ)
```

Connection security profiles for the channel initiator:

Profiles for checking connections from the channel initiator are composed of the queue manager or queue-sharing group name followed by the word *CHIN*. Give the user ID used by the channel initiator started task address space READ access to the connection profile.

Profiles for checking connections from the channel initiator take the form:

```
hlq.CHIN
```

where *hlq* can be either *qmgr-name* (queue manager name) or *qsg-name* (queue-sharing group name). If you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name. If it does not find one, it looks for a profile prefixed by the queue-sharing group name. If it fails to find either profile, the connection request fails

For connection requests by the channel initiator, define access to the connection profile for the user ID used by the channel initiator started task address space.

For example, the following RACF commands allow the channel initiator address space running with user ID *DQCTRL* to connect to the queue manager *TQM1*:

```
RDEFINE MQCONN TQM1.CHIN UACC(NONE)
PERMIT TQM1.CHIN CLASS(MQCONN) ID(DQCTRL) ACCESS(READ)
```

Profiles for queue security:

If queue security is active, you must define profiles in the appropriate classes and permit the necessary groups or user IDs access to these profiles. Queue security profiles are named after the queue manager or queue-sharing group, and the queue to be opened.

If queue security is active, you must:

- Define profiles in the **MQQUEUE** or **GMQUEUE** classes if using uppercase profiles.
- Define profiles in the **MXQUEUE** or **GMXQUEUE** classes if using mixed case profiles.
- Permit the necessary groups or user IDs access to these profiles, so that they can issue WebSphere MQ API requests that use queues.

Profiles for queue security take the form:

hlq.queueename

where hlq can be either qmgr-name (queue manager name) or qsg-name (queue-sharing group name), and queueename is the name of the queue being opened, as specified in the object descriptor on the **MQOPEN** or **MQPUT1** call.

A profile prefixed by the queue manager name controls access to a single queue on that queue manager. A profile prefixed by the queue-sharing group name controls access to access to one or more queues with that queue name on all queue managers within the queue-sharing group, or access to a shared queue by any queue manager within the group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that queue on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

If you are using shared queues, you are recommended to use queue-sharing group level security.

For details of how queue security operates when the queue name is that of an alias or a model queue, see “Considerations for alias queues” on page 535 and “Considerations for model queues” on page 536.

The RACF access required to open a queue depends on the **MQOPEN** or **MQPUT1** options specified. If more than one of the **MQOO_*** and **MQPMO_*** options is coded, the queue security check is performed for the highest RACF authority required.

Table 38. Access levels for queue security using the MQOPEN or MQPUT1 calls

MQOPEN or MQPUT1 option	RACF access level required to access hlq.queueename
MQOO_BROWSE	READ
MQOO_INQUIRE	READ
MQOO_BIND_*	UPDATE
MQOO_INPUT_*	UPDATE
MQOO_OUTPUT or MQPUT1	UPDATE

Table 38. Access levels for queue security using the MQOPEN or MQPUT1 calls (continued)

MQOPEN or MQPUT1 option	RACF access level required to access hlq.queueuname
MQOO_PASS_ALL_CONTEXT MQPMO_PASS_ALL_CONTEXT	UPDATE
MQOO_PASS_IDENTITY_CONTEXT MQPMO_PASS_IDENTITY_CONTEXT	UPDATE
MQOO_SAVE_ALL_CONTEXT	UPDATE
MQOO_SET_IDENTITY_CONTEXT MQPMO_SET_IDENTITY_CONTEXT	UPDATE
MQOO_SET_ALL_CONTEXT MQPMO_SET_ALL_CONTEXT	UPDATE
MQOO_SET	ALTER

For example, on WebSphere MQ queue manager QM77, all user IDs in the RACF group PAYGRP are to be given access to get messages from or put messages to all queues with names beginning with 'PAY.'. You can do this using these RACF commands:

```
RDEFINE MQQUEUE QM77.PAY.** UACC(NONE)
PERMIT QM77.PAY.** CLASS(MQQUEUE) ID(PAYGRP) ACCESS(UPDATE)
```

Also, all user IDs in the PAYGRP group must have access to put messages on queues that do not follow the PAY naming convention. For example:

```
REQUEST_QUEUE_FOR_PAYROLL
SALARY.INCREASE.SERVER
REPLIES.FROM.SALARY.MODEL
```

You can do this by defining profiles for these queues in the GMQQUEUE class and giving access to that class as follows:

```
RDEFINE GMQQUEUE PAYROLL.EXTRAS UACC(NONE)
      ADDMEM(QM77.REQUEST_QUEUE_FOR_PAYROLL,
            QM77.SALARY.INCREASE.SERVER,
            QM77.REPLIES.FROM.SALARY.MODEL)
PERMIT PAYROLL.EXTRAS CLASS(GMQQUEUE) ID(PAYGRP) ACCESS(UPDATE)
```

Note:

1. If the RACF access level that an application has to a queue security profile is changed, the changes only take effect for any new object handles obtained (that is, new **MQOPENS**) for that queue. Those handles already in existence at the time of the change retain their existing access to the queue. If an application is required to use its changed access level to the queue rather than its existing access level, it must close and reopen the queue for each object handle that requires the change.
2. In the example, the queue manager name QM77 could also be the name of a queue-sharing group.

Other types of security checks might also occur at the time the queue is opened depending on the open options specified and the types of security that are active. See also “Profiles for context security” on page 549 and “Profiles for alternate user security” on page 547. For a summary table showing the open options and the security authorization needed when queue, context, and alternate user security are all active, see Table 43 on page 540.

If you are using publish/subscribe you must consider the following. When an MQSUB request is processed a security check is performed to ensure that the user ID making the request has the required access to put messages to the target WebSphere MQ queue as well as the required access to subscribe to the WebSphere MQ topic.

Table 39. Access levels for queue security using the MQSUB call

MQSUB option	RACF access level required to access hlq.queueName
MQSO_ALTER, MQSO_CREATE, and MQSO_RESUME	UPDATE

Note:

1. The hlq.queueName is the destination queue for publications. When this is a managed queue, you need access to the appropriate model queue to be used for the managed queue and the dynamic queue that are created.
2. You can use a technique like this for the destination queue you provide on an MQSUB API call if you want to distinguish between the users making the subscriptions, and the users retrieving the publications from the destination queue.

Considerations for alias queues:

When you issue an **MQOPEN** or **MQPUT1** call for an alias queue, WebSphere MQ makes a resource check against the queue name specified in the object descriptor (MQOD) on the call. It does not check if the user is allowed access to the target queue name.

For example, an alias queue called PAYROLL.REQUEST resolves to a target queue of PAY.REQUEST. If queue security is active, you need only be authorized to access the queue PAYROLL.REQUEST. No check is made to see if you are authorized to access the queue PAY.REQUEST.

Using alias queues to distinguish between MQGET and MQPUT requests:

The range of MQI calls available in one access level can cause a problem if you want to restrict access to a queue to allow only the **MQPUT** call or only the **MQGET** call. A queue can be protected by defining two aliases that resolve to that queue: one that enables applications to get messages from the queue, and one that enable applications to put messages on the queue.

The following text gives you an example of how you can define your queues to WebSphere MQ:

```
DEFINE QLOCAL(MUST_USE_ALIAS_TO_ACCESS) GET(ENABLED)
      PUT(ENABLED)
```

```
DEFINE QALIAS(USE_THIS_ONE_FOR_GETS) GET(ENABLED)
      PUT(DISABLED) TARGQ(MUST_USE_ALIAS_TO_ACCESS)
```

```
DEFINE QALIAS(USE_THIS_ONE_FOR_PUTS) GET(DISABLED)
      PUT(ENABLED) TARGQ(MUST_USE_ALIAS_TO_ACCESS)
```

You must also make the following RACF definitions:

```
RDEFINE MQQUEUE hlq.MUST_USE_ALIAS_TO_ACCESS UACC(NONE)
RDEFINE MQQUEUE hlq.USE_THIS_ONE_FOR_GETS UACC(NONE)
RDEFINE MQQUEUE hlq.USE_THIS_ONE_FOR_PUTS UACC(NONE)
```

Then you ensure that no users have access to the queue hlq.MUST_USE_ALIAS_TO_ACCESS, and give the appropriate users or groups access to the alias. You can do this using the following RACF commands:

```
PERMIT hlq.USE_THIS_ONE_FOR_GETS CLASS(MQQUEUE)
      ID(GETUSER,GETGRP) ACCESS(UPDATE)
PERMIT hlq.USE_THIS_ONE_FOR_PUTS CLASS(MQQUEUE)
      ID(PUTUSER,PUTGRP) ACCESS(UPDATE)
```

This means user ID GETUSER and user IDs in the group GETGRP are only allowed to get messages on MUST_USE_ALIAS_TO_ACCESS through the alias queue USE_THIS_ONE_FOR_GETS; and user ID PUTUSER and user IDs in the group PUTGRP are only allowed to put messages through the alias queue USE_THIS_ONE_FOR_PUTS.

Note:

1. If you want to use a technique like this, you must inform your application developers, so that they can design their programs appropriately.
2. You can use a technique like this for the destination queue you provide on and MQSUB API request if you want to distinguish between the users making the subscriptions and the users 'getting' the publications from the destination queue.

Considerations for model queues:

To open a model queue, you must be able to open both the model queue itself and the dynamic queue to which it resolves. Define generic RACF profiles for dynamic queues, including dynamic queues used by WebSphere MQ utilities.

When you open a model queue, WebSphere MQ security makes two queue security checks:

1. Are you authorized to access the model queue?
2. Are you authorized to access the dynamic queue to which the model queue resolves?

If the dynamic queue name contains a trailing asterisk (*) character, this * is replaced by a character string generated by WebSphere MQ, to create a dynamic queue with a unique name. However, because the whole name, including this generated string, is used for checking authority, you should define generic profiles for these queues.

For example, an **MQOPEN** call uses a model queue name of CREDIT.CHECK.REPLY.MODEL and a dynamic queue name of CREDIT.REPLY.* on queue manager (or queue-sharing group) MQSP.

To do this, you must issue the following RACF commands to define the necessary queue profiles:

```
RDEFINE MQQUEUE MQSP.CREDIT.CHECK.REPLY.MODEL
RDEFINE MQQUEUE MQSP.CREDIT.REPLY.**
```

You must also issue the corresponding RACF PERMIT commands to allow the user access to these profiles.

A typical dynamic queue name created by an **MQOPEN** is something like CREDIT.REPLY.A346EF00367849A0. The precise value of the last qualifier is unpredictable; this is why you should use generic profiles for such queue names.

A number of WebSphere MQ utilities put messages on dynamic queues. You should define profiles for the following dynamic queue names, and provide RACF UPDATE access to the relevant user IDs (see "User IDs for security checking" on page 568 for the correct user IDs):

SYSTEM.CSQUTIL.*	(used by CSQUTIL)
SYSTEM.CSQOREXX.*	(used by the operations and control panels)
SYSTEM.CSQXCMD.*	(used by the channel initiator when processing CSQINPX)
CSQ4SAMP.*	(used by the WebSphere MQ supplied samples)

You might also consider defining a profile to control use of the dynamic queue name used by default in the application programming copy members. The WebSphere MQ-supplied copybooks contain a default *DynamicQName*, which is CSQ.*. This enables an appropriate RACF profile to be established.

Note: Do not allow application programmers to specify a single * for the dynamic queue name. If you do, you must define an hlq.** profile in the MQQUEUE class, and you would have to give it wide-ranging access. This means that this profile could also be used for other non-dynamic queues that do not have a more specific RACF profile. Your users could, therefore, gain access to queues you do not want them to access.

Close options on permanent dynamic queues:

If an application opens a permanent dynamic queue that was created by another application and then attempts to delete that queue with an **MQCLOSE** option, some extra security checks are applied when the attempt is made.

Table 40. Access levels for close options on permanent dynamic queues

MQCLOSE option	RACF access level required to hlq.queueName
MQCO_DELETE	ALTER
MQCO_DELETE_PURGE	ALTER

Security and remote queues:

When a message is put on a remote queue, the queue security that is implemented by the local queue manager depends on how the remote queue is specified when it is opened.

The following rules are applied:

1. If the remote queue has been defined on the local queue manager through the WebSphere MQ **DEFINE QREMOTE** command, the queue that is checked is the name of the remote queue. For example, if a remote queue is defined on queue manager MQS1 as follows:

```
DEFINE QREMOTE(BANK7.CREDIT.REFERENCE)
      RNAME(CREDIT.SCORING.REQUEST)
      RQMNAME(BNK7)
      XMITQ(BANK1.TO.BANK7)
```

In this case, a profile for BANK7.CREDIT.REFERENCE must be defined in the MQQUEUE class.

2. If the *ObjectQMgrName* for the request does not resolve to the local queue manager, a security check is carried out against the resolved (remote) queue manager name except in the case of a cluster queue where the check is made against the cluster queue name.

For example, the transmission queue BANK1.TO.BANK7 is defined on queue manager MQS1. An **MQPUT1** request is then issued on MQS1 specifying *ObjectName* as BANK1.INTERBANK.TRANSFERS and an *ObjectQMgrName* of BANK1.TO.BANK7. In this case, the user performing the request must have access to BANK1.TO.BANK7.

3. If you make an **MQPUT** request to a queue and specify *ObjectQMgrName* as the name of an alias of the local queue manager, only the queue name is checked for security, not that of the queue manager.

When the message gets to the remote queue manager it might be subject to additional security processing. For more information, see "Security for remote messaging" on page 430.

Dead-letter queue security:

Special considerations apply to the dead-letter queue, because many users must be able to put messages on it, but access to retrieve messages must be tightly restricted. You can achieve this by applying different RACF authorities to the dead-letter queue and an alias queue.

Undelivered messages can be put on a special queue called the dead-letter queue. If you have sensitive data that could possibly end up on this queue, you must consider the security implications of this because you do not want unauthorized users to retrieve this data.

Each of the following must be allowed to put messages onto the dead-letter queue:

- Application programs.
- The channel initiator address space and any MCA user IDs. (If the RESLEVEL profile is not present, or is defined so that channel user IDs are checked, the channel user ID also needs authority to put messages on the dead-letter queue.)
- CKTI, the WebSphere MQ-supplied CICS task initiator.
- CSQQTRMN, the WebSphere MQ-supplied IMS trigger monitor.

The only application that can retrieve messages from the dead-letter queue should be a 'special' application that processes these messages. However, a problem arises if you give applications RACF UPDATE authority to the dead-letter queue for **MQPUTs** because they can then automatically retrieve messages from the queue using **MQGET** calls. You cannot disable the dead-letter queue for get operations because, if you do, not even the 'special' applications could retrieve the messages.

One solution to this problem is set up a two-level access to the dead-letter queue. CKTI, message channel agent transactions or the channel initiator address space, and 'special' applications have direct access; other applications can only access the dead-letter queue through an alias queue. This alias is defined to allow applications to put messages on the dead-letter queue, but not to get messages from it.

This is how it might work:

1. Define the real dead-letter queue with attributes PUT(ENABLED) and GET(ENABLED), as shown in the sample thlqual.SCSQPROC(CSQ4INYG).
2. Give RACF UPDATE authority for the dead-letter queue to the following user IDs:
 - User IDs that the CKTI and the MCAs or channel initiator address space run under.
 - The user IDs associated with the 'special' dead-letter queue processing application.
3. Define an alias queue that resolves to the real dead-letter queue, but give the alias queue these attributes: PUT(ENABLED) and GET(DISABLED). Give the alias queue a name with the same stem as the dead-letter queue name but append the characters ".PUT" to this stem. For example, if the dead-letter queue name is hlq.DEAD.QUEUE, the alias queue name would be hlq.DEAD.QUEUE.PUT.
4. To put a message on the dead-letter queue, an application uses the alias queue. This is what your application must do:
 - Retrieve the name of the real dead-letter queue. To do this, it opens the queue manager object using **MQOPEN** and then issues an **MQINQ** to get the dead-letter queue name.
 - Build the name of the alias queue by appending the characters '.PUT' to this name, in this case, hlq.DEAD.QUEUE.PUT.
 - Open the alias queue, hlq.DEAD.QUEUE.PUT.
 - Put the message on the real dead-letter queue by issuing an **MQPUT** against the alias queue.
5. Give the user ID associated with the application RACF UPDATE authority to the alias, but no access (authority NONE) to the real dead-letter queue. This means that:
 - The application can put messages onto the dead-letter queue using the alias queue.

- The application cannot get messages from the dead-letter queue using the alias queue because the alias queue is disabled for get operations.

The application cannot get any messages from the real dead-letter queue either because it does not have the correct RACF authority.

Table 41 summarizes the RACF authority required for the various participants in this solution.

Table 41. RACF authority to the dead-letter queue and its alias

Associated user IDs	Real dead-letter queue (hlq.DEAD.QUEUE)	Alias dead-letter queue (hlq.DEAD.QUEUE.PUT)
MCA or channel initiator address space and CKTI	UPDATE	NONE
'Special' application (for dead-letter queue processing)	UPDATE	NONE
User-written application user IDs	NONE	UPDATE

If you use this method, the application cannot determine the maximum message length (MAXMSGL) of the dead-letter queue. This is because the MAXMSGL attribute cannot be retrieved from an alias queue. Therefore, your application should assume that the maximum message length is 100 MB, the maximum size WebSphere MQ for z/OS supports. The real dead-letter queue should also be defined with a MAXMSGL attribute of 100 MB.

Note: User-written application programs do not normally use alternate user authority to put messages on the dead-letter queue. This reduces the number of user IDs that have access to the dead-letter queue.

System queue security:

You must set up RACF access to allow certain user IDs access to particular system queues.

Many of the system queues are accessed by the ancillary parts of WebSphere MQ:

- The CSQUTIL utility
- The operations and control panels
- The channel initiator address space (including the Queued Pub/Sub Daemon)

The user IDs under which these run must be given RACF access to these queues, as shown in Table 42.

Table 42. Access required to the SYSTEM queues by WebSphere MQ

SYSTEM queue	CSQUTIL	Operations and control panels	Channel initiator for distributed queuing
SYSTEM.ADMIN.CHANNEL.EVENT	–	–	UPDATE
SYSTEM.BROKER.ADMIN.STREAM	–	–	ALTER
SYSTEM.BROKER.CONTROL.QUEUE	–	–	ALTER
SYSTEM.BROKER.DEFAULT.STREAM	–	–	ALTER
SYSTEM.BROKER.INTER.BROKER.COMMUNICATIONS	–	–	UPDATE
SYSTEM.CHANNEL.INITQ	–	–	UPDATE
SYSTEM.CHANNEL.SYNCQ	–	–	UPDATE
SYSTEM.CLUSTER.COMMAND.QUEUE	–	–	ALTER
SYSTEM.CLUSTER.REPOSITORY.QUEUE	–	–	UPDATE
SYSTEM.CLUSTER.TRANSMIT.QUEUE	–	–	ALTER

Table 42. Access required to the SYSTEM queues by WebSphere MQ (continued)

SYSTEM queue	CSQUTIL	Operations and control panels	Channel initiator for distributed queuing
SYSTEM.COMMAND.INPUT	UPDATE	UPDATE	UPDATE
SYSTEM.COMMAND.REPLY.*	-	-	UPDATE
SYSTEM.COMMAND.REPLY.MODEL	UPDATE	UPDATE	UPDATE
SYSTEM.CSQOREXX.*	-	UPDATE	-
SYSTEM.CSQUTIL.*	UPDATE	-	-
SYSTEM.CSQXCMD.*	-	-	UPDATE
SYSTEM.HIERARCHY.STATE	-	-	UPDATE
SYSTEM.INTER.QMGR.CONTROL	-	-	UPDATE
SYSTEM.INTER.QMGR.PUBS	-	-	UPDATE
SYSTEM.INTER.QMGR.FANREQ	-	-	UPDATE
SYSTEM.QSG.CHANNEL.SYNCQ	-	-	UPDATE
SYSTEM.QSG.TRANSMIT.QUEUE	-	-	UPDATE

API-resource security access quick reference:

A summary of the **MQOPEN**, **MQPUT1**, **MQSUB**, and **MQCLOSE** options and the access required by the different resource security types.

Table 43. MQOPEN, MQPUT1, MQSUB, and MQCLOSE options and the security authorization required. Callouts shown like this (1) refer to the notes following this table.

RACF class: RACF profile:	Minimum RACF access level required			
	MXTOPIC (15 or 16)	MQQUEUE or MXQUEUE(1) (2)	MQADMIN or MXADMIN (3)	MQADMIN or MXADMIN (4)
MQOPEN option				
MQOO_INQUIRE		READ (5)	No check	No check
MQOO_BROWSE		READ	No check	No check
MQOO_INPUT_*		UPDATE	No check	No check
MQOO_SAVE_ALL_CONTEXT (6)		UPDATE	No check	No check
MQOO_OUTPUT (USAGE=NORMAL) (7)		UPDATE	No check	No check
MQOO_PASS_IDENTITY_CONTEXT (8)		UPDATE	READ	No check
MQOO_PASS_ALL_CONTEXT (8) (9)		UPDATE	READ	No check
MQOO_SET_IDENTITY_CONTEXT (8) (9)		UPDATE	UPDATE	No check
MQOO_SET_ALL_CONTEXT (8) (10)		UPDATE	CONTROL	No check
MQOO_OUTPUT (USAGE (XMITQ) (11)		UPDATE	CONTROL	No check
MQOO_OUTPUT (topic object)	UPDATE (16)			
MQOO_OUTPUT (alias queue to topic object)	UPDATE (16)	UPDATE		
MQOO_SET		ALTER	No check	No check
MQOO_ALTERNATE_USER_AUTHORITY		(12)	(12)	UPDATE
MQPUT1 option				

Table 43. MQOPEN, MQPUT1, MQSUB, and MQCLOSE options and the security authorization required (continued). Callouts shown like this (1) refer to the notes following this table.

RACF class: RACF profile:	Minimum RACF access level required			
	MXTOPIC (15 or 16)	MQQUEUE or MXQUEUE(1) (2)	MQADMIN or MXADMIN (3)	MQADMIN or MXADMIN (4)
Put on a normal queue (7)		UPDATE	No check	No check
MQPMO_PASS_IDENTITY_CONTEXT		UPDATE	READ	No check
MQPMO_PASS_ALL_CONTEXT		UPDATE	READ	No check
MQPMO_SET_IDENTITY_CONTEXT		UPDATE	UPDATE	No check
MQPMO_SET_ALL_CONTEXT		UPDATE	CONTROL	No check
MQOO_OUTPUT Put on a transmission queue (11)		UPDATE	CONTROL	No check
MQOO_OUTPUT (topic object)	UPDATE (16)			
MQOO_OUTPUT (alias queue to topic object)	UPDATE (16)	UPDATE		
MQPMO_ALTERNATE_USER_AUTHORITY		(13)	(13)	UPDATE
MQCLOSE option				
MQCO_DELETE (14)		ALTER	No check	No check
MQCO_DELETE_PURGE (14)		ALTER	No check	No check
MQCO_REMOVE_SUB	ALTER (15)			
MQSUB option				
MQSO_CREATE	ALTER (15)	(17)	(18)	
MQSO_ALTER	ALTER (15)	(17)	(18)	
MQSO_RESUME	READ (15)	(17)	No check	
MQSO_ALTERNATE_USER_AUTHORITY				UPDATE
MQSO_SET_IDENTITY_CONTEXT			(18)	

Note:

1. This option is not restricted to queues. Use the MQNLIST or MXNLIST class for namelists, and the MQPROC or MXPROC class for processes.
2. Use RACF profile: hlq.resourcename
3. Use RACF profile: hlq.CONTEXT.queueuname
4. Use RACF profile: hlq.ALTERNATE.USER.alternateuserid
alternateuserid is the user identifier that is specified in the *AlternateUserId* field of the object descriptor. Note that up to 12 characters of the *AlternateUserId* field are used for this check, unlike other checks where only the first 8 characters of a user identifier are used.
5. No check is made when opening the queue manager for inquiries.
6. MQOO_INPUT_* must be specified as well. This is valid for a local, model or alias queue.
7. This check is done for a local or model queue that has a *Usage* queue attribute of MQUS_NORMAL, and also for an alias or remote queue (that is defined to the connected queue manager.) If the queue is a remote queue that is opened specifying an *ObjectQMgrName* (not the name of the connected queue manager) explicitly, the check is carried out against the queue with the same name as *ObjectQMgrName* (which must be a local queue with a *Usage* queue attribute of MQUS_TRANSMISSION).
8. MQOO_OUTPUT must be specified as well.

9. MQOO_PASS_IDENTITY_CONTEXT is implied as well by this option.
10. MQOO_PASS_IDENTITY_CONTEXT, MQOO_PASS_ALL_CONTEXT and MQOO_SET_IDENTITY_CONTEXT are implied as well by this option.
11. This check is done for a local or model queue that has a *Usage* queue attribute of MQUS_TRANSMISSION, and is being opened directly for output. It does not apply if a remote queue is being opened.
12. At least one of MQOO_INQUIRE, MQOO_BROWSE, MQOO_INPUT_*, MQOO_OUTPUT or MQOO_SET must be specified as well. The check carried out is the same as that for the other options specified.
13. The check carried out is the same as that for the other options specified.
14. This applies only for permanent dynamic queues that have been opened directly, that is, not opened through a model queue. No security is required to delete a temporary dynamic queue.
15. Use RACF profile hlq.SUBSCRIBE.topicname.
16. Use RACF profile hlq.PUBLISH.topicname.
17. If on the MQSUB request you specified a destination queue for the publications to be sent to, then a security check is carried out against that queue to ensure that you have put authority to that queue.
18. If on the MQSUB request, with MQSO_CREATE or MQSO_ALTER options specified, you want to set any of the identity context fields in the MQSD structure, you also need to specify the MQSO_SET_IDENTITY_CONTEXT option and you also need the appropriate authority to the context profile for the destination queue.

Profiles for topic security:

If topic security is active, you must define profiles in the appropriate classes and permit the necessary groups or user IDs access to those profiles.

The concept of topic security within a topic tree is described in Publish/subscribe security.

If topic security is active, you must perform the following actions:

- Define profiles in the **MXTOPIC** or **GMXTOPIC** classes.
- Permit the necessary groups or user IDs access to these profiles, so that they can issue IBM WebSphere MQ API requests that use topics.

Profiles for topic security take the form:

```
hlq.SUBSCRIBE.topicname
hlq.PUBLISH.topicname
```

where

- hlq is either qmgr-name (queue manager name) or qsg-name (queue-sharing group name).
- topicname is the name of the topic administration node in the topic tree, associated either with the topic being subscribed to through an MQSUB call, or being published to through an MQOPEN call.

A profile prefixed by the queue manager name controls access to a single topic on that queue manager. A profile prefixed by the queue-sharing group name controls access to one or more topics with that topic name on all queue managers within the queue-sharing group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that topic on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, IBM WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

Subscribe

To subscribe to a topic, you need access to both the topic you are trying to subscribe to, and the target queue for the publications.

When you issue an MQSUB request, the following security checks take place:

- Whether you have the appropriate level of access to subscribe to that topic, and also that the target queue (if specified) is opened for output
- Whether you have the appropriate level of access to that target queue.

Table 44. Access level required for topic security to subscribe

Action	Access level required
MQSUB to a topic	RACF access required to <i>hlq.SUBSCRIBE.topicname</i> profile in MXTOPIC class
MQSO_CREATE and MQSO_ALTER	ALTER
MQSO_RESUME	READ
MQSUB - additional authority to non-managed destination queues.	RACF access required to <i>hlq.CONTEXT.queueName</i> profile in MQADMIN or MXADMIN class
MQSO_CREATE, MQSO_ALTER, and MQSO_RESUME	UPDATE
	RACF access required to <i>hlq.queueName</i> profile in MQQUEUE or MXQUEUE class
MQSO_CREATE and MQSO_ALTER	UPDATE
	RACF access required to <i>hlq.ALTERNATE.USER.alternateuserid</i> profile in MQADMIN or MXADMIN class
MQSO_ALTERNATE_USER_AUTHORITY	UPDATE

Considerations for managed queues for subscriptions

A security check is carried out to see if you are allowed to subscribe to the topic. However, no security checks are carried out when the managed queue is created, or to determine if you have access to put messages to this destination queue.

You cannot close delete a managed queue.

The model queues used are:SYSTEM.DURABLE.MODEL.QUEUE and SYSTEM.NDURABLE.MODEL.QUEUE.

The managed queues created from these model queues are of the form SYSTEM.MANAGED.DURABLE.A346EF00367849A0 and SYSTEM.MANAGED.NDURABLE.A346EF0036785EA0 where the last qualifier is unpredictable.

Do not give any user access to these queues. The queues can be protected using generic profiles of the form SYSTEM.MANAGED.DURABLE.* and SYSTEM.MANAGED.NDURABLE.* with no authorities granted.

Messages can be retrieved from these queues using the handle returned on the MQSUB request.

If you explicitly issue an MQCLOSE call for a subscription with the MQCO_REMOVE_SUB option specified, and you did not create the subscription you are closing under this handle, a security check is performed at the time of closure to ensure that you have the correct authority to perform the operation.

Table 45. Access level required to profiles for topic security for closure of a subscribe operation

Action	Access level required
MQCLOSE (of a subscription)	RACF access required to <i>hlq.SUBSCRIBE.topicname</i> profile in MXTOPIC class
MQCO_REMOVE_SUB	ALTER

Publish

To publish on a topic you need access to the topic and, if you are using alias queues, to the alias queue as well.

Table 46. Access level required to profiles for topic security for a publish operation

Action	Access level required
MQOPEN (of a topic)	RACF access required to <i>hlq.PUBLISH.topicname</i> profile in MXTOPIC class
MQOO_OUTPUT or MQPUT1	UPDATE
MQOPEN (Alias queue to topic)	RACF access required to <i>hlq.queueName</i> profile in MQQUEUE or MXQUEUE class for the alias queue
MQOO_OUTPUT or MQPUT1	UPDATE

For details of how topic security operates when an alias queue that resolves to a topic name is opened for publish, see “Considerations for alias queues that resolve to topics for a publish operation.”

When you consider alias queues used for destination queues for PUT or GET restrictions, see “Considerations for alias queues” on page 535.

If the RACF access level that an application has to a topic security profile is changed, the changes take effect only for any new object handles obtained (that is, a new MQSUB or MQOPEN) for that topic. Those handles already in existence at the time of the change retain their existing access to the topic. Also, existing subscribers retain their access to any subscriptions that they have already made.

Considerations for alias queues that resolve to topics for a publish operation

When you issue an MQOPEN or MQPUT1 call for an alias queue that resolves to a topic, IBM WebSphere MQ makes two resource checks:

- The first one against the alias queue name specified in the object descriptor (MQOD) on the MQOPEN or MQPUT1 call.
- The second against the topic to which the alias queue resolves

You must be aware that this behavior is different from the behavior you get when alias queues resolve to other queues. You need the correct access to both profiles in order for the publish action to proceed.

System topic security

Many of the system topics are accessed by the ancillary parts of IBM WebSphere MQ, for example the channel initiator address space.

The user IDs under which this runs must be given RACF access to these queues, as shown in Table 47 on page 545.

Table 47. Access required to the SYSTEM topics

SYSTEM topic	Profile	Channel Initiator for Distributed queuing
SYSTEM.BROKER.ADMIN.STREAM	PUBLISH.topicname	UPDATE
SYSTEM.BROKER.ADMIN.STREAM	SUBSCRIBE.topicname	ALTER

Profiles for processes:

If process security is active, you must define profiles in the appropriate classes and permit the necessary groups or user IDs access to those profiles.

If process security is active, you must:

- Define profiles in the **MQPROC** or **GMQPROC** classes if using uppercase profiles.
- Define profiles in the **MXPROC** or **GMXPROC** classes if using mixed case profiles.
- Permit the necessary groups or user IDs access to these profiles, so that they can issue WebSphere MQ API requests that use processes.

Profiles for processes take the form:

where hlq can be either qmgr-name (queue manager name) or qsg-name (queue-sharing group name), and

hlq.processname

processname is the name of the process being opened.

A profile prefixed by the queue manager name controls access to a single process definition on that queue manager. A profile prefixed by the queue-sharing group name controls access to one or more process definitions with that name on all queue managers within the queue-sharing group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that process definition on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

The following table shows the access required for opening a process.

Table 48. Access levels for process security

MQOPEN option	RACF access level required to hlq.processname
MQOO_INQUIRE	READ

For example, on queue manager MQS9, the RACF group INQVPRC must be able to inquire (**MQINQ**) on all processes starting with the letter V. The RACF definitions for this would be:

```
RDEFINE MQPROC MQS9.V* UACC(NONE)
PERMIT MQS9.V* CLASS(MQPROC) ID(INQVPRC) ACCESS(READ)
```

Alternate user security might also be active, depending on the open options specified when a process definition object is opened.

Profiles for namelists:

If namelist security is active, you define profiles in the appropriate classes and give the necessary groups or user IDs access to these profiles.

If namelist security is active, you must:

- Define profiles in the **MQNLIST** or **GMQNLIST** classes if using uppercase profiles.
- Define profiles in the **MXNLIST** or **GMXNLIST** classes if using mixed case profiles.
- Permit the necessary groups or user IDs access to these profiles.

Profiles for namelists take the form:

hlq.namelistname

where hlq can be either qmgr-name (queue manager name) or qsg-name (queue-sharing group name), and namelistname is the name of the namelist being opened.

A profile prefixed by the queue manager name controls access to a single namelist on that queue manager. A profile prefixed by the queue-sharing group name controls access to access to one or more namelists with that name on all queue managers within the queue-sharing group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that namelist on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

The following table shows the access required for opening a namelist.

Table 49. Access levels for namelist security

MQOPEN option	RACF access level required to hlq.namelistname
MQOO_INQUIRE	READ

For example, on queue manager (or queue-sharing group) PQM3, the RACF group DEPT571 must be able to inquire (**MQINQ**) on these namelists:

- All namelists starting with "DEPT571".
- PRINTER/DESTINATIONS/DEPT571
- AGENCY/REQUEST/QUEUES
- WAREHOUSE.BROADCAST

The RACF definitions to do this are:

```

RDEFINE MQNLIST PQM3.DEPT571.** UACC(NONE)
PERMIT PQM3.DEPT571.** CLASS(MQNLIST) ID(DEPT571) ACCESS(READ)

RDEFINE GMQNLIST NLISTS.FOR.DEPT571 UACC(NONE)
  ADDMEM(PQM3.PRINTER/DESTINATIONS/DEPT571,
    PQM3.AGENCY/REQUEST/QUEUES,
    PQM3.WAREHOUSE.BROADCAST)
PERMIT NLISTS.FOR.DEPT571 CLASS(GMQNLIST) ID(DEPT571) ACCESS(READ)

```

Alternate user security might be active, depending on the options specified when a namelist object is opened.

System namelist security

Many of the system namelists are accessed by the ancillary parts of WebSphere MQ:

- The CSQUTIL utility
- The operations and control panels
- The channel initiator address space (including the Queued Publish/Subscribe Daemon)

The user IDs under which these run must be given RACF access to these namelists, as shown in Table 50.

Table 50. Access required to the SYSTEM namelists by WebSphere MQ

SYSTEM namelist	CSQUTIL	Operations and control panels	Channel initiator for distributed queuing
SYSTEM.QPUBSUB.QUEUE.NAMELIST	–	–	READ
SYSTEM.QPUBSUB.SUBPOINT.NAMELIST	–	–	READ

Profiles for alternate user security:

If alternate user security is active, you must define profiles in the appropriate classes and permit the necessary groups or user IDs access to those profiles.

If alternate user security is active, you must:

- Define profiles in the MQADMIN or GMQADMIN classes if you are using uppercase profiles.
- Define profiles in the MXADMIN or GMXADMIN classes if you are using mixed case profiles.

Permit the necessary groups or user IDs access to these profiles, so that they can use the ALTERNATE_USER_AUTHORITY options when the object is opened.

Profiles for alternate user security can be specified at subsystem level or at queue-sharing group level and take the following form:

```
hlq.ALTERNATE.USER.alternateuserid
```

Where hlq can be either qmgr-name (queue manager name) or qsg-name (queue-sharing group name), and alternateuserid is the value of the *AlternateUserId* field in the object descriptor.

A profile prefixed by the queue manager name controls use of an alternative user ID on that queue manager. A profile prefixed by the queue-sharing group name controls use of an alternative user ID on all queue managers within the queue-sharing group. This alternative user ID can be used on any queue manager within the queue-sharing group by a user that has the correct access. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that alternative user ID on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

The following table shows the access when specifying an alternative user option.

Table 51. Access levels for alternate user security

MQOPEN, MQSUB, or MQPUT1 option	RACF access level required
MQOO_ALTERNATE_USER_AUTHORITY MQSO_ALTERNATE_USER_AUTHORITY MQPMO_ALTERNATE_USER_AUTHORITY	UPDATE

In addition to alternate user security checks, other security checks for queue, process, namelist, and context security can also be made. The alternative user ID, if provided, is only used for security checks on queue, process definition, or namelist resources. For alternate user and context security checks, the user ID requesting that the check is used. For details about how user IDs are handled, see “User IDs for security checking” on page 568. For a summary table showing the open options and the security checks required when queue, context and alternate user security are all active, see Table 43 on page 540.

An alternative user profile gives the requesting user ID access to resources associated with the user ID specified in the alternative user ID. For example, the payroll server running under user ID PAYSERV on queue manager QMPY processes requests from personnel user IDs, all of which start with PS. To cause the work performed by the payroll server to be carried out under the user ID of the requesting user, alternative user authority is used. The payroll server knows which user ID to specify as the alternative user ID because the requesting programs generate messages using the MQPMO_DEFAULT_CONTEXT put message option. See “User IDs for security checking” on page 568 for more details about from where alternative user IDs are obtained.

The following example RACF definitions enable the server program to specify alternative user IDs starting with the characters PS:

```
RDEFINE MQADMIN QMPY.ALTERNATE.USER.PS* UACC(NONE)
PERMIT QMPY.ALTERNATE.USER.PS* CLASS(MQADMIN) ID(PAYSERV) ACCESS(UPDATE)
```

Note:

1. The *AlternateUserId* fields in the object descriptor and subscription descriptor are 12 bytes long. All 12 bytes are used in the profile checks, but only the first 8 bytes are used as the user ID by WebSphere MQ. If this user ID truncation is not desirable, application programs making the request must translate any alternative user ID over 8 bytes into something more appropriate.
2. If you specify MQOO_ALTERNATE_USER_AUTHORITY, MQSO_ALTERNATE_USER_AUTHORITY, or MQPMO_ALTERNATE_USER_AUTHORITY and you do not specify an *AlternateUserId* field in the object descriptor, a user ID of blanks is used. For the purposes of the alternate user security check the user ID used for the *AlternateUserId* qualifier is -BLANK-. For example RDEF MQADMIN h1q.ALTERNATE.USER.-BLANK-.

If the user is allowed to access this profile, all further checks are made with a user ID of blanks. For details of blank user IDs, see “Blank user IDs and UACC levels” on page 576.

The administration of alternative user IDs is easier if you have a naming convention for user IDs that enables you to use generic alternative user profiles. If they do not, you can use the RACF RACVARS feature. For details about using RACVARS, see the *z/OS SecureWay Security Server RACF Security Administrator’s Guide*.

When a message is put to a queue that has been opened with alternative user authority and the context of the message has been generated by the queue manager, the MQMD_USER_IDENTIFIER field is set to the alternative user ID.

Profiles for context security:

WebSphere MQ uses profiles for controlling access to the context information specific to a particular message. The context is contained within the message descriptor (MQMD).

Using profiles for context security

If context security is active, you must:

- Define a profile in the **MQADMIN** class if using uppercase profiles.
- Define profile in the **MXADMIN** class if using mixed case profiles.

The profile is called hlq.CONTEXT.queueName, where:

hlq Can be either qmgr-name (queue manager name) or qsg-name (queue-sharing group name).

queueName

Can be either the full name of the queue you want to define the context profile for, or a generic profile.

Special considerations apply if you are migrating from a previous version; see  Migrating from a previous version.

A profile prefixed by the queue manager name, and with ** specified as the queue name, allows control for context security on all queues belonging to that queue manager. This can be overridden on an individual queue by defining a queue level profile for context on that queue.

A profile prefixed by the queue-sharing group name, and with ** specified as the queue name, allows control for context on all queues belonging to the queue managers within the queue-sharing group. This can be overridden on an individual queue manager by defining a queue-manager level profile for context on that queue manager, by specifying a profile prefixed by the queue manager name. It can also be overridden on an individual queue by specifying a profile suffixed with the queue name.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

You must give the necessary groups or user IDs access to this profile. The following table shows the access level required, depending on the specification of the context options when the queue is opened.

Table 52. Access levels for context security

MQOPEN or MQPUT1 option	RACF access level required to hlq.CONTEXT.queueName
MQPMO_NO_CONTEXT	No context security check
MQPMO_DEFAULT_CONTEXT	No context security check
MQOO_SAVE_ALL_CONTEXT	No context security check
MQOO_PASS_IDENTITY_CONTEXT MQPMO_PASS_IDENTITY_CONTEXT	READ
MQOO_PASS_ALL_CONTEXT MQPMO_PASS_ALL_CONTEXT	READ

Table 52. Access levels for context security (continued)

MQOPEN or MQPUT1 option	RACF access level required to hlq.CONTEXT.queueName
MQOO_SET_IDENTITY_CONTEXT MQPMO_SET_IDENTITY_CONTEXT	UPDATE
MQOO_SET_ALL_CONTEXT MQPMO_SET_ALL_CONTEXT	CONTROL
MQOO_OUTPUT or MQPUT1 (USAGE(XMITQ))	CONTROL
MQSUB option	
MQSO_SET_IDENTITY_CONTEXT(Note 2)	UPDATE
Note:	
<ol style="list-style-type: none"> 1. The user IDs used for distributed queuing require CONTROL access to hlq.CONTEXT.queueName to put messages on the destination queue. See "User IDs used by the channel initiator" on page 571 for information about the user IDs used. 2. If on the MQSUB request, with MQSO_CREATE or MQSO_ALTER options specified, you want to set any of the identity context fields in the MQSD structure, you need to specify the MQSO_SET_IDENTITY_CONTEXT option. You require also, the appropriate authority to the context profile for the destination queue. 	

If you put commands on the system-command input queue, use the default context put message option to associate the correct user ID with the command.

For example, the WebSphere MQ-supplied utility program CSQUTIL can be used to offload and reload messages in queues. When offloaded messages are restored to a queue, the CSQUTIL utility uses the MQOO_SET_ALL_CONTEXT option to return the messages to their original state. In addition to the queue security required by this open option, context authority is also required. For example, if this authority is required by the group BACKGRP on queue manager MQS1, this would be defined by:

```
RDEFINE MQADMIN MQS1.CONTEXT.** UACC(NONE)
PERMIT MQS1.CONTEXT.** CLASS(MQADMIN) ID(BACKGRP) ACCESS(CONTROL)
```

Depending on the options specified, and the types of security performed, other types of security checks might also occur when the queue is opened. These include queue security (see "Profiles for queue security" on page 533), and alternate user security (see "Profiles for alternate user security" on page 547). For a summary table showing the open options and the security checks required when queue, context and alternate user security are all active, see Table 43 on page 540.

System queue context security

Many of the system queues are accessed by the ancillary parts of WebSphere MQ, for example the channel initiator address space.

The user IDs under which this runs must be given RACF access to these queues, as shown in Table 53 on page 551.

Table 53. Access required to the SYSTEM queues for context operations

SYSTEM queue	Channel Initiator for Distributed queuing
SYSTEM.BROKER.CONTROL.QUEUE	CONTROL
SYSTEM.BROKER.INTER.BROKER.COMMUNICATIONS	CONTROL
SYSTEM.CHANNEL.SYNCQ	CONTROL
SYSTEM.CLUSTER.TRANSMIT.QUEUE	CONTROL

Profiles for command security:

To enable security checking for commands, add profiles to the MQCMD5 class. The profile names are based on the MQSC commands but control both MQSC and PCF commands. Profiles can apply to a queue manager or a queue-sharing group.

If you want security checking for commands (so you have not defined the command security switch profile `hlq.NO.COMD.CHECKS`) you must add profiles to the MQCMD5 class.

The same security profiles control both MQSC and PCF commands. The names of the RACF profiles for command security checking are based on the MQSC command names themselves. These profiles take the form:

`hlq.verb.pkw`

Where `hlq` can be either `qmgr-name` (queue manager name) or `qsg-name` (queue-sharing group name), `verb` is the verb part of the command name, for example `ALTER`, and `pkw` is the object type, for example `QLOCAL` for a local queue.

Thus, the profile name for the `ALTER QLOCAL` command in subsystem `CSQ1` is:

`CSQ1.ALTER.QLOCAL`

You can use generic profiles to protect sets of commands so that you have fewer profiles to maintain and, therefore, fewer access lists. Consider creating a generic profile that applies to all commands not protected by a more specific profile. Define this profile with `UACC(NONE)` and grant `ALTER` access only to the RACF groups containing administrators. You might then create a generic profile applicable to all `DISPLAY` commands and grant widespread access to it. Between these extremes, you might identify groups of users needing access to certain sets of commands, in which case you can create profiles for those sets and grant access to RACF groups representing those classes of user. Avoid giving users access to commands they do not require: Apply the principle of least privilege, so that users only have access to the commands that are required for their jobs.

A profile prefixed by the queue manager name controls the use of the command on that queue manager. A profile prefixed by the queue-sharing group name controls the use of the command on all queue managers within the queue-sharing group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that command on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

By setting up command profiles at queue manager level, a user can be restricted from issuing commands on a particular queue manager. Alternatively, you can define one profile for a queue-sharing group for each command verb, and all security checks take place against that profile instead of individual queue managers.

If both subsystem security and queue-sharing group security are active and a local profile is not found, a command security check is performed to see if the user has access to a queue-sharing group profile.

If you use the CMDSCOPE attribute to route a command to other queue managers in a queue-sharing group, security is checked on each queue manager where the command is run, but not necessarily on the queue manager where the command is entered.

Table 54 shows, for each WebSphere MQ MQSC command, the profiles required for command security checking to be carried out, and the corresponding access level for each profile in the MQCMD5 class.

Table 55 on page 556 shows, for each WebSphere MQ PCF command, the profiles required for command security checking to be carried out, and the corresponding access level for each profile in the MQCMD5 class.

Table 54. MQSC commands, profiles, and their access levels

Command	Command profile for MQCMD5	Access level for MQCMD5	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
ALTER AUTHINFO	hlq.ALTER.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
ALTER BUFFPOOL	hlq.ALTER.BUFFPOOL	ALTER	No check	–
ALTER CFSTRUCT	hlq.ALTER.CFSTRUCT	ALTER	No check	–
ALTER CHANNEL	hlq.ALTER.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
ALTER NAMELIST	hlq.ALTER.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
ALTER PROCESS	hlq.ALTER.PROCESS	ALTER	hlq.PROCESS.process	ALTER
ALTER PSID	hlq.ALTER.PSID	ALTER	No check	–
ALTER QALIAS	hlq.ALTER.QALIAS	ALTER	hlq.QUEUE.queue	ALTER
ALTER QLOCAL	hlq.ALTER.QLOCAL	ALTER	hlq.QUEUE.queue	ALTER
ALTER QMGR	hlq.ALTER.QMGR	ALTER	No check	–
ALTER QMODEL	hlq.ALTER.QMODEL	ALTER	hlq.QUEUE.queue	ALTER
ALTER QREMOTE	hlq.ALTER.QREMOTE	ALTER	hlq.QUEUE.queue	ALTER
ALTER SECURITY	hlq.ALTER.SECURITY	ALTER	No check	–
ALTER SMDS	hlq.ALTER.SMDS	ALTER	No check	–
ALTER STGCLASS	hlq.ALTER.STGCLASS	ALTER	No check	–
ALTER SUB	hlq.ALTER.SUB	ALTER	No check	–
ALTER TOPIC	hlq.ALTER.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
ALTER TRACE	hlq.ALTER.TRACE	ALTER	No check	–
ARCHIVE LOG	hlq.ARCHIVE.LOG	CONTROL	No check	–
BACKUP CFSTRUCT	hlq.BACKUP.CFSTRUCT	CONTROL	No check	–
CLEAR QLOCAL	hlq.CLEAR.QLOCAL	ALTER	hlq.QUEUE.queue	ALTER
CLEAR TOPICSTR 3 on page 556	hlq.CLEAR.TOPICSTR	ALTER	hlq.TOPIC.topic	ALTER
DEFINE AUTHINFO	hlq.DEFINE.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
DEFINE BUFFPOOL	hlq.DEFINE.BUFFPOOL	ALTER	No check	–
DEFINE CFSTRUCT	hlq.DEFINE.CFSTRUCT	ALTER	No check	–
DEFINE CHANNEL	hlq.DEFINE.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER

Table 54. MQSC commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
DEFINE LOG	hlq.DEFINE.LOG	ALTER	No check	–
DEFINE MAXSMGS	hlq.DEFINE.MAXSMGS	ALTER	No check	–
DEFINE NAMELIST	hlq.DEFINE.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
DEFINE PROCESS	hlq.DEFINE.PROCESS	ALTER	hlq.PROCESS.process	ALTER
DEFINE PSID	hlq.DEFINE.PSID	ALTER	No check	–
DEFINE QALIAS	hlq.DEFINE.QALIAS	ALTER	hlq.QUEUE.queue	ALTER
DEFINE QLOCAL	hlq.DEFINE.QLOCAL	ALTER	hlq.QUEUE.queue	ALTER
DEFINE QMODEL	hlq.DEFINE.QMODEL	ALTER	hlq.QUEUE.queue	ALTER
DEFINE QREMOTE	hlq.DEFINE.QREMOTE	ALTER	hlq.QUEUE.queue	ALTER
DEFINE STGCLASS	hlq.DEFINE.STGCLASS	ALTER	No check	–
DEFINE SUB	hlq.DEFINE.SUB	ALTER	No check	–
DEFINE TOPIC	hlq.DEFINE.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
DELETE AUTHINFO	hlq.DELETE.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
DELETE BUFFPOOL	hlq.DELETE.BUFFPOOL	ALTER	No check	–
DELETE CFSTRUCT	hlq.DELETE.CFSTRUCT	ALTER	No check	–
DELETE CHANNEL	hlq.DELETE.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
DELETE NAMELIST	hlq.DELETE.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
DELETE PROCESS	hlq.DELETE.PROCESS	ALTER	hlq.PROCESS.process	ALTER
DELETE PSID	hlq.DELETE.PSID	ALTER	No check	–
DELETE QALIAS	hlq.DELETE.QALIAS	ALTER	hlq.QUEUE.queue	ALTER
DELETE QLOCAL	hlq.DELETE.QLOCAL	ALTER	hlq.QUEUE.queue	ALTER
DELETE QMODEL	hlq.DELETE.QMODEL	ALTER	hlq.QUEUE.queue	ALTER
DELETE QREMOTE	hlq.DELETE.QREMOTE	ALTER	hlq.QUEUE.queue	ALTER
DELETE STGCLASS	hlq.DELETE.STGCLASS	ALTER	No check	–
DELETE SUB	hlq.DELETE.SUB	ALTER	No check	–
DELETE TOPIC	hlq.DELETE.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
DISPLAY ARCHIVE 1 on page 556	hlq.DISPLAY.ARCHIVE	READ	No check	–
DISPLAY AUTHINFO	hlq.DISPLAY.AUTHINFO	READ	No check	–
DISPLAY CFSTATUS	hlq.DISPLAY.CFSTATUS	READ	No check	–
DISPLAY CFSTRUCT	hlq.DISPLAY.CFSTRUCT	READ	No check	–
DISPLAY CHANNEL	hlq.DISPLAY.CHANNEL	READ	No check	–
DISPLAY CHINIT	hlq.DISPLAY.CHINIT	READ	No check	–
DISPLAY CHLAUTH	hlq.DISPLAY.CHLAUTH	READ	No check	–
DISPLAY CHSTATUS	hlq.DISPLAY.CHSTATUS	READ	No check	–
DISPLAY CLUSQMGR	hlq.DISPLAY.CLUSQMGR	READ	No check	–
DISPLAY CMDSERV	hlq.DISPLAY.CMDSERV	READ	No check	–

Table 54. MQSC commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
DISPLAY CONN 1 on page 556	hlq.DISPLAY.CONN	READ	No check	–
DISPLAY GROUP	hlq.DISPLAY.GROUP	READ	No check	–
DISPLAY LOG 1 on page 556	hlq.DISPLAY.LOG	READ	No check	–
DISPLAY MAXSMGS	hlq.DISPLAY.MAXSMGS	READ	No check	–
DISPLAY NAMELIST	hlq.DISPLAY.NAMELIST	READ	No check	–
DISPLAY PROCESS	hlq.DISPLAY.PROCESS	READ	No check	–
DISPLAY PUBSUB	hlq.DISPLAY.PUBSUB	READ	No check	–
DISPLAY QALIAS	hlq.DISPLAY.QALIAS	READ	No check	–
DISPLAY QCLUSTER	hlq.DISPLAY.QCLUSTER	READ	No check	–
DISPLAY QLOCAL	hlq.DISPLAY.QLOCAL	READ	No check	–
DISPLAY QMGR	hlq.DISPLAY.QMGR	READ	No check	–
DISPLAY QMODEL	hlq.DISPLAY.QMODEL	READ	No check	–
DISPLAY QREMOTE	hlq.DISPLAY.QREMOTE	READ	No check	–
DISPLAY QSTATUS	hlq.DISPLAY.QSTATUS	READ	No check	–
DISPLAY QUEUE	hlq.DISPLAY.QUEUE	READ	No check	–
DISPLAY SBSTATUS	hlq.DISPLAY.SBSTATUS	READ	No check	–
DISPLAY SMDS	hlq.DISPLAY.SMDS	READ	No check	–
DISPLAY SMDSCONN	hlq.DISPLAY.SMDSCONN	READ	No check	–
DISPLAY SUB	hlq.DISPLAY.SUB	READ	No check	–
DISPLAY SECURITY	hlq.DISPLAY.SECURITY	READ	No check	–
DISPLAY STGCLASS	hlq.DISPLAY.STGCLASS	READ	No check	–
DISPLAY SYSTEM 1 on page 556	hlq.DISPLAY.SYSTEM	READ	No check	–
DISPLAY THREAD	hlq.DISPLAY.THREAD	READ	No check	–
DISPLAY TPSTATUS	hlq.DISPLAY.TPSTATUS	READ	No check	–
DISPLAY TOPIC	hlq.DISPLAY.TOPIC	READ	No check	–
DISPLAY TPSTATUS	hlq.DISPLAY.TPSTATUS	READ	No check	–
DISPLAY TRACE	hlq.DISPLAY.TRACE	READ	No check	–
DISPLAY USAGE 1 on page 556	hlq.DISPLAY.USAGE	READ	No check	–
MOVE QLOCAL	hlq.MOVE.QLOCAL	ALTER	hlq.QUEUE.from-queue hlq.QUEUE.to-queue	ALTER
PING CHANNEL	hlq.PING.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
RECOVER BSDS	hlq.RECOVER.BSDS	CONTROL	No check	–
RECOVER CFSTRUCT	hlq.RECOVER.CFSTRUCT	CONTROL	No check	–
REFRESH CLUSTER	hlq.REFRESH.CLUSTER	ALTER	No check	–
REFRESH QMGR	hlq.REFRESH.QMGR	ALTER	No check	–

Table 54. MQSC commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
REFRESH SECURITY	hlq.REFRESH.SECURITY	ALTER	No check	–
RESET CFSTRUCT	hlq.RESET.CFSTRUCT	CONTROL	No check	–
RESET CHANNEL	hlq.RESET.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
RESET CLUSTER	hlq.RESET.CLUSTER	CONTROL	No check	–
RESET QMGR	hlq.RESET.QMGR	CONTROL	No check	–
RESET QSTATS	hlq.RESET.QSTATS	CONTROL	hlq.QUEUE.queue	CONTROL
RESET SMDS	hlq.RESET.SMDS	CONTROL	No check	–
RESET TPIPE	hlq.RESET.TPIPE	CONTROL	No check	–
RESOLVE CHANNEL	hlq.RESOLVE.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
RESOLVE INDOUBT	hlq.RESOLVE.INDOUBT	CONTROL	No check	–
RESUME QMGR	hlq.RESUME.QMGR	CONTROL	No check	–
RVERIFY SECURITY	hlq.RVERIFY.SECURITY	ALTER	No check	–
SET ARCHIVE	hlq.SET.ARCHIVE	CONTROL	No check	–
SET CHLAUTH	hlq.SET.CHLAUTH	CONTROL	No check	–
SET LOG	hlq.SET.LOG	CONTROL	No check	–
SET SYSTEM	hlq.SET.SYSTEM	CONTROL	No check	–
START CHANNEL	hlq.START.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
START CHINIT 4 on page 556	hlq.START.CHINIT	CONTROL	No check	–
START CMDSERV	hlq.START.CMDSERV	CONTROL	No check	–
START LISTENER	hlq.START.LISTENER	CONTROL	No check	–
START QMGR	None 2 on page 556	–	–	–
START SMDSCONN	hlq.START.SMDSCONN	CONTROL	No check	–
START TRACE	hlq.START.TRACE	CONTROL	No check	–
STOP CHANNEL	hlq.STOP.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
STOP CHINIT	hlq.STOP.CHINIT	CONTROL	No check	–
STOP CMDSERV	hlq.STOP.CMDSERV	CONTROL	No check	–
STOP LISTENER	hlq.STOP.LISTENER	CONTROL	No check	–
STOP QMGR	hlq.STOP.QMGR	CONTROL	No check	–
STOP SMDSCONN	hlq.STOP.SMDSCONN	CONTROL	No check	–
STOP TRACE	hlq.STOP.TRACE	CONTROL	No check	–
SUSPEND QMGR	hlq.SUSPEND.QMGR	CONTROL	No check	–

Table 54. MQSC commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
Notes:				
<p>1. These commands might be issued internally by the queue manager; no authority is checked in these cases.</p> <p>2. WebSphere MQ does not check the authority of the user who issues the START QMGR command. However, you can use RACF, or your alternative security facilities to control access to the START xxxxMSTR command that is issued as a result of the START QMGR command. This is done by controlling access to the MVS.START.STC.xxxxMSTR profile in the RACF operator commands (OPERCMD) class. For details of this procedure, see the <i>z/OS SecureWay Security Server RACF Security Administrator's Guide</i>. If you use this technique, and an unauthorized user tries to start the queue manager, it terminates with a reason code of 00F30216.</p> <p>3. The hlq.TOPIC.topic resource refers to the Topic object derived from the TOPICSTR. For more details, see "Publish/subscribe security" on page 808</p> <p>4. At releases prior to WebSphere MQ for z/OS V6, the security check was for MVS.START.STC.CSQ1CHIN. At WebSphere MQ for z/OS V6 and later, the resource name has an additional JOBNAME qualifier appended to it. This can cause problems when starting the channel initiator.</p> <p>To resolve the problem replace MVS.START.STC.ssidCHIN with a profile for a resource named MVS.START.STC.ssidCHIN.* or MVS.START.STC.ssidCHIN.ssidCHIN where <i>ssid</i> is the subsystem id for the queue manager. This requires RACF UPDATE authority. For more details, see the  z/OS product documentation for <i>Operation planning, MVS Commands, RACF Access Authorities, and Resource Names</i>.</p> <p>The START for <i>ssidMSTR</i> does not include the JOBNAME= parameter. For consistency, you might want to update the profile for MVS.START.STC.ssidMSTR to MVS.START.STC.ssidMSTR.*.</p>				

Table 55. PCF commands, profiles, and their access levels

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
Backup CF Structure	hlq.BACKUP.CFSTRUCT	CONTROL	No check	–
Change Authentication Information Object	hlq.ALTER.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
Change CF Structure	hlq.ALTER.CFSTRUCT	ALTER	No check	–
Change Channel	hlq.ALTER.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
Change Namelist	hlq.ALTER.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
Change Process	hlq.ALTER.PROCESS	ALTER	hlq.PROCESS.process	ALTER
Change Queue	hlq.ALTER.QUEUE	ALTER	hlq.QUEUE.queue	ALTER
Change Queue Manager	hlq.ALTER.QMGR	ALTER	No check	–
Change Security	hlq.ALTER.SECURITY	ALTER	No check	–
Change SMDS	hlq.ALTER.SMDS	ALTER	No check	–
Change Storage Class	hlq.ALTER.STGCLASS	ALTER	No check	–
Change Subscription	hlq.ALTER.SUB	ALTER	No check	–
Change Topic	hlq.ALTER.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
Clear Queue	hlq.CLEAR.QLOCAL	ALTER	hlq.QUEUE.queue	ALTER
Clear Topic String 1 on page 559	hlq.CLEAR.TOPICSTR	ALTER	hlq.TOPIC.topic	ALTER

Table 55. PCF commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
Copy Authentication Information Object	hlq.DEFINE.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
Copy CF Structure	hlq.DEFINE.CFSTRUCT	ALTER	No check	–
Copy Channel	hlq.DEFINE.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
Copy Namelist	hlq.DEFINE.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
Copy Process	hlq.DEFINE.PROCESS	ALTER	hlq.PROCESS.process	ALTER
Copy Queue	hlq.DEFINE.QUEUE	ALTER	hlq.QUEUE.queue	ALTER
Copy Subscription	hlq.DEFINE.SUB	ALTER	No check	–
Copy Storage Class	hlq.DEFINE.STGCLASS	ALTER	No check	–
Copy Topic	hlq.DEFINE.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
Create Authentication Information Object	hlq.DEFINE.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
Create CF Structure	hlq.DEFINE.CFSTRUCT	ALTER	No check	–
Create Channel	hlq.DEFINE.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
Create Namelist	hlq.DEFINE.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
Create Process	hlq.DEFINE.PROCESS	ALTER	hlq.PROCESS.process	ALTER
Create Queue	hlq.DEFINE.QUEUE	ALTER	hlq.QUEUE.queue	ALTER
Create Storage Class	hlq.DEFINE.STGCLASS	ALTER	No check	–
Create Subscription	hlq.DEFINE.SUB	ALTER	No check	–
Create Topic	hlq.DEFINE.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
Delete Authentication Information Object	hlq.DELETE.AUTHINFO	ALTER	hlq.AUTHINFO.resourcename	ALTER
Delete CF Structure	hlq.DELETE.CFSTRUCT	ALTER	No check	–
Delete Channel	hlq.DELETE.CHANNEL	ALTER	hlq.CHANNEL.channel	ALTER
Delete Namelist	hlq.DELETE.NAMELIST	ALTER	hlq.NAMELIST.namelist	ALTER
Delete Process	hlq.DELETE.PROCESS	ALTER	hlq.PROCESS.process	ALTER
Delete Queue	hlq.DELETE.QUEUE	ALTER	hlq.QUEUE.queue	ALTER
Delete Storage Class	hlq.DELETE.STGCLASS	ALTER	No check	–
Delete Subscription	hlq.DELETE.SUB	ALTER	No check	–
Delete Topic	hlq.DELETE.TOPIC	ALTER	hlq.TOPIC.topic	ALTER
Inquire Archive	hlq.DISPLAY.ARCHIVE	READ	No check	–
Inquire Authentication Information Object	hlq.DISPLAY.AUTHINFO	READ	No check	–
Inquire Authentication Information Object Names	hlq.DISPLAY.AUTHINFO	READ	No check	–
Inquire CF Structure	hlq.DISPLAY.CFSTRUCT	READ	No check	–
Inquire CF Structure Names	hlq.DISPLAY.CFSTRUCT	READ	No check	–

Table 55. PCF commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
Inquire CF Structure Status	hlq.DISPLAY.CFSTATUS	READ	No check	–
Inquire Channel	hlq.DISPLAY.CHANNEL	READ	No check	–
Inquire Channel Authentication Records	hlq.DISPLAY.CHLAUTH	READ	No check	–
Inquire Channel Initiator	hlq.DISPLAY.CHINIT	READ	No check	–
Inquire Channel Names	hlq.DISPLAY.CHANNEL	READ	No check	–
Inquire Channel Status	hlq.DISPLAY.CHSTATUS	READ	No check	–
Inquire Cluster Queue Manager	hlq.DISPLAY.CLUSQMGR	READ	No check	–
Inquire Connection	hlq.DISPLAY.CONNPCF	READ	No check	–
Inquire Group	hlq.DISPLAY.GROUP	READ	No check	–
Inquire Log	hlq.DISPLAY.LOG	READ	No check	–
Inquire Namelist	hlq.DISPLAY.NAMELIST	READ	No check	–
Inquire Namelist Names	hlq.DISPLAY.NAMELIST	READ	No check	–
Inquire Process	hlq.DISPLAY.PROCESS	READ	No check	–
Inquire Process Names	hlq.DISPLAY.PROCESS	READ	No check	–
Inquire Pub/Sub Status	hlq.DISPLAY.PUBSUB	READ	No check	–
Inquire Queue	hlq.DISPLAY.QUEUE	READ	No check	–
Inquire Queue Manager	hlq.DISPLAY.QMGR	READ	No check	–
Inquire Queue Names	hlq.DISPLAY.QUEUE	READ	No check	–
Inquire Queue Status	hlq.DISPLAY.QSTATUS	READ	No check	–
Inquire Security	hlq.DISPLAY.SECURITY	READ	No check	–
Inquire SMDS	hlq.DISPLAY.SMDS	READ	No check	–
Inquire SMDSCONN	hlq.DISPLAY.SMDSCONN	READ	No check	–
Inquire Storage Class	hlq.DISPLAY.STGCLASS	READ	No check	–
Inquire Storage Class Names	hlq.DISPLAY.STGCLASS	READ	No check	–
Inquire Subscription	hlq.INQUIRE.SUB	READ	No check	–
Inquire Subscription Status	hlq.INQUIRE.SBSTATUS	READ	No check	–
Inquire System	hlq.DISPLAY.SYSTEM	READ	No check	–
Inquire Topic	hlq.DISPLAY.TOPIC	READ	No check	–
Inquire Topic Names	hlq.DISPLAY.TOPIC	READ	No check	–
Inquire Topic Status	hlq.DISPLAY.TPSTATUS	READ	No check	–
Inquire Usage	hlq.DISPLAY.USAGE	READ	No check	–

Table 55. PCF commands, profiles, and their access levels (continued)

Command	Command profile for MQCMDS	Access level for MQCMDS	Command resource profile for MQADMIN or MXADMIN	Access level for MQADMIN or MXADMIN
Move Queue	hlq.MOVE.QLOCAL	ALTER	hlq.QUEUE.from-queue hlq.QUEUE.to-queue	ALTER
Ping Channel	hlq.PING.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
Recover CF Structure	hlq.RECOVER.CFSTRUCT	CONTROL	No check	–
Refresh Cluster	hlq.REFRESH.CLUSTER	ALTER	No check	–
Refresh Queue Manager	hlq.REFRESH.QMGR	ALTER	No check	–
Refresh Security	hlq.REFRESH.SECURITY	ALTER	No check	–
Reset CF Structure	hlq.RESET.CFSTRUCT	CONTROL	No check	–
Reset Channel	hlq.RESET.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
Reset Cluster	hlq.RESET.CLUSTER	CONTROL	No check	–
Reset Queue Manager	hlq.RESET.QMGR	CONTROL	No check	–
Reset Queue Statistics	hlq.RESET.QSTATS	CONTROL	hlq.QUEUE.queue	CONTROL
Reset SMDS	hlq.RESET.SMDS	CONTROL	No check	–
Resolve Channel	hlq.RESOLVE.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
Resume Queue Manager	hlq.RESUME.QMGR	CONTROL	No check	–
Resume Queue Manager Cluster	hlq.RESUME.QMGR	CONTROL	No check	–
Reverify Security	hlq.RVERIFY.SECURITY	ALTER	No check	–
Set Archive	hlq.SET.ARCHIVE	CONTROL	No check	–
Set Channel Authentication Record	hlq.SET.CHLAUTH	CONTROL	No check	–
Set Log	hlq.SET.LOG	CONTROL	No check	–
Set System	hlq.SET.SYSTEM	CONTROL	No check	–
Start Channel	hlq.START.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
Start Channel Initiator	hlq.START.CHINIT	CONTROL	No check	–
Start Channel Listener	hlq.START.LISTENER	CONTROL	No check	–
Start SMDS Connection	hlq.START.SMDSCONN	CONTROL	No check	–
Stop Channel	hlq.STOP.CHANNEL	CONTROL	hlq.CHANNEL.channel	CONTROL
Stop Channel Initiator	hlq.STOP.CHINIT	CONTROL	No check	–
Stop Channel Listener	hlq.STOP.LISTENER	CONTROL	No check	–
Stop SMDS Connection	hlq.STOP.SMDSCONN	CONTROL	No check	–
Suspend Queue Manager	hlq.SUSPEND.QMGR	CONTROL	No check	–
Suspend Queue Manager Cluster	hlq.SUSPEND.QMGR	CONTROL	No check	–
Notes:				
1. The hlq.TOPIC.topic resource refers to the Topic object derived from the TOPICSTR. For more details, see “Publish/subscribe security” on page 808				

Profiles for command resource security:

If you have not defined the command resource security switch profile, because you want security checking for resources associated with commands, you must add resource profiles for each resource to the appropriate class. The same security profiles control both MQSC and PCF commands.

If you have not defined the command resource security switch profile, `hlq.NO.CMD.RESC.CHECKS`, because you want security checking for resources associated with commands, you must:

- Add a resource profile in the **MQADMIN** class, if using uppercase profiles, for each resource.
- Add a resource profile in the **MXADMIN** class, if using mixed case profiles, for each resource.

The same security profiles control both MQSC and PCF commands.

Profiles for command resource security checking take the form:

```
hlq.type.resourcename
```

where `hlq` can be either `qmgr-name` (queue manager name) or `qsg-name` (queue-sharing group name).

A profile prefixed by the queue manager name controls access to the resources associated with commands on that queue manager. A profile prefixed by the queue-sharing group name controls access to the resources associated with commands on all queue managers within the queue-sharing group. This access can be overridden on an individual queue manager by defining a queue-manager level profile for that command resource on that queue manager.

If your queue manager is a member of a queue-sharing group and you are using both queue manager and queue-sharing group level security, WebSphere MQ checks for a profile prefixed by the queue manager name first. If it does not find one, it looks for a profile prefixed by the queue-sharing group name.

For example, the RACF profile name for command resource security checking against the model queue `CREDIT.WORTHY` in subsystem `CSQ1` is:

```
CSQ1.QUEUE.CREDIT.WORTHY
```

Because the profiles for all types of command resource are held in the **MQADMIN** class, the “type” part of the profile name is needed in the profile to distinguish between resources of different types that have the same name. The “type” part of the profile name can be `CHANNEL`, `QUEUE`, `TOPIC`, `PROCESS`, or `NAMELIST`. For example, a user might be authorized to define `hlq.QUEUE.PAYROLL.ONE`, but not authorized to define `hlq.PROCESS.PAYROLL.ONE`.

If the resource type is a queue, and the profile is a queue-sharing group level profile, it controls access to one or more local queues within the queue sharing group, or access to a single shared queue from any queue manager in the queue-sharing group.

MQSC commands, profiles, and their access levels shows, for each WebSphere MQ MQSC command, the profiles required for command security checking to be carried out, and the corresponding access level for each profile in the **MQCMDS** class.

PCF commands, profiles, and their access levels shows, for each WebSphere MQ PCF command, the profiles required for command security checking to be carried out, and the corresponding access level for each profile in the **MQCMDS** class.

Command resource security checking for alias queues and remote queues:

Alias queue and remote queues both provide indirection to another queue. Additional points apply when you consider security checking for these queues.

Alias queues

When you define an alias queue, command resource security checks are only performed against the name of the alias queue, not against the name of the target queue to which the alias resolves.

Alias queues can resolve to both local and remote queues. If you do not want to permit users access to certain local or remote queues, you must do both of the following:

1. Do not allow the users access to these local and remote queues.
2. Restrict the users from being able to define aliases for these queues. That is, prevent them from being able to issue DEFINE QALIAS and ALTER QALIAS commands.

Remote queues

When you define a remote queue, command resource security checks are performed only against the name of the remote queue. No checks are performed against the names of the queues specified in the RNAME or XMITQ attributes in the remote queue object definition.

Related reference:



START CHINIT (*WebSphere MQ V7.1 Reference*)

The RESLEVEL security profile

You can define a special profile in the MQADMIN or MXADMIN class to control the number of user IDs checked for API-resource security. This profile is called the RESLEVEL profile. How this profile affects API-resource security depends on how you access WebSphere MQ.

When an application tries to connect to WebSphere MQ, WebSphere MQ checks the access that the user ID associated with the connection has to a profile in the MQADMIN or MXADMIN class called:

hlq.RESLEVEL

Where hlq can be either ssid (subsystem ID) or qsg (queue-sharing group ID).

The user IDs associated with each connection type are:

- The user ID of the connecting task for batch connections
- The CICS address space user ID for CICS connections
- The IMS region address space user ID for IMS connections
- The channel initiator address space user ID for channel initiator connections

Attention: RESLEVEL is a very powerful option; it can cause the bypassing of all resource security checks for a particular connection.

If you do not have a RESLEVEL profile defined, you must be careful that no other profile in the MQADMIN class matches hlq.RESLEVEL. For example, if you have a profile in MQADMIN called hlq.** and no hlq.RESLEVEL profile, beware of the consequences of the hlq.** profile because it is used for the RESLEVEL check.

Define an hlq.RESLEVEL profile and set the UACC to NONE, rather than have no RESLEVEL profile at all. Have as few users or groups in the access list as possible. For details about how to audit RESLEVEL access, see “Auditing considerations on z/OS” on page 586.

If you are using queue manager level security only, WebSphere MQ performs RESLEVEL checks against the `qmgr-name.RESLEVEL` profile. If you are using queue-sharing group level security only, WebSphere MQ performs RESLEVEL checks against the `qsg-name.RESLEVEL` profile. If you are using a combination of both queue manager and queue-sharing group level security, WebSphere MQ first checks for the existence of a RESLEVEL profile at queue manager level. If it does not find one, it checks for a RESLEVEL profile at queue-sharing group level.

If it cannot find a RESLEVEL profile, WebSphere MQ enables checking of both the job and task (or alternate user) ID for a CICS or an IMS connection. For a batch connection, WebSphere MQ enables checking of the job (or alternate) user ID. For the channel initiator, WebSphere MQ enables checking of the channel user ID and the MCA (or alternate) user ID.

If there is a RESLEVEL profile, the level of checking depends on the environment and access level for the profile.

Remember that if your queue manager is a member of a queue-sharing group and you do not define this profile at queue-manager level, there might be one defined at queue-sharing group level that will affect the level of checking. To activate the checking of two user IDs, you define a RESLEVEL profile (prefixed with either the queue manager name or the queue-sharing group name) with a UACC(NONE) and ensure that the relevant users do not have access granted against this profile.

When you consider the access that the channel initiator's user ID has to RESLEVEL, remember that the connection established by the channel initiator is also the connection used by the channels. A setting that causes the bypassing of all resource security checks for the channel initiator's user ID effectively bypasses security checks for all channels. If the channel initiator's user ID access to RESLEVEL is something other than NONE, then only one user ID (for an access level of READ or UPDATE) or no user IDs (for an access level of CONTROL or ALTER) is checked for access. If you grant the channel initiator's user ID an access level other than NONE to RESLEVEL, be sure that you understand the effect of this setting on the security checks done for channels.

Using the RESLEVEL profile means that normal security audit records are not taken. For example, if you put UAUDIT on a user, the access to the `hlq.RESLEVEL` profile in MQADMIN is not audited.

If you use the RACF WARNING option on the `hlq.RESLEVEL` profile, no RACF warning messages are produced for profiles in the RESLEVEL class.

Security checking for report messages such as CODs are controlled by the RESLEVEL profile associated with the originating application. For example, if a batch job's userid has CONTROL or ALTER authority to a RESLEVEL profile, then all resource checking performed by the batch job are bypassed, including the security check of report messages.

If you change the RESLEVEL profile, users must disconnect and connect again before the change takes place. (This includes stopping and restarting the channel initiator if the access that the distributed queuing address space user ID has to the RESLEVEL profile is changed.)

To switch RESLEVEL auditing off, use the RESAUDIT system parameter.

RESLEVEL and batch connections:

By default, when a WebSphere MQ resource is being accessed through batch and batch-type connections, the user must be authorized to access that resource for the particular operation. You can bypass the security check by setting up an appropriate RESLEVEL definition.

Whether the user is checked or not is based on the user ID used at connect time, the same user ID used for the connection check.

For example, you can set up RESLEVEL so that when a user you trust accesses certain resources through a batch connection, no API-resource security checks are done; but when a user you do not trust tries to access the same resources, security checks are carried out as normal. You should set up RESLEVEL checking to bypass API-resource security checks only when you sufficiently trust the user and the programs run by that user.

The following table shows the checks made for batch connections.

Table 56. Checks made at different RACF access levels for batch connections

RACF access level	Level of checking
NONE	Resource checks performed
READ	Resource checks performed
UPDATE	Resource checks performed
CONTROL	No check.
ALTER	No check.

RESLEVEL and system functions:

The application of RESLEVEL to the operation and control panels, and to CSQUTIL.

The operation and control panels and the CSQUTIL utility are batch-type applications that make requests to the queue manager's command server, and so they are subject to the considerations described in "RESLEVEL and batch connections." You can use RESLEVEL to bypass security checking for the SYSTEM.COMMAND.INPUT and SYSTEM.COMMAND.REPLY.MODEL queues that they use, but not for the dynamic queues SYSTEM.CSQXCMD.*, SYSTEM.CSQOREXX.*, and SYSTEM.CSQUTIL.*.

The command server is an integral part of the queue manager and so does not have connection or RESLEVEL checking associated with it. To maintain security, therefore, the command server must confirm that the user ID of the requesting application has authority to open the queue being used for replies. For the operations and control panels, this is SYSTEM.CSQOREXX.*. For CSQUTIL, it is SYSTEM.CSQUTIL.*. Users must be authorized to use these queues, as described in "System queue security" on page 539, in addition to any RESLEVEL authorization they are given.

For other applications using the command server, it is the queue they name as their reply-to queue. Such other applications might deceive the command server into placing messages on unauthorized queues by passing (in the message context) a more trusted user ID than its own to the command server. To prevent this, use a CONTEXT profile to protect the identity context of messages placed on SYSTEM.COMMAND.INPUT.

RESLEVEL and CICS connections:

By default, when an API-resource security check is made on a CICS connection, two user IDs are checked. You can change which user IDs are checked by setting up a RESLEVEL profile.

The first user ID checked is that of the CICS address space. This is the user ID on the job card of the CICS job, or the user ID assigned to the CICS started task by the z/OS STARTED class or the started procedures table. (It is not the CICS DFLTUSER.)

The second user ID checked is the user ID associated with the CICS transaction.

If one of these user IDs does not have access to the resource, the request fails with a completion code of MQRC_NOT_AUTHORIZED. Both the CICS address space user ID and the user ID of the person running the CICS transaction must have access to the resource at the correct level.

How RESLEVEL can affect the checks made

Depending on how you set up your RESLEVEL profile, you can change which user IDs are checked when access to a resource is requested. See Table 57 for more information.

The user IDs checked depend on the user ID used at connection time, that is, the CICS address space user ID. This control enables you to bypass API-resource security checking for WebSphere MQ requests coming from one system (for example, a test system, TESTCICS,) but to implement them for another (for example, a production system, PRODCICS).

Note: If you set up your CICS address space user ID with the “trusted” attribute in the STARTED class or the RACF started procedures table ICHRIN03, this overrides any user ID checks for the CICS address space established by the RESLEVEL profile for your queue manager (that is, the queue manager does not perform the security checks for the CICS address space). For more information, see the *CICS Transaction Server for z/OS V3.2 RACF Security Guide*.

The following table shows the checks made for CICS connections.

Table 57. Checks made at different RACF access levels for CICS connections

RACF access level	Level of checking
NONE	WebSphere MQ checks the CICS address space user ID and the transaction user ID.
READ	WebSphere MQ checks the CICS address space user ID only.
UPDATE	If the transaction is defined to CICS with RESSEC(YES), WebSphere MQ checks the CICS address space user ID and the transaction user ID.
UPDATE	If the transaction is defined to CICS with RESSEC(NO), WebSphere MQ checks the CICS address space user ID only.
CONTROL or ALTER	WebSphere MQ does not check any user IDs.

RESLEVEL and IMS connections:

By default, when an API-resource security check is made for an IMS connection, two user IDs are checked. You can change which user IDs are checked by setting up a RESLEVEL profile.

By default, when an API-resource security check is made for an IMS connection, two user IDs are checked to see if access is allowed to the resource.

The first user ID checked is that of the address space of the IMS region. This is taken from either the USER field from the job card or the user ID assigned to the region from the z/OS STARTED class or the started procedures table (SPT).

The second user ID checked is associated with the work being done in the dependent region. It is determined according to the type of the dependent region as shown in How the second user ID is determined for the IMS(tm) connection.

If either the first or second IMS user ID does not have access to the resource, the request fails with a completion code of MQRC_NOT_AUTHORIZED.

The setting of WebSphere MQ RESLEVEL profiles cannot alter the user ID under which IMS transactions are scheduled from the IBM-supplied MQ-IMS trigger monitor program CSQQTRMN. This user ID is the PSBNAME of that trigger monitor, which by default is CSQQTRMN.

How RESLEVEL can affect the checks made

Depending on how you set up your RESLEVEL profile, you can change which user IDs are checked when access to a resource is requested. The possible checks are:

- Check the IMS region address space user ID and the second user ID or alternate user ID.
- Check IMS region address space user ID only.
- Do not check any user IDs.

The following table shows the checks made for IMS connections.

Table 58. Checks made at different RACF access levels for IMS connections

RACF access level	Level of checking
NONE	Check the IMS address space user ID and the IMS second user ID or alternate user ID.
READ	Check the IMS address space user ID.
UPDATE	Check the IMS address space user ID.
CONTROL	No check.
ALTER	No check.

RESLEVEL and the channel initiator connection:

By default, when an API-resource security check is made by the channel initiator, two user IDs are checked. You can change which user IDs are checked by setting up a RESLEVEL profile.

By default, when an API-resource security check is made by the channel initiator, two user IDs are checked to see if access is allowed to the resource.

The user IDs checked can be that specified by the MCAUSER channel attribute, that received from the network, that of the channel initiator address space, or the alternate user ID for the message descriptor. Which user IDs are checked depends on the communication protocol you are using and the setting of the PUTAUT channel attribute. See "User IDs used by the channel initiator" on page 571 for more information.

If one of these user IDs does not have access to the resource, the request fails with a completion code of MQRC_NOT_AUTHORIZED.

How RESLEVEL can affect the checks made

Depending on how you set up your RESLEVEL profile, you can change which user IDs are checked when access to a resource is requested, and how many are checked.

The following table shows the checks made for the channel initiator's connection, and for all channels since they use this connection.

Table 59. Checks made at different RACF access levels for channel initiator connections

RACF access level	Level of checking
NONE	Check two user IDs.
READ	Check one user ID.
UPDATE	Check one user ID.
CONTROL	No check.
ALTER	No check.
Note: See "User IDs used by the channel initiator" on page 571 for a definition of the user IDs checked	

RESLEVEL and intra-group queuing:

By default, when an API-resource security check is made by the intra-group queuing agent, two user IDs are checked. You can change which user IDs are checked by setting up a RESLEVEL profile.

By default, when an API-resource security check is made by the intra-group queuing agent, two user IDs are checked to see if access is allowed to the resource.

The user IDs checked can be the user ID determined by the IGQUSER attribute of the receiving queue manager, the user ID of the queue manager within the queue-sharing group that put the message on to the SYSTEM.QSG.TRANSMIT.QUEUE, or the alternate user ID specified in the *UserIdentifier* field of the message descriptor of the message. See "User IDs used by the intra-group queuing agent" on page 575 for more information.

Because the intra-group queuing agent is an internal queue manager task, it does not issue an explicit connect request and runs under the user ID of the queue manager. The intra-group queuing agent starts at queue manager initialization. During the initialization of the intra-group queuing agent, WebSphere MQ checks the access that the user ID associated with the queue manager has to a profile in the MQADMIN class called:

hlq.RESLEVEL

This check is always performed unless the hlq.NO.SUBSYS.SECURITY switch has been set.

If there is no RESLEVEL profile, WebSphere MQ enables checking for two user IDs. If there is a RESLEVEL profile, the level of checking depends on the access level granted to the user ID of the queue manager for the profile. Checks made at different RACF(r) access levels for the intra-group queuing agent shows the checks made for the intra-group queuing agent.

Table 60. Checks made at different RACF access levels for the intra-group queuing agent

RACF access level	Level of checking
NONE	Check two user IDs.
READ	Check one user ID.
UPDATE	Check one user ID.
CONTROL	No check.
ALTER	No check.
Note: See "User IDs used by the intra-group queuing agent" on page 575 for a definition of the user IDs checked	

If the permissions granted to the RESLEVEL profile for the queue manager's user ID are changed, the intra-group queuing agent must be stopped and restarted to pick up the new permissions. Because there is no way to independently stop and restart the intra-group queuing agent, the queue manager must be stopped and restarted to achieve this.

RESLEVEL and the user IDs checked:

Example of setting a RESLEVEL profile and granting access to it.

User ID checking against profile name for batch connections through User IDs checked against profile name for LU 6.2 and TCP/IP server-connection channels show how RESLEVEL affects which user IDs are checked for different MQI requests.

For example, you have a queue manager called QM66 with the following requirements:

- User WS21B is to be exempt from resource security.
- CICS started task WXNCICS running under address space user ID CICSWXN is to perform full resource checking only for transactions defined with RESSEC(YES).

To define the appropriate RESLEVEL profile, issue the following RACF command:

```
RDEFINE MQADMIN QM66.RESLEVEL UACC(NONE)
```

Then give the users access to this profile, using the following commands:

```
PERMIT QM66.RESLEVEL CLASS(MQADMIN) ID(WS21B) ACCESS(CONTROL)  
PERMIT QM66.RESLEVEL CLASS(MQADMIN) ID(CICSWXN) ACCESS(UPDATE)
```

If you make these changes while the user IDs are connected to queue manager QM66, the users must disconnect and connect again before the change takes place.

If subsystem security is not active when a user connects but, while this user is still connected, subsystem security becomes active, full resource security checking is applied to the user. The user must reconnect to get the correct RESLEVEL processing.

User IDs for security checking

WebSphere MQ initiates security checks based on user IDs associated with users, terminals, applications, and other resources. This collection of topics lists which user IDs are used for each type of security check.

User IDs for connection security:

The user ID used for connection security depends on the type of connection.

Connection type	User ID contents
Batch connection	The user ID of the connecting task. For example: <ul style="list-style-type: none"> • The TSO user ID • The user ID assigned to a batch job by the USER JCL parameter • The user ID assigned to a started task by the STARTED class or the started procedures table
CICS connection	The CICS address space user ID.
IMS connection	The IMS region address space user ID.
Channel initiator connection	The channel initiator address space user ID.

User IDs for command security and command resource security:

The user ID used for command security or command resource security depends on where the command is issued from.

Issued from...	User ID contents
CSQINP1, CSQINP2, or CSQINPT	No check is made.
System command input queue	The user ID found in the <i>UserIdentifier</i> of the message descriptor of the message that contains the command. If the message does not contain a <i>UserIdentifier</i> , a user ID of blanks is passed to the security manager.
Console	The user ID signed onto the console. If the console is not signed on, the default user ID set by the CMDUSER system parameter in CSQ6SYSP. To issue commands from a console, the console must have the z/OS SYS AUTHORITY attribute.
SDSF/TSO console	TSO or job user ID.
Operations and control panels	TSO user ID. If you are going to use the operations and control panels, you must have the appropriate authority to issue the commands corresponding to the actions that you choose. In addition, you must have READ access to all the hlq.DISPLAY.object profiles in the MQCMDS class because the panels use the various DISPLAY commands to gather the information that they present.
MGCRE	If MGCRE is used with UTOKEN, the user ID in the UTOKEN. If MGCRE is issued without the UTOKEN, the TSO or job user ID is used.
CSQUTIL	Job user ID.
CSQINPX	User ID of the channel initiator address space.

User IDs for resource security (MQOPEN, MQSUB, and MQPUT1):

This information shows the contents of the user IDs for normal and alternate user IDs for each type of connection. The number of checks is defined by the RESLEVEL profile. The user ID checked is that used for **MQOPEN**, **MQSUB**, or **MQPUT1** calls.

Note: All user ID fields are checked exactly as they are received. No conversions take place, and, for example, three user ID fields containing “Bob”, “BOB”, and “bob” are not equivalent.

User IDs checked for batch connections:

The user ID checked for a batch connection depends on how the task is run and whether an alternate user ID has been specified.

Table 61. User ID checking against profile name for batch connections

Alternate user ID specified on open?	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueName profile	hlq.resourcename profile
No	–	JOB	JOB
Yes	JOB	JOB	ALT

Key:

ALT Alternate user ID.

JOB

- The user ID of a TSO or USS sign-on.
- The user ID assigned to a batch job.
- The user ID assigned to a started task by the STARTED class or the started procedures table.
- The user ID associated with the executing Db2 stored procedure

A Batch job is performing an **MQPUT1** to a queue called Q1 with RESLEVEL set to READ and alternate user ID checking turned off.

Checks made at different RACF(r) access levels for batch connections and User ID checking against profile name for batch connections show that the job user ID is checked against profile hlq.Q1.

User IDs checked for CICS connections:

The user IDs checked for CICS connections depend on whether one or two checks are to be carried out, and whether an alternate user ID is specified.

Table 62. User ID checking against profile name for CICS-type user IDs

Alternate user ID specified on open?	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueName profile	hlq.resourcename profile
No, 1 check	–	ADS	ADS
No, 2 checks	–	ADS+TXN	ADS+TXN
Yes, 1 check	ADS	ADS	ADS
Yes, 2 checks	ADS+TXN	ADS+TXN	ADS+ALT

Table 62. User ID checking against profile name for CICS-type user IDs (continued)

Alternate user ID specified on open?	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueuname profile	hlq.resourcename profile
Key:			
ALT	Alternate user ID		
ADS	The user ID associated with the CICS batch job or, if CICS is running as a started task, through the STARTED class or the started procedures table.		
TXN	The user ID associated with the CICS transaction. This is normally the user ID of the terminal user who started the transaction. It can be the CICS DFLTUSER, a PRESET security terminal, or a manually signed-on user.		

Determine the user IDs checked for the following conditions:

- The RACF access level to the RESLEVEL profile, for a CICS address space user ID, is set to NONE.
- An **MQOPEN** call is made against a queue with MQOO_OUTPUT and MQOO_PASS_IDENTITY_CONTEXT.

First, see how many CICS user IDs are checked based on the CICS address space user ID access to the RESLEVEL profile. From Table 57 on page 564 in topic “RESLEVEL and CICS connections” on page 564, two user IDs are checked if the RESLEVEL profile is set to NONE. Then, from Table 62 on page 569 on, these checks are carried out:

- The hlq.ALTERNATE.USER.userid profile is not checked.
- The hlq.CONTEXT.queueuname profile is checked with both the CICS address space user ID and the CICS transaction user ID.
- The hlq.resourcename profile is checked with both the CICS address space user ID and the CICS transaction user ID.

This means that four security checks are made for this **MQOPEN** call.

User IDs checked for IMS connections:

The user IDs checked for IMS connections depend on whether one or two checks are to be performed, and whether an alternate user ID is specified. If a second user ID is checked, it depends on the type of dependent region and on which user IDs are available.

Table 63. User ID checking against profile name for IMS-type user IDs

Alternate user ID specified on open?	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueuname profile	hlq.resourcename profile
<i>No, 1 check</i>	–	REG	REG
<i>No, 2 checks</i>	–	REG+SEC	REG+SEC
<i>Yes, 1 check</i>	REG	REG	REG
<i>Yes, 2 checks</i>	REG+SEC	REG+SEC	REG+ALT
Key:			
ALT	Alternate user ID.		
REG	The user ID is normally set through the STARTED class or the started procedures table or, if IMS is running, from a submitted job, by the USER JCL parameter.		
SEC	The second user ID is associated with the work being done in a dependent region. It is determined according to Table 64 on page 571.		

Table 64. How the second user ID is determined for the IMS connection

Types of dependent region	Hierarchy for determining the second user ID
<ul style="list-style-type: none"> BMP message driven and successful GET UNIQUE issued. IFP and GET UNIQUE issued. MPP. 	User ID associated with the IMS transaction if the user is signed on. LTERM name if available. PSBNAME.
<ul style="list-style-type: none"> BMP message driven and successful GET UNIQUE not issued. BMP not message driven. IFP and GET UNIQUE not issued. 	User ID associated with the IMS dependent region address space if this is not all blanks or all zeros. PSBNAME.

User IDs used by the channel initiator:

This collection of topics describes the user IDs used and checked for receiving channels and for client MQI requests issued over server-connection channels. Information is provided for TCP/IP and for LU6.2

You can use the PUTAUT parameter of the receiving channel definition to determine the type of security checking used. To get consistent security checking throughout your WebSphere MQ network, you can use the ONLYMCA and ALTMCA options.

You can use the DISPLAY CHSTATUS command to determine the user identifier used by the MCA.

Related reference:

 DISPLAY CHSTATUS (*WebSphere MQ V7.1 Reference*)

 DEFINE CHANNEL (*WebSphere MQ V7.1 Reference*)

Receiving channels using TCP/IP:

The user IDs checked depend on the PUTAUT option of the channel and on whether one or two checks are to be performed.

Table 65. User IDs checked against profile name for TCP/IP channels

PUTAUT option specified on receiver or requester channel	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueprofile	hlq.resourcename profile
DEF, 1 check	-	CHL	CHL
DEF, 2 checks	-	CHL + MCA	CHL + MCA
CTX, 1 check	CHL	CHL	CHL
CTX, 2 checks	CHL + MCA	CHL + MCA	CHL + ALT
ONLYMCA, 1 check	-	MCA	MCA
ONLYMCA, 2 checks	-	MCA	MCA
ALTMCA, 1 check	MCA	MCA	MCA
ALTMCA, 2 checks	MCA	MCA	MCA + ALT

Table 65. User IDs checked against profile name for TCP/IP channels (continued)

PUTAUT option specified on receiver or requester channel	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueprofile	hlq.resourcename profile
<p>Key:</p> <p>MCA (MCA user ID) The user ID specified for the MCAUSER channel attribute at the receiver; if blank, the channel initiator address space user ID of the receiver or requester side is used.</p> <p>CHL (Channel user ID) On TCP/IP, security is not supported by the communication system for the channel. If the Secure Sockets Layer (SSL) is being used and a digital certificate has been flowed from the partner, the user ID associated with this certificate (if installed), or the user ID associated with a matching filter found by using RACF's Certificate Name Filter (CNF), is used. If no associated user ID is found, or if SSL is not being used, the user ID of the channel initiator address space of the receiver or requester end is used as the channel user ID on channels defined with the PUTAUT parameter set to DEF or CTX. Note: The use of RACF's Certificate Name Filter (CNF) allows you to assign the same RACF user ID to multiple remote users, for example all the users in the same organization unit, who would naturally all have the same security authority. This means that the server does not have to have a copy of the certificate of every possible remote end user across the world and greatly simplifies certificate management and distribution. If the PUTAUT parameter is set to ONLYMCA or ALTMCA for the channel, the channel user ID is ignored and the MCA user ID of the receiver or requester is used. This also applies to TCP/IP channels using SSL.</p> <p>ALT (Alternate user ID) The user ID from the context information (that is, the <i>UserIdentifier</i> field) within the message descriptor of the message. This user ID is moved into the <i>AlternateUserID</i> field in the object descriptor before an MQOPEN or MQPUT1 call is issued for the target destination queue.</p>			

Receiving channels using LU 6.2:

The user IDs checked depend on the PUTAUT option of the channel and on whether one or two checks are to be performed.

Table 66. User IDs checked against profile name for LU 6.2 channels

PUTAUT option specified on receiver or requester channel	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueprofile	hlq.resourcename profile
DEF, 1 check	-	CHL	CHL
DEF, 2 checks	-	CHL + MCA	CHL + MCA
CTX, 1 check	CHL	CHL	CHL
CTX, 2 checks	CHL + MCA	CHL + MCA	CHL + ALT
ONLYMCA, 1 check	-	MCA	MCA
ONLYMCA, 2 checks	-	MCA	MCA
ALTMCA, 1 check	MCA	MCA	MCA
ALTMCA, 2 checks	MCA	MCA	MCA + ALT

Table 66. User IDs checked against profile name for LU 6.2 channels (continued)

PUTAUT option specified on receiver or requester channel	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueprofile	hlq.resourcename profile
<p>Key:</p> <p>MCA (MCA user ID) The user ID specified for the MCAUSER channel attribute at the receiver; if blank, the channel initiator address space user ID of the receiver or requester side is used.</p> <p>CHL (Channel user ID)</p> <p>Requester-server channels If the channel is started from the requester, there is no opportunity to receive a network user ID (the channel user ID). If the PUTAUT parameter is set to DEF or CTX on the requester channel, the channel user ID is that of the channel initiator address space of the requester because no user ID is received from the network. If the PUTAUT parameter is set to ONLYMCA or ALTMCA, the channel user ID is ignored and the MCA user ID of the requester is used.</p> <p>Other channel types If the PUTAUT parameter is set to DEF or CTX on the receiver or requester channel, the channel user ID is the user ID received from the communications system when the channel is initiated.</p> <ul style="list-style-type: none"> • If the sending channel is on z/OS, the channel user ID received is the channel initiator address space user ID of the sender. • If the sending channel is on a different platform (for example, AIX or HP-UX), the channel user ID received is typically provided by the USERID parameter of the channel definition. <p>If the user ID received is blank, or no user ID is received, a channel user ID of blanks is used.</p> <p>ALT (Alternate user ID) The user ID from the context information (that is, the <i>UserIdentifier</i> field) within the message descriptor of the message. This user ID is moved into the <i>AlternateUserID</i> field in the object descriptor before an MQOPEN or MQPUT1 call is issued for the target destination queue.</p>			

Client MQI requests:

Various user IDs can be used, depending on which user IDs and environment variables have been set. These user IDs are checked against various profiles, depending on the PUTAUT option used and whether an alternate user ID is specified.

This section describes the user IDs checked for client MQI requests issued over server-connection channels for TCP/IP and LU 6.2. The MCA user ID and channel user ID are as for the TCP/IP and LU 6.2 channels described in the previous sections.

For server-connection channels, the user ID received from the client is used if the MCAUSER attribute is blank.

See “Access control for clients” on page 432 for more information.

For client MQOPEN, MQSUB, and MQPUT1 requests, use the following rules to determine the profile that is checked:

- If the request specifies alternate-user authority, a check is made against the *hlq.ALTERNATE.USER.userid* profile.
- If the request specifies context authority, a check is made against the *hlq.CONTEXT.queueprofile* profile.
- For all MQOPEN, MQSUB, and MQPUT1 requests, a check is made against the *hlq.resourcename* profile.

When you have determined which profiles are checked, use the following table to determine which user IDs are checked against these profiles.

Table 67. User IDs checked against profile name for LU 6.2 and TCP/IP server-connection channels

PUTAUT option specified on server-connection channel	Alternate user ID specified on open?	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queuname profile	hlq.resourcename profile
DEF, 1 check	No	-	CHL	CHL
DEF, 1 check	Yes	CHL	CHL	CHL
DEF, 2 checks	No	-	CHL + MCA	CHL + MCA
DEF, 2 checks	Yes	CHL + MCA	CHL + MCA	CHL + ALT
ONLYMCA, 1 check	No	-	MCA	MCA
ONLYMCA, 1 check	Yes	MCA	MCA	MCA
ONLYMCA, 2 checks	No	-	MCA	MCA
ONLYMCA, 2 checks	Yes	MCA	MCA	MCA + ALT

Key:

MCA (MCA user ID)

The user ID specified for the MCAUSER channel attribute at the server-connection; if blank, the channel initiator address space user ID is used.

CHL (Channel user ID)

On TCP/IP, security is not supported by the communication system for the channel. If the Secure Sockets Layer (SSL) is being used and a digital certificate has been flowed from the partner, the user ID associated with this certificate (if installed), or the user ID associated with a matching filter found by using RACF Certificate Name Filtering (CNF), is used. If no associated user ID is found, or if SSL is not being used, the user ID of the channel initiator address space is used as the channel user ID on channels defined with the PUTAUT parameter set to DEF or CTX.

Note: The use of RACF Certificate Name Filtering (CNF) allows you to assign the same RACF user ID to multiple remote users, for example all the users in the same organization unit, who would naturally all have the same security authority. This means that the server does not have to have a copy of the certificate of every possible remote end user across the world and greatly simplifies certificate management and distribution.

If the PUTAUT parameter is set to ONLYMCA or ALTMCA for the channel, the channel user ID is ignored and the MCA user ID of the server-connection channel is used. This also applies to TCP/IP channels using SSL.

ALT (Alternate user ID)

The user ID from the context information (that is, the *UserIdentifier* field) within the message descriptor of the message. This user ID is moved into the *AlternateUserID* field in the object or subscription descriptor before an MQOPEN, MQSUB or MQPUT1 call is issued on behalf of the client application.

Channel initiator example:

An example of how user IDs are checked against RACF profiles.

A user performs an **MQPUT1** operation to a queue on queue manager QM01 that resolves to a queue called QB on queue manager QM02. The message is sent on a TCP/IP channel called QM01.TO.QM02. RESLEVEL is set to NONE, and the open is performed with alternate user ID and context checking. The receiver channel definition has PUTAUT(CTX) and the MCA user ID is set. Which user IDs are used on the receiving channel to put the message to queue QB?

Answer: Table 59 on page 566 shows that two user IDs are checked because RESLEVEL is set to NONE.

Table 65 on page 571 shows that, with PUTAUT set to CTX and 2 checks, the following user IDs are checked:

- The channel initiator user ID and the MCAUSER user ID are checked against the hlq.ALTERNATE.USER.userid profile.
- The channel initiator user ID and the MCAUSER user ID are checked against the hlq.CONTEXT.queueusername profile.
- The channel initiator user ID and the alternate user ID specified in the message descriptor (MQMD) are checked against the hlq.Q2 profile.

User IDs used by the intra-group queuing agent:

The user IDs that are checked when the intra-group queuing agent opens destination queues are determined by the values of the IGQAUT and IGQUSER queue manager attributes.

The possible user IDs are:

Intra-group queuing user ID (IGQ)

The user ID determined by the IGQUSER attribute of the receiving queue manager. If this is set to blanks, the user ID of the receiving queue manager is used. However, because the receiving queue manager has authority to access all queues defined to it, security checks are not performed for the receiving queue manager's user ID. In this case:

- If only one user ID is to be checked and the user ID is that of the receiving queue manager, no security checks take place. This can occur when IGAUT is set to ONLYIGQ or ALTIGQ.
- If two user IDs are to be checked and one of the user IDs is that of the receiving queue manager, security checks take place for the other user ID only. This can occur when IGAUT is set to DEF, CTX, or ALTIGQ.
- If two user IDs are to be checked and both user IDs are that of the receiving queue manager, no security checks take place. This can occur when IGAUT is set to ONLYIGQ.

Sending queue manager user ID (SND)

The user ID of the queue manager within the queue-sharing group that put the message on to the SYSTEM.QSG.TRANSMIT.QUEUE.

Alternate user ID (ALT)

The user ID specified in the *UserIdentifier* field in the message descriptor of the message.

Table 68. User IDs checked against profile name for intra-group queuing

IGQAUT option specified on receiving queue manager	hlq.ALTERNATE.USER.userid profile	hlq.CONTEXT.queueuname profile	hlq.resourcename profile
DEF, 1 check	–	SND	SND
DEF, 2 checks	–	SND +IGQ	SND +IGQ
CTX, 1 check	SND	SND	SND
CTX, 2 checks	SND + IGQ	SND +IGQ	SND + ALT
ONLYIGQ, 1 check	–	IGQ	IGQ
ONLYIGQ, 2 checks	–	IGQ	IGQ
ALTIGQ, 1 check	–	IGQ	IGQ
ALTIGQ, 2 checks	IGQ	IGQ	IGQ + ALT

Key:

ALT Alternate user ID.

IGQ IGQ user ID.

SND Sending queue manager user ID.

Blank user IDs and UACC levels:

If a blank user ID occurs, a RACF undefined user is signed on. Do not grant wide-ranging access to the undefined user.

Blank user IDs can exist when a user is manipulating messages using context or alternate-user security, or when WebSphere MQ is passed a blank user ID. For example, a blank user ID is used when a message is written to the system-command input queue without context.

Note: A user ID of "* " (that is, an asterisk character followed by seven spaces) is treated as an undefined user ID.

WebSphere MQ passes the blank user ID to RACF and a RACF undefined user is signed on. All security checks then use the universal access (UACC) for the relevant profile. Depending on how you have set your access levels, the UACC might give the undefined user a wide-ranging access.

For example, if you issue this RACF command from TSO:

```
RDEFINE MQQUEUE Q.AVAILABLE.TO.EVERYONE UACC(UPDATE)
```

you define a profile that enables both z/OS-defined user IDs (that have not been put in the access list) and the RACF undefined user ID to put messages on, and get messages from, that queue.

To protect against blank user IDs you must plan your access levels carefully, and limit the number of people who can use context and alternate-user security. You must prevent people using the RACF undefined user ID from getting access to resources that they must not access. However, at the same time, you must allow access to people with defined user IDs. To do this, you can specify a user ID of asterisk (*) in a RACF command PERMIT, giving access to resources for all defined user IDs. Therefore all undefined user IDs (such as "* ") are denied access. For example, these RACF commands prevent the RACF undefined user ID from gaining access to the queue to put or get messages:

```
RDEFINE MQQUEUE Q.AVAILABLE.TO.RACF.DEFINED.USERS.ONLY UACC(NONE)
PERMIT Q.AVAILABLE.TO.RACF.DEFINED.USERS.ONLY CLASS(MQQUEUE) ACCESS(UPDATE) ID(*)
```

WebSphere MQ for z/OS security management

WebSphere MQ uses an in-storage table to hold information relating to each user and the access requests made by each user. To manage this table efficiently and to reduce the number of requests made from WebSphere MQ to the external security manager (ESM), a number of controls are available.

These controls are available through both the operations and control panels and WebSphere MQ commands.

User ID reverification:

If the RACF definition of a user who is using WebSphere MQ resources has been changed—for example, by connecting the user to a new group—you can tell the queue manager to sign this user on again the next time it tries to access a WebSphere MQ resource. You can do this by using the WebSphere MQ command RVERIFY SECURITY.

- User HX0804 is getting and putting messages to the PAYROLL queues on queue manager PRD1. However HX0804 now requires access to some of the PENSION queues on the same queue manager (PRD1).
- The data security administrator connects user HX0804 to the RACF group that allows access to the PENSION queues.
- So that HX0804 can access the PENSION queues immediately—that is, without shutting down queue manager PRD1, or waiting for HX0804 to time out—you must use the WebSphere MQ command:
RVERIFY SECURITY(HX0804)

Note: If you turn off user ID timeout for long periods of time (days or even weeks) while the queue manager is running, you must remember to run the RVERIFY SECURITY command for any users that have been revoked or deleted in that time.

User ID timeouts:

You can make WebSphere MQ sign a user off a queue manager after a period of inactivity.

When a user accesses a WebSphere MQ resource, the queue manager tries to sign this user on to the queue manager (if subsystem security is active). This means that the user is authenticated to the ESM. This user remains signed on to WebSphere MQ until either the queue manager is shut down, or until the user ID is *timed out* (the authentication lapses) or reverified (reauthenticated).

When a user is timed out, the user ID is *signed off* within the queue manager and any security-related information retained for this user is discarded. The signing on and off of the user within the queue manager is not apparent to the application program or to the user.

Users are eligible for timeout when they have not used any WebSphere MQ resources for a predetermined amount of time. This time period is set by the MQSC ALTER SECURITY command.

Two values can be specified in the ALTER SECURITY command:

TIMEOUT

The time period in minutes that an unused user ID and its associated resources can remain within the WebSphere MQ queue manager.

INTERVAL

The time period in minutes between checks for user IDs and their associated resources, to determine whether the *TIMEOUT* has expired.

For example, if the *TIMEOUT* value is 30 and the *INTERVAL* value is 10, every 10 minutes WebSphere MQ checks user IDs and their associated resources to determine whether any have not been used for 30 minutes. If a timed-out user ID is found, that user ID is signed off within the queue manager. If any

timed-out resource information associated with non-timed-out user IDs is found, that resource information is discarded. If you do not want to time out user IDs, set the *INTERVAL* value to zero. However, if the *INTERVAL* value is zero, storage occupied by user IDs and their associated resources is not freed until you issue a **REFRESH SECURITY** or **RVERIFY SECURITY** command.

Tuning this value can be important if you have many one-off users. If you set small interval and timeout values, resources that are no longer required are freed.

Note: If you use values for *INTERVAL* or *TIMEOUT* other than the defaults, you must reenter the command at every queue manager startup. You can do this automatically by putting the **ALTER SECURITY** command in the CSQINP1 data set for that queue manager.

Related reference:



ALTER SECURITY (*WebSphere MQ V7.1 Reference*)

Refreshing queue manager security on z/OS:

IBM WebSphere MQ for z/OS caches RACF data to improve performance. When you change certain security classes, you must refresh this cached information. Refresh security infrequently, for performance reasons. You can also choose to refresh only SSL security information.

When a queue is opened for the first time (or for the first time since a security refresh) WebSphere MQ performs a RACF check to obtain the user's access rights and places this information in the cache. The cached data includes user IDs and resources on which security checking has been performed. If the queue is opened again by the same user, the presence of the cached data means that WebSphere MQ does not have to issue RACF checks, which improves performance. The action of a security refresh is to discard any cached security information and so force WebSphere MQ to make a new check against RACF. Whenever you add, change or delete a RACF resource profile that is held in the MQADMIN, MXADMIN, MQPROC, MXPROC, MQQUEUE, MXQUEUE, MQNLIST, MXNLIST, or MXTOPIC class, you must tell the queue managers that use this class to refresh the security information that they hold. To do this, issue the following commands:

- The RACF SETROPTS RACLIST(classname) REFRESH command to refresh at the RACF level.
- The IBM WebSphere MQ REFRESH SECURITY command to refresh the security information held by the queue manager. This command needs to be issued by each queue manager that accesses the profiles that have changed. If you have a queue-sharing group, you can use the command scope attribute to direct the command to all the queue managers in the group.

If you are using generic profiles in any of the IBM WebSphere MQ classes, you must also issue normal RACF refresh commands if you change, add, or delete any generic profiles. For example, SETROPTS GENERIC(classname) REFRESH.

However, if a RACF resource profile is added, changed or deleted, and the resource to which it applies has not yet been accessed (so no information is cached), IBM WebSphere MQ uses the new RACF information without a REFRESH SECURITY command being issued.

If RACF auditing is turned on, (for example, by using the RACF RALTER AUDIT(access-attempt (audit_access_level)) command), no caching takes place, and therefore WebSphere MQ refers directly to the RACF dataspace for every check. Changes are therefore picked up immediately and REFRESH SECURITY is not necessary to access the changes. You can confirm whether RACF auditing is on by using the RACF RLIST command. For example, you could issue the command

```
RLIST MQQUEUE (qmgr.SYSTEM.COMMAND.INPUT) GEN
```

and receive the results

```

CLASS      NAME
-----
MQQUEUE   QP*.SYSTEM.COMMAND.*.** (G)
          AUDITING
          -----
          FAILURES(READ)

```

This indicates that auditing is set on. For more information, refer to the *z/OS Security Server RACF Auditor's Guide* and the *z/OS Security Server RACF Command Language Reference*.

Figure 83 summarizes the situations in which security information is cached and in which cached information is used.

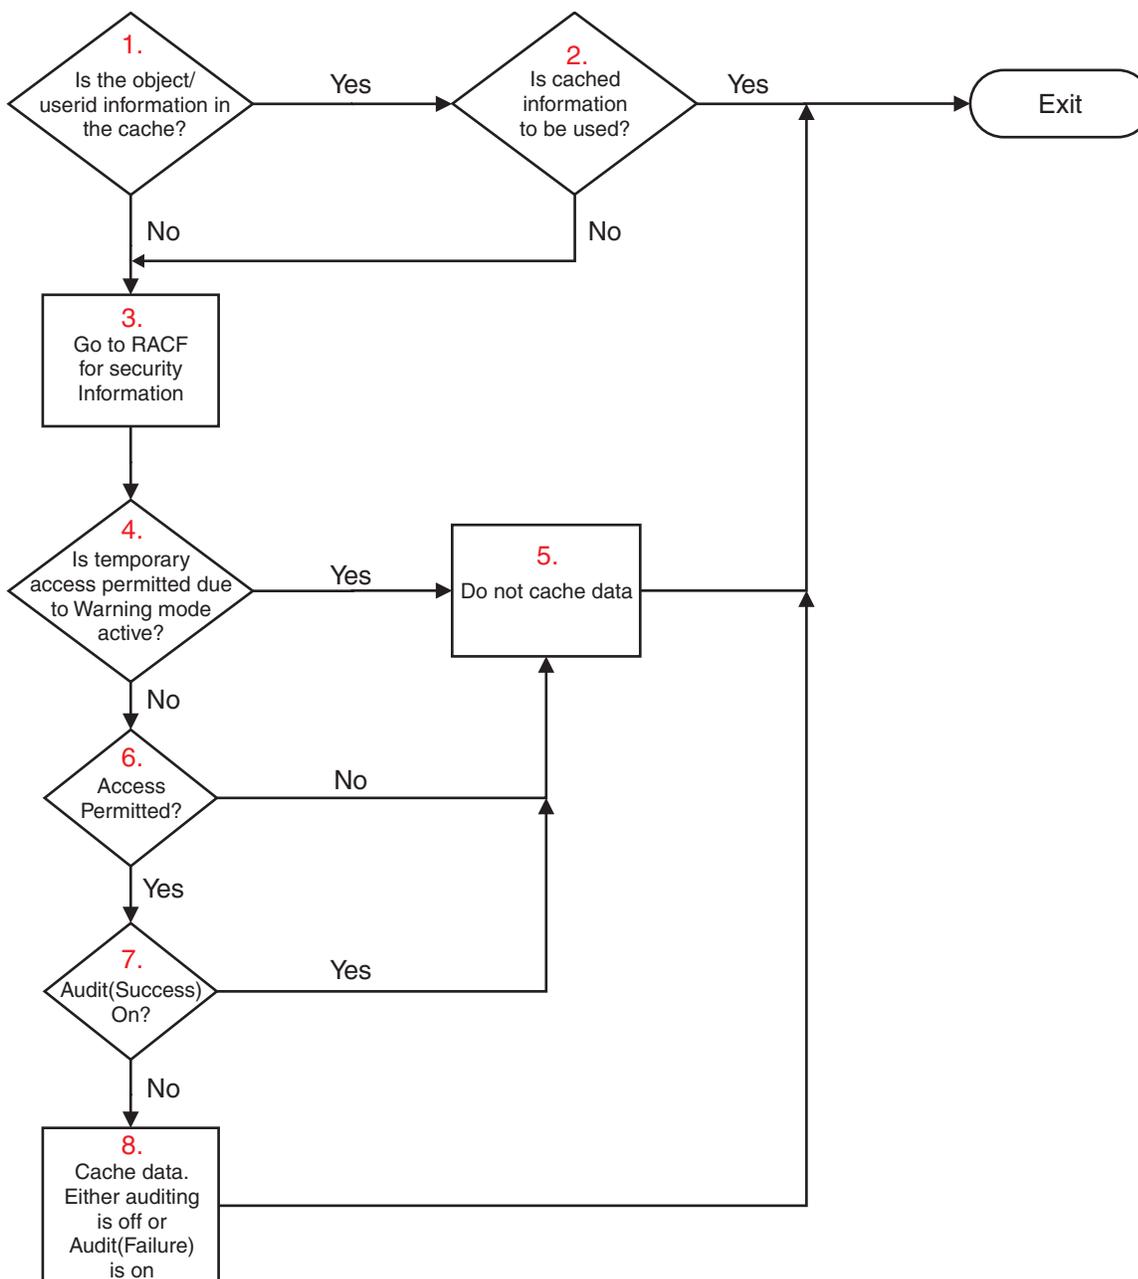


Figure 83. Logic flow for WebSphere MQ security caching

If you change your security settings by adding or deleting switch profiles in the MQADMIN or MXADMIN classes, use one of these commands to pick up these changes dynamically:

```
REFRESH SECURITY(*)
REFRESH SECURITY(MQADMIN)
REFRESH SECURITY(MXADMIN)
```

This means you can activate new security types, or deactivate them without having to restart the queue manager.

For performance reasons, these are the only classes affected by the REFRESH SECURITY command. You do *not* need to use REFRESH SECURITY if you change a profile in either the MQCONN or MQCMDS classes.

Note: A refresh of the MQADMIN or MXADMIN class is not required if you change a RESLEVEL security profile.

For performance reasons, use REFRESH SECURITY as infrequently as possible, ideally at off-peak times. You can minimize the number of security refreshes by connecting users to RACF groups that are already in the access list for WebSphere MQ profiles, rather than putting individual users in the access lists. In this way, you change the user rather than the resource profile. You can also RVERIFY SECURITY the appropriate user instead of refreshing security.

As an example of REFRESH SECURITY, suppose you define the new profiles to protect access to queues starting with INSURANCE.LIFE on queue manager PRMQ. You use these RACF commands:

```
RDEFINE MQQUEUE PRMQ.INSURANCE.LIFE.** UACC(NONE)
PERMIT PRMQ.INSURANCE.LIFE.** ID(LIFEGRP) ACCESS(UPDATE)
```

You must issue the following command to tell RACF to refresh the security information that it holds, for example:

```
SETROPTS RACLIST(MQQUEUE) REFRESH
```

Because these profiles are generic, you must tell RACF to refresh the generic profiles for MQQUEUE. For example:

```
SETROPTS GENERIC(MQQUEUE) REFRESH
```

Then you must use this command to tell queue manager PRMQ that the queue profiles have changed:

```
REFRESH SECURITY(MQQUEUE)
```

Refreshing SSL security

To refresh the cached view of the SSL Key Repository, issue the REFRESH SECURITY command with the option TYPE(SSL). This enables you to update some of your SSL settings without having to restart your channel initiator.

Related reference:



REFRESH SECURITY (*WebSphere MQ V7.1 Reference*)

Displaying security status:

To display the status of the security switches, and other security controls, issue the MQSC DISPLAY SECURITY command.

The following figure shows typical output of the DISPLAY SECURITY ALL command.

```

CSQH015I +CSQ1 Security timeout = 54 MINUTES
CSQH016I +CSQ1 Security interval = 12 MINUTES
CSQH030I +CSQ1 Security switches ...
CSQH034I +CSQ1 SUBSYSTEM: ON, 'SQ05.NO.SUBSYS.SECURITY' not found
CSQH032I +CSQ1 QMGR: ON, 'CSQ1.YES.QMGR.CHECKS' found
CSQH031I +CSQ1 QSG: OFF, 'SQ05.NO.QSG.CHECKS' found
CSQH031I +CSQ1 CONNECTION: OFF, 'CSQ1.NO.CONNECT.CHECKS' found
CSQH034I +CSQ1 COMMAND: ON, 'CSQ1.NO.COMMAND.CHECKS' not found
CSQH031I +CSQ1 CONTEXT: OFF, 'CSQ1.NO.CONTEXT.CHECKS' found
CSQH034I +CSQ1 ALTERNATE USER: ON, 'CSQ1.NO.ALTERNATE.USER.CHECKS' not found
CSQH034I +CSQ1 PROCESS: ON, 'CSQ1.NO.PROCESS.CHECKS' not found
CSQH034I +CSQ1 NAMELIST: ON, 'CSQ1.NO.NLIST.CHECKS' not found
CSQH034I +CSQ1 QUEUE: ON, 'CSQ1.NO.QUEUE.CHECKS' not found
CSQH034I +CSQ1 TOPIC: ON, 'CSQ1.NO.TOPIC.CHECKS' not found
CSQH031I +CSQ1 COMMAND RESOURCES: OFF, 'CSQ1.NO.CMD.RESC.CHECKS' found
CSQ9022I +CSQ1 CSQHPDTC ' DISPLAY SECURITY' NORMAL COMPLETION

```

Figure 84. Typical output from the DISPLAY SECURITY command

The example shows that the queue manager that replied to the command has subsystem, command, alternate user, process, namelist, and queue security active at queue manager level but not at queue-sharing group level. Connection, command resource, and context security are not active. It also shows that user ID timeouts are active, and that every 12 minutes the queue manager checks for user IDs that have not been used in this queue manager for 54 minutes and removes them.

Note: This command shows the current security status. It does not necessarily reflect the current status of the switch profiles defined to RACF, or the status of the RACF classes. For example, the switch profiles might have been changed since the last restart of this queue manager or REFRESH SECURITY command.

Related reference:



DISPLAY SECURITY (*WebSphere MQ V7.1 Reference*)

Security installation tasks for z/OS

After installing and customizing WebSphere MQ, authorize started task procedures to RACF, authorize access to various resources, and set up RACF definitions. Optionally, configure your system for SSL.

When WebSphere MQ is first installed and customized, you must perform these security-related tasks:

1. Set up WebSphere MQ data set and system security by:
 - Authorizing the queue manager started-task procedure xxxxMSTR and the distributed queuing started-task procedure xxxxCHIN to run under RACF.
 - Authorizing access to queue manager data sets.
 - Authorizing access to resources for those user IDs that will use the queue manager and utility programs.
 - Authorizing access for those queue managers that will use the coupling facility list structures.
 - Authorizing access for those queue managers that will use Db2.
2. Set up RACF definitions for WebSphere MQ security.
3. If you want to use the Secure Sockets Layer (SSL), prepare your system to use certificates and keys.

Setting up WebSphere MQ for z/OS data set security:

There are many types of WebSphere MQ user. Use RACF to control their access to system data sets.

The possible users of WebSphere MQ data sets include the following entities:

- The queue manager itself.
- The channel initiator
- WebSphere MQ administrators, who need to create WebSphere MQ data sets, run utility programs, and similar tasks.
- Application programmers who need to use the WebSphere MQ-supplied copybooks, include data sets, macros, and similar resources.
- Applications involving one or more of:
 - Batch jobs
 - TSO users
 - CICS regions
 - IMS regions
- Data sets CSQOUTX and CSQSNAP
- Dynamic queues SYSTEM.CSQXCMD.*

For all these potential users, protect the WebSphere MQ data sets with RACF.

You must also control access to all your 'CSQINP' data sets.

RACF authorization of started-task procedures:

Some IBM WebSphere MQ data sets are for the exclusive use of the queue manager. If you protect your IBM WebSphere MQ data sets using RACF, you must also authorize the queue manager started-task procedure xxxxMSTR, and the distributed queuing started-task procedure xxxxCHIN, using RACF. To do this, use the STARTED class. Alternatively, you can use the started procedures table (ICHRIN03), but then you must perform an IPL of your z/OS system before the changes take effect.

For more information, see the *z/OS Security Server RACF System Programmer's Guide*.

The RACF user ID identified must have the required access to the data sets in the started-task procedure. For example, if you associate a queue manager started task procedure called CSQ1MSTR with the RACF user ID QMGRCSQ1, the user ID QMGRCSQ1 must have access to the z/OS resources accessed by the CSQ1 queue manager.

Also, the content of the GROUP field in the user ID of the queue manager must be the same as the content of the GROUP field in the STARTED profile for that queue manager. If the content in each GROUP field does not match then the appropriate user ID is prevented from entering the system. This situation causes IBM WebSphere MQ to run with an undefined user ID and consequently close due to a security violation.

The RACF user IDs associated with the queue manager and channel initiator started task procedures must not have the TRUSTED attribute set.

Authorizing access to data sets:

The WebSphere MQ data sets should be protected so that no unauthorized user can run a queue manager instance, or gain access to any queue manager data. To do this, use normal z/OS RACF data set protection.

Table 69 summarizes the RACF access that the queue manager started task procedure must have to the different data sets.

Table 69. RACF access to data sets associated with a queue manager

RACF access	Data sets
READ	<ul style="list-style-type: none">• thlqual.SCSQAUTH and thlqual.SCSQANLx (where x is the language letter for your national language).• The data sets referred to by CSQINP1, CSQINP2 and CSQXLIB in the queue manager's started task procedure.
UPDATE	<ul style="list-style-type: none">• All page sets and log and BSDS data sets.
ALTER	<ul style="list-style-type: none">• All archive log data sets.

Table 70 summarizes the RACF access that the started task procedure for distributed queuing must have to the different data sets.

Table 70. RACF access to data sets associated with distributed queuing

RACF access	Data sets
READ	<ul style="list-style-type: none">• thlqual.SCSQAUTH, thlqual.SCSQANLx (where x is the language letter for your national language), and thlqual.SCSQMVR1.• LE library data sets.• The data sets referred to by CSQXLIB and CSQINPX in the distributed queuing started task procedure.
UPDATE	<ul style="list-style-type: none">• Data sets CSQOUTX and CSQSNAP• Dynamic queues SYSTEM.CSQXCMD.*

For more information, see the *z/OS Security Server RACF Security Administrator's Guide*.

Setting up WebSphere MQ for z/OS resource security:

There are many types of WebSphere MQ user. Use RACF to control their access to WebSphere MQ resources.

The possible users of WebSphere MQ resources, such as queues and channels include the following entities:

- The queue manager itself.
- The channel initiator
- WebSphere MQ administrators, who need to create WebSphere MQ data sets, run utility programs, and similar tasks
- Application programmers who need to use the WebSphere MQ-supplied copybooks, include data sets, macros, and similar resources.
- Applications involving one or more of:
 - Batch jobs
 - TSO users

- CICS regions
- IMS regions
- Data sets CSQOUTX and CSQSNAP
- Dynamic queues SYSTEM.CSQXCMD.*

For all these potential users, protect the WebSphere MQ resources with RACF. In particular, note that the channel initiator needs access to various resources, as described in “Security considerations for the channel initiator on z/OS” on page 590, and so the user ID under which it runs must be authorized to access these resources.

If you are using a queue-sharing group, the queue manager might issue various commands internally, so the user ID it uses must be authorized to issue such commands. The commands are:

- DEFINE, ALTER, and DELETE for every object that has QSGDISP(GROUP)
- START and STOP CHANNEL for every channel used with CHLDISP(SHARED)

Configuring your system to use the Secure Sockets Layer (SSL):

Use this topic as example of how to configure WebSphere MQ for z/OS with SSL using RACF commands.

If you want to use SSL for channel security, there are a number of tasks you need to perform on your system. (For details on using RACF commands for certificates and key repositories (key rings), see Working with SSL or TLS on z/OS.)

1. Create a key ring in RACF to hold all the keys and certificates for your system, using the RACF RACDCERT command. For example:

```
RACDCERT ID(CHINUSER) ADDRING(QM1RING)
```

The ID must be either the channel initiator address space user ID or the user ID you want to own the key ring if it is to be a shared key ring.

2. Create a digital certificate for each queue manager, using the RACF RACDCERT command. The label of the certificate must be of the form *ibmWebSphereMQqmgr-name* or *ibmWebSphereMQqsg-name*: in this example it is *ibmWebSphereMQQM1*.

For example:

```
RACDCERT ID(USERID) GENCERT
SUBJECTSDN(CN('username') O('IBM') OU('departmentname') C('England'))
WITHLABEL('ibmWebSphereMQQM1')
```

3. Connect the certificate in RACF to the key ring, using the RACF RACDCERT command. For example:

```
RACDCERT CONNECT(ID(USERID) LABEL('ibmWebSphereMQQM1') RING(QM1RING))
CONNECT ID(CHINUSER)
```

You also need to connect any relevant signer certificates (from a certificate authority) to the key ring. That is, all certificate authorities for the SSL certificate of this queue manager and all certificate authorities for all SSL certificates that this queue manager communicates with. For example:

```
RACDCERT ID(CHINUSER)
CONNECT(CERTAUTH LABEL('My CA') RING(QM1RING) USAGE(CERTAUTH))
```

4. On each of your queue managers, use the WebSphere MQ ALTER QMGR command to specify the key repository that the queue manager needs to point to. For example, if the key ring is owned by the channel initiator address space:

```
ALTER QMGR SSLKEYR(QM1RING)
```

or if you are using a shared key ring:

```
ALTER QMGR SSLKEYR(userid/QM1RING)
```

where *userid* is the user ID that owns the shared key ring.

5. Certificate Revocation Lists (CRLs) allow the certificate authorities to revoke certificates that can no longer be trusted. CRLs are stored in LDAP servers. To access this list on the LDAP server, you first need to create an AUTHINFO object of AUTHTYPE CRLLDAP, using the WebSphere MQ DEFINE AUTHINFO command. For example:

```
DEFINE AUTHINFO(LDAP1)
AUTHTYPE(CRLLDAP)
CONNAME(ldap.server(389))
LDAPUSER('')
LDAPPWD('')
```

In this example, the certificate revocation list is stored in a public area of the LDAP server, so the LDAPUSER and LDAPPWD fields are not necessary.

Next, put your AUTHINFO object into a namelist, using the WebSphere MQ DEFINE NAMELIST command. For example:

```
DEFINE NAMELIST(LDAPNL) NAMES(LDAP1)
```

Finally, associate the namelist with each queue manager, using the WebSphere MQ ALTER QMGR command. For example:

```
ALTER QMGR SSLCRLNL(LDAPNL)
```

6. Set up your queue manager to run SSL calls, using the WebSphere MQ ALTER QMGR command. This defines server subtasks that handle SSL calls only, which leaves the normal dispatchers to continue processing as normal without being affected by any SSL calls. You must have at least two of these subtasks. For example:

```
ALTER QMGR SSLTASKS(8)
```

This change only takes effect when the channel initiator is restarted.

7. Specify the cipher specification to be used for each channel, using the WebSphere MQ DEFINE CHANNEL or ALTER CHANNEL command. For example:

```
ALTER CHANNEL(LDAPCHL)
CHLTYPE(SDR)
SSLCIPH(RC4_MD5_US)
```

Both ends of the channel must specify the same cipher specification.

Managing channel authentication records in a queue-sharing group

Channel authentication records apply to the queue manager that they are created on, they are not shared throughout the queue-sharing group. Therefore if all the queue managers in the queue sharing group are required to have the same rules, some management needs to be carried out to keep all the rules the consistent.

1. Always add the CMDSCOPE(*) option to all SET CHLAUTH commands. This will send the command to all running queue managers in the queue-sharing group
2. Use the DISPLAY CHLAUTH command with the CMDSCOPE(*) option and then analyze the responses to see if the records are the same from all queue managers. When an inconsistency is found a SET CHLAUTH command can be issued containing the same rule with CMDSCOPE(*) or CMDSCOPE(*qmgr-name*).
3. Add a member to the queue manager's CSQINP2 concatenation (see "Initialization commands" on page 245 for details) that has the full set of rules. These will be read as part of the queue manager's initialization process. If the SET CHLAUTH command uses ACTION(ADD) the rule will only be added if it didn't exist. Using ACTION(REPLACE) will replace an existing rule if it already exists or add it if it does not. The same member could then be placed in the CSQINP2 concatenation of all queue managers in the queue-sharing group.
4. Use the CSQUTIL utility (see  Issuing commands to IBM WebSphere MQ (COMMAND) (*WebSphere MQ V7.1 Reference*) for details) to extract the rules from one queue manager using either the MAKEDEF or MAKEREP option. Then replay the output using CSQUTIL into the target queue manager.

Related concepts:

Channel authentication records

Auditing considerations on z/OS

The normal RACF auditing controls are available for conducting a security audit of a queue manager. WebSphere MQ does not gather any security statistics of its own. The only statistics are those that can be created by auditing.

RACF auditing can be based upon:

- User IDs
- Resource classes
- Profiles

For more details, see the *z/OS Security Server RACF Auditor's Guide*.

Note: Auditing degrades performance; the more auditing you implement, the more performance is degraded. This is also a consideration for the use of the RACF WARNING option.

Auditing RESLEVEL:

Use the RESAUDIT system parameter to control the production of RESLEVEL audit records. RACF GENERAL audit records are produced.

Produce RESLEVEL audit records by setting the RESAUDIT system parameter to YES. If the RESAUDIT parameter is set to NO, audit records are not produced. For more details about setting this parameter, see

 Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*).

If RESAUDIT is set to YES, no normal RACF audit records are taken when the RESLEVEL check is made to see what access an address space user ID has to the hlq.RESLEVEL profile. Instead, WebSphere MQ requests that RACF create a GENERAL audit record (event number 27). These checks are only carried out at connect time, so the performance cost is minimal.

You can report the WebSphere MQ general audit records using the RACF report writer (RACFRW). You could use the following RACFRW commands to report the RESLEVEL access:

```
RACFRW
SELECT PROCESS
EVENT GENERAL
LIST
END
```

A sample report from RACFRW, excluding the *Date*, *Time*, and *SYSID* fields, is shown in Figure 85.

```

          RACF REPORT - LISTING OF PROCESS RECORDS                                PAGE 4
          E
          V Q
          E U
*JOB/USER *STEP/  --TERMINAL--  N A
  NAME    GROUP   ID    LVL  T  L
WS21B    MQMGRP  IGJZM000  0   27 0  JOBID=(WS21B 05.111 09:44:57),USERDATA=()
  TRUSTED USER                                AUTH=(NONE),REASON=(NONE)
                                                SESSION=TSOLOGON,TERMINAL=IGJZM000,
                                                LOGSTR='CSQH RESLEVEL CHECK PERFORMED AGAINST PROFILE(QM66.RESLEVEL),
                                                CLASS(MQADMIN), ACCESS EQUATES TO (CONTROL)',RESULT=SUCCESS,MQADMIN
```

Figure 85. Sample output from RACFRW showing RESLEVEL general audit records

From checking the LOGSTR data in the output above, you can see that TSO user WS21B has CONTROL access to QM66.RESLEVEL. This means that all resource security checks are bypassed when user WS21B access QM66 resources.

For more information about using RACFRW, see the *z/OS Security Server RACF Auditor's Guide*.

Customizing security

If you want to change the way WebSphere MQ security operates, you must do this through the SAF exit (ICHRFR00), or exits in your external security manager.

To find out more about RACF exits, see the *z/OS Security Server RACROUTE Macro Reference* manual.

Note: Because WebSphere MQ optimizes calls to the ESM, RACROUTE requests might not be made on, for example, every open for a particular queue by a particular user.

Security violation messages on z/OS

A security violation is indicated by the return code MQRC_NOT_AUTHORIZED in an application program or by a message in the job log.

A return code of MQRC_NOT_AUTHORIZED can be returned to an application program for the following reasons:

- A user is not allowed to connect to the queue manager. In this case, you get an ICH408I message in the Batch/TSO, CICS, or IMS job log.
- A user sign-on to the queue manager has failed because, for example, the job user ID is not valid or appropriate, or the task user ID or alternate user ID is not valid. One or more of these user IDs might not be valid because they have been revoked or deleted. In this case, you get an ICHxxxx message and possibly an IRRxxxx message in the queue manager job log giving the reason for the sign-on failure. For example:

```
ICH408I USER(NOTDFND ) GROUP(          ) NAME(???          )
LOGON/JOB INITIATION - USER AT TERMINAL          NOT RACF-DEFINED
IRR012I VERIFICATION FAILED. USER PROFILE NOT FOUND
```

- An alternate user has been requested, but the job or task user ID does not have access to the alternate user ID. For this failure, you get a violation message in the job log of the relevant queue manager.
- A context option has been used or is implied by opening a transmission queue for output, but the job user ID or, where applicable, the task or alternate user ID does not have access to the context option. In this case, a violation message is put in the job log of the relevant queue manager.
- An unauthorized user has attempted to access a secured queue manager object, for example, a queue. In this case, an ICH408I message for the violation is put in the job log of the relevant queue manager. This violation might be due to the job or, when applicable, the task or alternate user ID.

Violation messages for command security and command resource security can also be found in the job log of the queue manager.

If the ICH408I violation message shows the queue manager jobname rather than a user ID, this is normally the result of a blank alternate user ID being specified. For example:

```
ICH408I JOB(MQS1MSTR) STEP(MQS1MSTR)
MQS1.PAYROLL.REQUEST CL(MQQUEUE)
INSUFFICIENT ACCESS AUTHORITY
ACCESS INTENT(UPDATE ) ACCESS ALLOWED(NONE )
```

You can find out who is allowed to use blank alternate user IDs by checking the access list of the MQADMIN profile hlq.ALTERNATE.USER.-BLANK-.

An ICH408I violation message can also be generated by:

- A command being sent to the system-command input queue without context. User-written programs that write to the system-command input queue should always use a context option. For more information, see “Profiles for context security” on page 549.
- When the job accessing the WebSphere MQ resource does not have a user ID associated with it, or when a WebSphere MQ adapter cannot extract the user ID from the adapter environment.

Violation messages might also be issued if you are using both queue-sharing group and queue manager level security. You might get messages indicating that no profile has been found at queue manager level, but still be granted access because of a queue-sharing group level profile.

```

ICH408I JOB(MQS1MSTR) STEP(MQS1MSTR)
      MQS1.PAYROLL.REQUEST CL(MQQUEUE)
      PROFILE NOT FOUND - REQUIRED FOR AUTHORITY CHECKING
      ACCESS INTENT(UPDATE ) ACCESS ALLOWED(NONE )

```

What to do if access is allowed or disallowed incorrectly

In addition to the steps detailed in the *z/OS Security Server RACF Security Administrator's Guide*, use this checklist if access to a resource appears to be incorrectly controlled.

- Are the switch profiles correctly set?
 - Is RACF active?
 - Are the WebSphere MQ RACF classes installed and active?
Use the RACF command, SETROPTS LIST, to check this.
 - Use the WebSphere MQ DISPLAY SECURITY command to display the current switch status from the queue manager.
 - Check the switch profiles in the MQADMIN class.
Use the RACF commands, SEARCH and RLIST, for this.
 - Recheck the RACF switch profiles by issuing the WebSphere MQ REFRESH SECURITY(MQADMIN) command.
- Has the RACF resource profile changed? For example, has universal access on the profile changed or has the access list of the profile changed?
 - Is the profile generic?
If it is, issue the RACF command, SETROPTS GENERIC(classname) REFRESH.
 - Have you refreshed the security on this queue manager?
If required, issue the RACF command SETROPTS RACLIST(classname) REFRESH.
If required, issue the WebSphere MQ REFRESH SECURITY(*) command.
- Has the RACF definition of the user changed? For example, has the user been connected to a new group or has the user access authority been revoked?
 - Have you reverified the user by issuing the WebSphere MQ RVERIFY SECURITY(userid) command?
- Are security checks being bypassed due to RESLEVEL?
 - Check the connecting user ID's access to the RESLEVEL profile. Use the RACF audit records to determine what the RESLEVEL is set to.
 - For channels, remember that the access level that the channel initiator's userid has to RESLEVEL is inherited by all channels, so an access level, such as ALTER, that causes all checks to be bypassed causes security checks to be bypassed for all channels.
 - If you are running from CICS, check the transaction's RESSEC setting.
 - If RESLEVEL has been changed while a user is connected, they must disconnect and reconnect before the new RESLEVEL setting takes effect.
- Are you using queue-sharing groups?
 - If you are using both queue-sharing group and queue manager level security, check that you have defined all the correct profiles. If queue manager profile is not defined, a message is sent to the log stating that the profile was not found.
 - Have you used a combination of switch settings that is not valid so that full security checking has been set on?
 - Do you need to define security switches to override some of the queue-sharing group settings for your queue manager?
 - Is a queue manager level profile taking precedence over a queue-sharing group level profile?

Security considerations for the channel initiator on z/OS

If you are using resource security in a distributed queuing environment, the Channel initiator address space needs appropriate access to various WebSphere MQ resources.

If you are using resource security, consider the following points if you are using distributed queuing:

System queues

The channel initiator address space needs RACF UPDATE access to the system queues listed at “System queue security” on page 539, and to all the user destination queues and the dead-letter queue (but see “Dead-letter queue security” on page 538).

Transmission queues

The channel initiator address space needs ALTER access to all the user transmission queues.

Context security

The channel user ID (and the MCA user ID if one has been specified) need RACF CONTROL access to the hlq.CONTEXT.queuname profiles in the MQADMIN class. Depending on the RESLEVEL profile, the channel user ID might also need CONTROL access to these profiles.

All channels need CONTROL access to the MQADMIN hlq.CONTEXT. dead-letter-queue profile. All channels (whether initiating or responding) can generate reports, and consequently they need CONTROL access to the hlq.CONTEXT.reply-q profile.

SENDER, CLUSSDR, and SERVER channels need CONTROL access to the hlq.CONTEXT.xmit-queue-name profiles since messages can be put onto the transmission queue to wake up the channel to end gracefully.

Note: If the channel user ID, or a RACF group to which the channel user ID is connected, has CONTROL or ALTER access to the hlq.RESLEVEL, then there are no resource checks for the channel initiator or any of its channels.

See “Profiles for context security” on page 549 “RESLEVEL and the channel initiator connection” on page 566 and “User IDs for security checking” on page 568 for more information.

CSQINPX

If you are using the CSQINPX input data set, the channel initiator also needs READ access to CSQINPX, and UPDATE access to data set CSQOUTX and dynamic queues SYSTEM.CSQXCMD.*.

Connection security

The channel initiator address space connection requests use a connection type of CHIN, for which appropriate access security must be set, see “Connection security profiles for the channel initiator” on page 532.

Data sets

The channel initiator address space needs appropriate access to queue manager data sets, see “Authorizing access to data sets” on page 583.

Commands

The distributed queuing commands (for example, DEFINE CHANNEL, START CHINIT, START LISTENER, and other channel commands) must have appropriate command security set, see Table 54 on page 552.

If you are using a queue-sharing group, the channel initiator might issue various commands internally, so the user ID it uses must be authorized to issue such commands. These commands are START and STOP CHANNEL for every channel used with CHLDISP(SHARED).

If the PSMODE of the queue manager is not DISABLED, the channel initiator must have READ access to the DISPLAY PUBSUB command.

Channel security

Channels, particularly receivers and server-connections, need appropriate security to be set up; see “User IDs for security checking” on page 568 for more information.

You can also use the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocols to provide security on channels. See “WebSphere MQ support for SSL and TLS” on page 386 for more information about using TLS and SSL with IBM WebSphere MQ.

See also “Access control for clients” on page 432 for information about server-connection security.

User IDs

The user IDs described in “User IDs used by the channel initiator” on page 571 and “User IDs used by the intra-group queuing agent” on page 575 need the following access:

- RACF UPDATE access to the appropriate destination queues and the dead-letter queue
- RACF CONTROL access to the hlq.CONTEXT.queueName profile if context checking is performed at the receiver
- Appropriate access to the hlq.ALTERNATE.USER.userid profiles they might need to use.
- For clients, the appropriate RACF access to the resources to be used.

APPC security

Set appropriate APPC security if you are using the LU 6.2 transmission protocol. (Use the APPCLU RACF class for example.) For information about setting up security for APPC, see the following manuals:

- *z/OS V1R2.0 MVS Planning: APPC Management*
- *Multiplatform APPC Configuration Guide*, an IBM Redbooks publication

Outbound transmissions use the “SECURITY(SAME)” APPC option. As a result, the user ID of the channel initiator address space and its default profile (RACF GROUP) are flowed across the network to the receiver with an indicator that the user ID has already been verified (ALREADYV).

If the receiving side is also z/OS, the user ID and profile are verified by APPC and the user ID is presented to the receiver channel and used as the channel user ID.

In an environment where the queue manager is using APPC to communicate with another queue manager on the same or another z/OS system, you need to ensure that either:

- The VTAM[®] definition for the communicating LU specifies SETACPT(ALREADYV)
- There is a RACF APPCLU profile for the connection between LUs that specifies CONVSEC(ALREADYV)

Changing security settings

If the RACF access level that either the channel user ID or MCA user ID has to a destination queue is changed, this change takes effect only for new object handles (that is, new **MQOPENS**) for the destination queue. The times when MCAs open and close queues is variable; if a channel is already running when such an access change is made, the MCA can continue to put messages on the destination queue using the existing security access of the user IDs rather than the updated security access. Stopping and restarting the channels to enforce the updated access level avoids this scenario.

Automatic restart

If you are using the z/OS Automatic Restart Manager (ARM) to restart the channel initiator, the user ID associated with the XCFAS address space must be authorized to issue the WebSphere MQ START CHINIT command.

Security in queue manager clusters on z/OS

Security considerations for clusters are the same for queue managers and channels that are not clustered. The channel initiator needs access to some additional system queues, and some additional commands need appropriate security set.

You can use the MCA user ID, channel authentication records, SSL or TLS, and security exits to authenticate cluster channels (as with conventional channels). The channel authentication records or security exit relating to the cluster-receiver channel must check that the remote queue manager is permitted access to the server queue manager's cluster queues. You can start to use IBM WebSphere MQ cluster support without changing your existing queue access security. You must, however, allow other queue managers in the cluster to write to the `SYSTEM.CLUSTER.COMMAND.QUEUE` if they are to join the cluster.

IBM WebSphere MQ cluster support does not provide a mechanism to limit a member of a cluster to the client role only. As a result, you must be sure that you trust any queue managers that you allow into the cluster. If any queue manager in the cluster creates a queue with a particular name, it can receive messages for that queue, regardless of whether the application putting messages to that queue intended this or not.

To restrict the membership of a cluster, take the same action that you would take to prevent queue managers connecting to receiver channels. You restrict the membership of a cluster by using channel authentication records or by writing a security exit program on the receiver channel. You can also write an exit program to prevent unauthorized queue managers from writing to the `SYSTEM.CLUSTER.COMMAND.QUEUE`.

Note: It is not advisable to permit applications to open the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` directly. It is also not advisable to permit an application to open any other transmission queue directly.

If you are using resource security, consider the following points in addition to the considerations contained in "Security considerations for the channel initiator on z/OS" on page 590:

System queues

The channel initiator needs RACF ALTER access to the following system queues:

- `SYSTEM.CLUSTER.COMMAND.QUEUE`
- `SYSTEM.CLUSTER.TRANSMIT.QUEUE`.

and UPDATE access to `SYSTEM.CLUSTER.REPOSITORY.QUEUE`

It also needs READ access to any namelists used for clustering.

Commands

Set appropriate command security (as described in Table 54 on page 552) for the cluster support commands (`REFRESH` and `RESET CLUSTER`, `SUSPEND`, and `RESUME QMGR`).

Security considerations for using WebSphere MQ with CICS

The CICS adapter provides information to WebSphere MQ for use in security.

The CICS adapter provides the following information to WebSphere MQ specifically for use in WebSphere MQ security:

- Whether CICS resource-level security is active for this transaction—as specified on the `RESSEC` or `RSLC` operand of the `RDO TRANSACTION` definition.
- User IDs.

For terminal tasks where a user has not signed on, the user ID is the CICS user ID associated with the terminal and is either:

- The default CICS user ID as specified on the `CICS` parameter `DFLTUSER SIT`
- A preset security user ID specified on the terminal definition

For non-terminal tasks, the CICS adapter tries to get a user ID with an EXEC CICS ASSIGN command. If this is unsuccessful, the adapter tries to get the user ID using EXEC CICS INQUIRE TASK. If security is active in CICS, and the non-terminal attached transaction is defined with CMDSEC(YES), the CICS adapter passes a user ID of blanks to WebSphere MQ.

For details of security considerations, see:

-  Security for the CICS-WebSphere MQ adapter.
-  Security for the CICS-WebSphere MQ bridge.

Controlling the security of CICS transactions supplied by WebSphere MQ:

If you want a user to administer the CICS adapter, grant the user authorization to certain transactions.

The CKTI and CKAM transactions are designed to be run without a terminal; no user should have access to these transactions. These transactions are examples of what the *CICS RACF Security Guide* calls “category 2 transactions”. For information about how to set up these transactions in CICS and RACF, see the information about category 2 transactions in the *CICS RACF Security Guide*.

If you want a user to administer the CICS adapter, you must grant the user authorization to these transactions:

CKQC	Controls the CICS adapter functions
CKBM	Controls the CICS adapter functions
CKRT	Controls the CICS adapter functions
CKCN	Connect
CKSD	Disconnect
CKRS	Statistics
CKDP	Full screen display
CKDL	Line mode display
CKSQ	CKTI START/STOP

If required, you can restrict access to specific functions of the adapter. For example, if you want to allow users to display the status of the adapter through the full screen interface, but nothing else, give them access to CKQC, CKBM, CKRT, and CKDP only.

Define these transactions to CICS with RESSEC(NO) and CMDSEC(NO). For more details, see the *CICS RACF Security Guide*.

The CICS adapter user ID:

The user ID associated with the CICS adapter is that of the WebSphere MQ-supplied task initiator transaction, CKTI. This has a number of implications.

User ID checking for WebSphere MQ resources during PLTPI and PLTSD

If a WebSphere MQ resource is accessed during the CICS PLTPI phase, the user ID passed to WebSphere MQ is blanks. If a WebSphere MQ resource is accessed during the CICS PLTSD phase, the user ID passed to WebSphere MQ is the user ID associated with the shutdown transaction.

If CKTI is started during the CICS PLTPI phase, the user ID of the CKTI task is the CICS sysidnt. This means that a user ID with the same name as the CICS sysidnt must be defined and given access to the required WebSphere MQ resources, for example, initiation queues.

Terminal user IDs

If CKTI is started from a terminal from the CKQC transaction or a user-written program that links to CSQCSSQ, the user ID that CKTI uses is the same as the user ID of the terminal that started CKTI.

Automating starting of CKTI

To automate the starting of CKTIs under a specific user ID, you can use an automation product, for example, Tivoli NetView[®] for z/OS. You can use this to sign on a CICS console and issue the STARTCKTI command.

You can also use preset security sequential terminals, which have been defined to emulate a CRLP terminal, with the sequential terminal input containing the CKQC STARTCKTI command.

However, when the CICS adapter alert monitor reconnects CICS to WebSphere MQ, after, for example, a WebSphere MQ restart, only the CKTI specified at the initial WebSphere MQ connection is restarted. You must automate starting any extra CKTIs yourself.

Propagating the CKTI user ID to other CICS transactions

If CKTI starts other CICS transactions, for example, message channel agents (MCAs) or user-written CICS applications, the user ID of CKTI is propagated to these applications. For example, if CKTI is running under user ID *USERNAME* and a trigger event occurs that requires the sender MCA transaction, *TRNA*, to be started, the *TRNA* transaction also runs under user ID *USERNAME*. Therefore user ID *USERNAME* must have access to the required transmission queue.

Security considerations for the CICS bridge:

When you run the CICS bridge, you can specify the level of authentication you want to take place.

If requested, the bridge checks the user ID and password extracted from the IBM WebSphere MQ request message before running the CICS program named in the request message. The queue manager uses the external security manager (ESM) (for example RACF) to do authentication. Therefore, user IDs in the request message must be defined to the ESM.

Note:

1. If you have not specified a user ID in the message descriptor (MQMD) or password in the CICS bridge header (MQCIH) of a message, the bridge task runs with the LOCAL level of authentication, even if you started the bridge monitor with a different authentication option.
2. The options that include password (or PassTicket) validation require an MQCIH to be provided. See  MQCIH – CICS bridge header (*WebSphere MQ V7.1 Reference*) for more information about the MQCIH header.
3. PassTicket validation is performed using IBM WebSphere MQ services, not EXEC CICS VERIFY, as the CICS service does not accept an APPLID.

The level of authentication you can use is described using the following keywords:

LOCAL

This is the default value. CICS programs run by the bridge task are started with the CICS DFLTUSER user ID, therefore run with the authority associated with this user ID. There is no checking of user IDs or passwords. If a CICS program is run that tries to access protected resources, it is likely to fail.

IDENTIFY

When you start the monitor task with the IDENTIFY authentication option, the bridge task is

started with the user ID specified in the message (MQMD). CICS programs run by the bridge run with the user ID from the MQMD. There is no password checking, the user ID is treated as trusted.

VERIFY_UOW

If MQMD.PutAppType is set to MQAT_NO_CONTEXT, the CICS DFLTUSER user ID is used, as for the LOCAL authentication option. Otherwise, the bridge monitor checks the user ID (in the MQMD) and password (in the CIH) before starting the bridge task, and CICS programs run by the bridge run with the user ID extracted from the MQMD. If the user ID or password is invalid, the request fails with return code MQCRC_SECURITY_ERROR. Subsequent messages processed by this transaction are not checked.

Bridge tasks that cannot be started because of a security failure are tried again if ROUTEMEM=N (the default) is specified. However, if ROUTEMEM=Y is specified the bridge messages are put to the dead-letter queue.

VERIFY_ALL

This value is the same as VERIFY_UOW except that the bridge task checks the user ID and password in every message. This option is not applicable for 3270 transactions when using CICS earlier than CICS Transaction Server Version 2 Release 2.

A PassTicket can be used in place of a password to avoid the need to flow passwords in messages (see Security Server RACF System Programmer's Guide). When generating a PassTicket an APPLID must be specified. If you are using a single bridge monitor, the APPLID is the CICS APPLID unless a different value was specified when the bridge was started. If you are using multiple bridge monitors for a queue, you must specify the APPLID to be used, using the PASSTKTA=*applid* parameter at bridge startup.

If you have not specified a user ID in a message, or you have not provided a password, the CICS program started by the CICS bridge runs with the user ID set to the user ID used to start the bridge monitor, regardless of the option requested. If you want more than one level of authentication checking performed, run a monitor task for each level you need.

When a CICS DPL request is read by the bridge monitor it starts the transaction specified in the CICS bridge header (MQCIH) or, if no transaction is specified, transaction CKBP. The user IDs under which the bridge monitor runs must have authority to start the various transactions that might be requested. The default transaction ID for the CICS bridge monitor is CKBR but you can change this or define additional transaction IDs if you want more granular access to queues and transactions. You can use CICS surrogate security to restrict which user ID and transaction combinations a bridge monitor transaction and user ID can start.

Table 71 and Table 72 on page 596 summarize the level of authority of the bridge monitor and the bridge tasks, and the use of the MQMD user ID.

Table 71. CICS bridge monitor security

Monitor started by	At a signed on terminal	Monitor authority
From a terminal or EXEC CICS LINK within a program	Yes	Signed on user ID
From a terminal or EXEC CICS LINK within a program	No	CICS default user ID
EXEC CICS START with user ID	–	User ID from START
EXEC CICS START without user ID	–	CICS default user ID
The IBM WebSphere MQ trigger monitor CKTI	–	CICS default user ID

Table 72. CICS bridge task security

AUTH	Bridge task authority
LOCAL	CICS default user ID
IDENTIFY	MQMD UserIdentifier
VERIFY_UOW	MQMD UserIdentifier
VERIFY_ALL	MQMD UserIdentifier

The options IDENTIFY, VERIFY_UOW, and VERIFY_ALL need the user ID of the bridge monitor defined to RACF as a surrogate of all the user IDs used in request messages. This requirement is in addition to the user ID in the message being defined to RACF. (A surrogate user is one who has the authority to start work on behalf of another user, without knowing the other user's password.)

For more information about surrogate user security, see the *CICS RACF Security Guide*.

Note: When IDENTIFY security is being used, you might see abend AICO for CKBP if you try to run with a user ID that has been revoked. The error reply has return code MQCRC_BRIDGE_ERROR with reason MQFB_CICS_BRIDGE_FAILURE.

Authority for the CICS bridge:

Components of the bridge need authority to either put to or get from the various WebSphere MQ queues.

In summary, the required authority is as follows:

- The monitor and all bridge tasks need authority to get messages from the bridge request queue.
- A bridge task needs authority to put messages to its own reply-to queue.
- To ensure that any error replies are received, the monitor needs authority to put messages to all reply-to queues.
- Bridge tasks need authority to put messages to the dead-letter queue.
- The monitor needs authority to put messages to the dead-letter queue, unless you want the bridge to stop if an error occurs.
- The monitor and all bridge tasks need authority to put messages to the backout requeue queue, if one is defined

See Table 71 on page 595 to determine the correlation between user IDs and authority.

Security considerations for using WebSphere MQ with IMS

Use this topic to plan your security requirements when you use WebSphere MQ with IMS.

Using the OPERCMDS class

If you are using RACF to protect resources in the OPERCMDS class, ensure that the userid associated with your WebSphere MQ queue manager address space has authority to issue the MODIFY command to any IMS system to which it can connect.

Security considerations for the IMS bridge

There are four aspects that you must consider when deciding your security requirements for the IMS bridge, these are:

- What security authorization is needed to connect WebSphere MQ to IMS
- How much security checking is performed on applications using the bridge to access IMS
- Which IMS resources these applications are allowed to use

- What authority is to be used for messages that are put and got by the bridge

When you define your security requirements for the IMS bridge you must consider the following:

- Messages passing across the bridge might have originated from applications on platforms that do not offer strong security features
- Messages passing across the bridge might have originated from applications that are not controlled by the same enterprise or organization

Security considerations for connecting to IMS:

Grant the user ID of the WebSphere MQ queue manager address space access to the OTMA group.

The IMS bridge is an OTMA client. The connection to IMS operates under the user ID of the WebSphere MQ queue manager address space. This is normally defined as a member of the started task group. This user ID must be granted access to the OTMA group (unless the /SECURE OTMA setting is NONE).

To do this, define the following profile in the FACILITY class:

```
IMSXCF.xcfgname.mqxcfmname
```

Where xcfgname is the XCF group name and mqxcfmname is the XCF member name of WebSphere MQ.

You must give your WebSphere MQ queue manager user ID read access to this profile.

Note:

1. If you change the authorities in the FACILITY class, you must issue the RACF command SETROPTS RACLST(FACILITY) REFRESH to activate the changes.
2. If profile hlq.NO.SUBSYS.SECURITY exists in the MQADMIN class, no user ID is passed to IMS and the connection fails unless the /SECURE OTMA setting is NONE.

Application access control for the IMS bridge:

Define a RACF profile in the FACILITY class for each IMS system. Grant an appropriate level of access to the WebSphere MQ queue manager user ID.

For each IMS system that the IMS bridge connects to, you can define the following RACF profile in the FACILITY class to determine how much security checking is performed for each message passed to the IMS system.

```
IMSXCF.xcfgname.imsxcfmname
```

Where xcfgname is the XCF group name and imsxcfmname is the XCF member name for IMS. (You need to define a separate profile for each IMS system.)

The access level you allow for the WebSphere MQ queue manager user ID in this profile is returned to WebSphere MQ when the IMS bridge connects to IMS, and indicates the level of security that is required on subsequent transactions. For subsequent transactions, WebSphere MQ requests the appropriate services from RACF and, where the user ID is authorized, passes the message to IMS.

OTMA does not support the IMS /SIGN command; however, WebSphere MQ allows you to set the access checking for each message to enable implementation of the necessary level of control.

The following access level information can be returned:

NONE or NO PROFILE FOUND

These values indicate that maximum security is required, that is, authentication is required for every transaction. A check is made to verify that the user ID specified in the *UserIdentifier* field of the MQMD structure, and the password or PassTicket in the *Authenticator* field of the MQIIIH structure are known to RACF, and are a valid combination. A UTOKEN is created with a password or PassTicket, and passed to IMS; the UTOKEN is not cached.

Note: If profile hlq.NO.SUBSYS.SECURITY exists in the MQADMIN class, this level of security overrides whatever is defined in the profile.

READ This value indicates that the same authentication is to be performed as for NONE under the following circumstances:

- The first time that a specific user ID is encountered
- When the user ID has been encountered before but the cached UTOKEN was not created with a password or PassTicket

WebSphere MQ requests a UTOKEN if required, and passes it to IMS.

Note: If a request to reverify security has been acted on, all cached information is lost and a UTOKEN is requested the first time each user ID is later encountered.

UPDATE

A check is made that the user ID in the *UserIdentifier* field of the MQMD structure is known to RACF.

A UTOKEN is built and passed to IMS; the UTOKEN is cached.

CONTROL/ALTER

These values indicate that no security UTOKENs need to be provided for any user IDs for this IMS system. (You would probably only use this option for development and test systems.)

Attention: Note that the user ID contained in the *UserIdentifier* field of the MQMD structure is still passed for **CONTROL/ALTER**.

Note:

1. This access is defined when WebSphere MQ connects to IMS, and lasts for the duration of the connection. To change the security level, the access to the security profile must be changed and then the bridge stopped and restarted (for example, by stopping and restarting OTMA).
2. If you change the authorities in the FACILITY class, you must issue the RACF command SETROPTS RACLIST(FACILITY) REFRESH to activate the changes.
3. You can use a password or a PassTicket, but you must remember that the IMS bridge does not encrypt data. For information about using PassTickets, see "Using RACF PassTickets in the IMS header" on page 600.
4. Some of these results might be affected by security settings in IMS, using the /SECURE OTMA command.
5. Cached UTOKEN information is held for the duration defined by the INTERVAL and TIMEOUT parameters of the WebSphere MQ ALTER SECURITY command.

Security checking on IMS:

Messages that pass across the bridge contain security information. The security checks made depend on the setting of the IMS command /SECURE OTMA.

Each WebSphere MQ message that passes across the bridge contains the following security information:

- A user ID contained in the *UserIdentifier* field of the MQMD structure
- The security scope contained in the *SecurityScope* field of the MQIHL structure (if the MQIHL structure is present)
- A UTOKEN (unless the WebSphere MQ sub system has CONTROL or ALTER access to the relevant IMSXCF.xcfcname.imsxcfmname profile)

The security checks made depend on the setting of the IMS command /SECURE OTMA, as follows:

/SECURE OTMA NONE

No security checks are made for the transaction.

/SECURE OTMA CHECK

The *UserIdentifier* field of the MQMD structure is passed to IMS for transaction or command authority checking.

An ACEE (Accessor Environment Element) is built in the IMS control region.

/SECURE OTMA FULL

The *UserIdentifier* field of the MQMD structure is passed to IMS for transaction or command authority checking.

An ACEE is built in the IMS dependent region as well as the IMS control region.

/SECURE OTMA PROFILE

The *UserIdentifier* field of the MQMD structure is passed to IMS for transaction or command authority checking

The *SecurityScope* field in the MQIHL structure is used to determine whether to build an ACEE in the IMS dependent region as well as the control region.

Note:

1. If you change the authorities in the TIMS or CIMS class, or the associated group classes GIMS or DIMS, you must issue the following IMS commands to activate the changes:
 - /MODIFY PREPARE RACF
 - /MODIFY COMMIT
2. If you do not use /SECURE OTMA PROFILE, any value specified in the *SecurityScope* field of the MQIHL structure is ignored.

Security checking done by the IMS bridge:

Different authorities are used depending on the action being performed.

When the bridge puts or gets a message, the following authorities are used:

Getting a message from the bridge queue

No security checks are performed.

Putting an exception, or COA report message

Uses the authority of the user ID in the *UserIdentifier* field of the MQMD structure.

Putting a reply message

Uses the authority of the user ID in the *UserIdentifier* field of the MQMD structure of the original message

Putting a message to the dead-letter queue

No security checks are performed.

Note:

1. If you change the WebSphere MQ class profiles, you must issue the WebSphere MQ REFRESH SECURITY(*) command to activate the changes.
2. If you change the authority of a user, you must issue the MQSC RVERIFY SECURITY command to activate the change.

Using RACF PassTickets in the IMS header:

You can use a PassTicket in place of a password in the IMS header.

If you want to use a PassTicket instead of a password in the IMS header (MQIIH), specify the application name against which the PassTicket is validated in the PASSTKTA attribute of the STGCLASS definition of the IMS Bridge queue to which the message is to be routed.

If the PASSTKTA value is left blank, you must arrange to have a PassTicket generated. The application name in this case must be of the form MVSxxxx, where xxxx is the SMFID of the z/OS system on which the target queue manager runs.

A PassTicket is built from a user ID, the target application name, and a secret key. It is an 8-byte value containing uppercase alphabetic and numeric characters. It can be used only once, and is valid for a 20 minute period. If a PassTicket is generated by a local RACF system, RACF only checks that the profile exists and not that the user has authority against the profile. If the PassTicket was generated on a remote system, RACF validates the access of the user ID to the profile. For full information about PassTickets, see the *z/OS SecureWay Security Server RACF Security Administrator's Guide*.

PassTickets in IMS headers are given to RACF by WebSphere MQ, not IMS.

Security scenario: two queue managers on z/OS

In this scenario, an application uses the **MQPUT1** call to put messages to queues on queue manager QM1. Some of the messages are then forwarded to queues on QM2, using TCP and LU 6.2 channels. The TCP channels can either use SSL or not. The application could be a batch application or a CICS application, and the messages are put using the MQPMO_SET_ALL_CONTEXT option.

This is illustrated in Figure 86 on page 601.

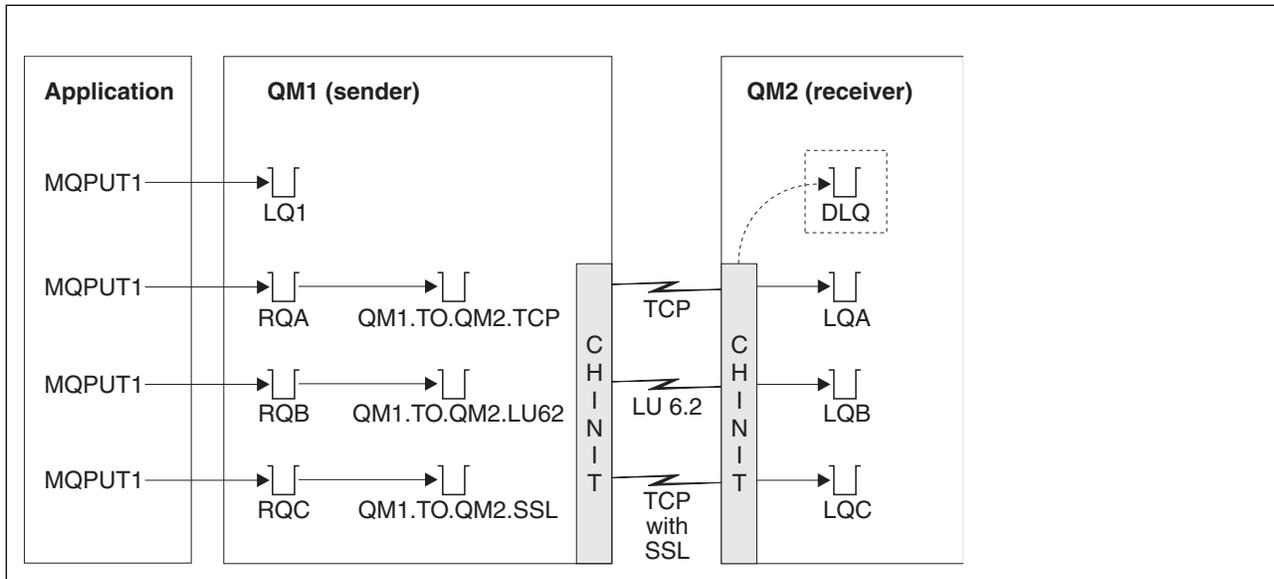


Figure 86. Example security scenario

The following assumptions are made about the queue managers:

- All the required WebSphere MQ definitions have been predefined or have been made through the CSQINP2 data set processed at queue manager startup.

If they have not, you need the appropriate access authority to the commands needed to define these objects.

- All the RACF profiles required have been defined and appropriate access authorities have been granted, before the queue manager and channel initiators started.

If they have not, you need the appropriate authority to issue the RACF commands required to define all the profiles needed and grant the appropriate access authorities to those profiles. You also need the appropriate authority to issue the MQSC security commands to start using the new security profiles.

- All digital certificates required have been created and connected to key rings. The digital certificate sent by QM1 as part of the SSL handshake is recognized by RACF on QM2's system, either because it is also installed in that RACF, or because a matching Certificate Name File (CNF) filter exists.

Security switch settings for two-queue-manager scenario:

Switch settings and RACF profiles.

The following security switches are set for both queue managers:

- Subsystem security on
- Queue security on
- Alternate user security on
- Context security on
- Process security off
- Namelist security off
- Connection security on
- Command security on
- Command resource security on

The following profiles are defined in the MQADMIN class to turn off process and namelist security:

QM1.NO.PROCESS.CHECKS
QM1.NO.NLIST.CHECKS
QM2.NO.PROCESS.CHECKS
QM2.NO.NLIST.CHECKS

Queue manager QM1 in two-queue-manager scenario:

Queues and channels for QM1.

The following queues are defined on queue manager QM1:

LQ1 A local queue.

RQA A remote queue definition, with the following attributes:

- RNAME(LQA)
- RQMNAME(QM2)
- XMITQ(QM1.TO.QM2.TCP)

RQB A remote queue definition, with the following attributes:

- RNAME(LQB)
- RQMNAME(QM2)
- XMITQ(QM1.TO.QM2.LU62)

RQC A remote queue definition, with the following attributes:

- RNAME(LQC)
- RQMNAME(QM2)
- XMITQ(QM1.TO.QM2.SSL)

QM1.TO.QM2.TCP

A transmission queue.

QM1.TO.QM2.LU62

A transmission queue.

QM1.TO.QM2.SSL

A transmission queue.

The following channels are defined on QM1:

QM1.TO.QM2.TCP

A sender channel definition, with the following attributes:

- CHLTYPE(SDR)
- TRPTYPE(TCP)
- XMITQ(QM1.TO.QM2.TCP)
- CONNAME(QM2TCP)

QM1.TO.QM2.LU62

A sender channel definition, with the following attributes:

- CHLTYPE(SDR)
- TRPTYPE(LU62)
- XMITQ(QM1.TO.QM2.LU62)
- CONNAME(QM2LU62)

(See "Security considerations for the channel initiator on z/OS" on page 590 for information about setting up APPC security.)

QM1.TO.QM2.SSL

A sender channel definition, with the following attributes:

- CHLTYPE(SDR)
- TRPTYPE(TCP)
- XMITQ(QM1.TO.QM2.SSL)
- CONNAME(QM2TCP)
- SSLCIPH(RC4_MD5_EXPORT)

Queue manager QM2 in two-queue-manager scenario:

Queues and channels for QM2.

The following queues have been defined on queue manager QM2:

LQA A local queue.

LQB A local queue.

LQC A local queue.

DLQ A local queue that is used as the dead-letter queue.

The following channels have been defined on QM2:

QM1.TO.QM2.TCP

A receiver channel definition, with the following attributes:

- CHLTYPE(RCVR)
- TRPTYPE(TCP)
- PUTAUT(CTX)
- MCAUSER(MCATCP)

QM1.TO.QM2.LU62

A receiver channel definition, with the following attributes:

- CHLTYPE(RCVR)
- TRPTYPE(LU62)
- PUTAUT(CTX)
- MCAUSER(MCALU62)

(See “Security considerations for the channel initiator on z/OS” on page 590 for information about setting up APPC security.)

QM1.TO.QM2.SSL

A receiver channel definition, with the following attributes:

- CHLTYPE(RCVR)
- TRPTYPE(TCP)
- PUTAUT(CTX)
- MCAUSER(MCASSL)
- SSLCIPH(RC4_MD5_EXPORT)

User IDs used in two-queue-manager scenario:

Explanation of the user IDs in the scenario.

The following user IDs are used:

BATCHID

Batch application (Job or TSO ID)

MSGUSR

UserIdentifier in MQMD (context user ID)

MOVER1

QM1 channel initiator address space user ID

MOVER2

QM2 channel initiator address space user ID

MCATCP

MCAUSER specified on the TCP/IP without SSL receiver channel definition

MCALU62

MCAUSER specified on the LU 6.2 receiver channel definition

MCASSL

MCAUSER specified on the TCP/IP with SSL receiver channel definition

CICSAD1

CICS address space ID

CICSTX1

CICS task user ID

CERTID

The user ID associated by RACF with the flowed certificate.

Security profiles and accesses required for the two-queue-manager scenario:

Security profiles and accesses for either a batch or CICS implementation of the two-queue-manager scenario.

Table 73, Table 74 on page 605, and Table 78 on page 607 show the security profiles that are required to enable the scenario to work:

Table 73. Security profiles for the example scenario

Class	Profile	User ID	Access
MQCONN	QM1.CHIN	MOVER1	READ
MQADMIN	QM1.RESLEVEL	BATCHID CICSAD1 MOVER1	NONE
MQADMIN	QM1.CONTEXT.**	MOVER1	CONTROL
MQQUEUE	QM1.SYSTEM.COMMAND.INPUT	MOVER1	UPDATE
MQQUEUE	QM1.SYSTEM.CHANNEL.SYNCQ	MOVER1	UPDATE
MQQUEUE	QM1.SYSTEM.CHANNEL.INITQ	MOVER1	UPDATE
MQQUEUE	QM1.SYSTEM.COMMAND.REPLY.MODEL	MOVER1	UPDATE
MQQUEUE	QM1.SYSTEM.ADMIN.CHANNEL.EVENT	MOVER1	UPDATE
MQQUEUE	QM1.QM1.TO.QM2.TCP	MOVER1	ALTER
MQQUEUE	QM1.QM1.TO.QM2.LU62	MOVER1	ALTER

Table 73. Security profiles for the example scenario (continued)

Class	Profile	User ID	Access
MQQUEUE	QM1.QM1.TO.QM2.SSL	MOVER1	ALTER
MQCONN	QM2.CHIN	MOVER2	READ
MQADMIN	QM2.RESLEVEL	MOVER2	NONE
MQADMIN	QM2.CONTEXT.**	MOVER2	CONTROL
MQQUEUE	QM2.SYSTEM.COMMAND.INPUT	MOVER2	UPDATE
MQQUEUE	QM2.SYSTEM.CHANNEL.SYNCQ	MOVER2	UPDATE
MQQUEUE	QM2.SYSTEM.CHANNEL.INITQ	MOVER2	UPDATE
MQQUEUE	QM2.SYSTEM.COMMAND.REPLY.MODEL	MOVER2	UPDATE
MQQUEUE	QM2.SYSTEM.ADMIN.CHANNEL.EVENT	MOVER2	UPDATE
MQQUEUE	QM2.DLQ	MOVER2	UPDATE

Security profiles required for a batch application:

Additional security profiles required for a batch implementation of the two-queue-manager scenario.

The batch application runs under user ID BATCHID on QM1. It connects to queue manager QM1 and puts messages to the following queues:

- LQ1
- RQA
- RQB
- RQC

It uses the MQPMO_SET_ALL_CONTEXT option. The user ID found in the *UserIdentifier* field of the message descriptor (MQMD) is MSGUSR.

The following profiles are required on queue manager QM1:

Table 74. Sample security profiles for the batch application on queue manager QM1

Class	Profile	User ID	Access
MQCONN	QM1.BATCH	BATCHID	READ
MQADMIN	QM1.CONTEXT.**	BATCHID	CONTROL
MQQUEUE	QM1.LQ1	BATCHID	UPDATE
MQQUEUE	QM1.RQA	BATCHID	UPDATE
MQQUEUE	QM1.RQB	BATCHID	UPDATE
MQQUEUE	QM1.RQC	BATCHID	UPDATE

The following profiles are required on queue manager QM2 for messages put to queue RQA on queue manager QM1 (for the TCP/IP channel not using SSL):

Table 75. Sample security profiles for queue manager QM2 using TCP/IP and not SSL

Class	Profile	User ID	Access
MQADMIN	QM2.ALTERNATE.USER.MSGUSR	MCATCP MOVER2	UPDATE
MQADMIN	QM2.CONTEXT.**	MCATCP MOVER2	CONTROL
MQQUEUE	QM2.LQA	MOVER2 MSGUSR	UPDATE
MQQUEUE	QM2.DLQ	MOVER2 MSGUSR	UPDATE

Notes:

1. The user ID passed in the MQMD of the message is used as the user ID for the **MQPUT1** on queue manager QM2 because the receiver channel was defined with PUTAUT(CTX) and MCAUSER(MCATCP).
2. The MCAUSER field of the receiver channel definition is set to MCATCP; this user ID is used in addition to the channel initiator address space user ID for the checks carried out against the alternate user ID and context profile.
3. The MOVER2 user ID and the *UserIdentifier* in the message descriptor (MQMD) are used for the resource checks against the queue.
4. The MOVER2 and MSGUSR user IDs both need access to the dead-letter queue so that messages that cannot be put to the destination queue can be sent there.
5. Two user IDs are checked on all three checks performed because RESLEVEL is set to NONE.

The following profiles are required on queue manager QM2 for messages put to queue RQB on queue manager QM1 (for the LU 6.2 channel):

Table 76. Sample security profiles for queue manager QM2 using LU 6.2

Class	Profile	User ID	Access
MQADMIN	QM2.ALTERNATE.USER.MSGUSR	MCALU62 MOVER1	UPDATE
MQADMIN	QM2.CONTEXT.**	MCALU62 MOVER1	CONTROL
MQQUEUE	QM2.LQB	MOVER1 MSGUSR	UPDATE
MQQUEUE	QM2.DLQ	MOVER1 MSGUSR	UPDATE

Notes:

1. The user ID passed in the MQMD of the message is used as the user ID for the **MQPUT1** on queue manager QM2 because the receiver channel was defined with PUTAUT(CTX) and MCAUSER(MCALU62).
2. The MCA user ID is set to the value of the MCAUSER field of the receiver channel definition (MCALU62).
3. Because LU 6.2 supports security on the communications system for the channel, the user ID received from the network is used as the channel user ID (MOVER1).
4. Two user IDs are checked on all three checks performed because RESLEVEL is set to NONE.
5. MCALU62 and MOVER1 are used for the checks performed against the alternate user ID and Context profiles, and MSGUSR and MOVER1 are used for the checks against the queue profile.
6. The MOVER1 and MSGUSR user IDs both need access to the dead-letter queue so that messages that cannot be put to the destination queue can be sent there.

The following profiles are required on queue manager QM2 for messages put to queue RQC on queue manager QM1 (for the TCP/IP channel using SSL):

Table 77. Sample security profiles for queue manager QM2 using TCP/IP and SSL

Class	Profile	User ID	Access
MQADMIN	QM2.ALTERNATE.USER.MSGUSR	MCASSL CERTID	UPDATE
MQADMIN	QM2.CONTEXT.**	MCASSL CERTID	CONTROL
MQQUEUE	QM2.LQC	CERTID MSGUSR	UPDATE
MQQUEUE	QM2.DLQ	CERTID MSGUSR	UPDATE

Notes:

1. The user ID passed in the MQMD of the message is used as the user ID for the **MQPUT1** on queue manager QM2 because the receiver channel was defined with PUTAUT(CTX) and MCAUSER(MCASSL).
2. The MCA user ID is set to the value of the MCAUSER field of the receiver channel definition (MCASSL).
3. Because the certificate flowed by the channel from QM1 as part of the SSL handshake might be installed on QM2's system, or might match a certificate name filter on QM2's system, the user ID found during that matching is used as the channel user ID (CERTID).
4. Two user IDs are checked on all three checks performed because RESLEVEL is set to NONE.
5. MCASSL and CERTID are used for the checks performed against the alternate user ID and Context profiles, and MSGUSR and MOVER1 are used for the checks against the queue profile.
6. The CERTID and MSGUSR user IDs both need access to the dead-letter queue so that messages that cannot be put to the destination queue can be sent there.

Security profiles required for a CICS application:

Additional security profiles required for a CICS implementation of the two-queue-manager scenario.

The CICS application uses a CICS address space user ID of CICSAD1 and a CICS task user ID of CICSTX1. The security profiles required on queue manager QM1 are different from those profiles required for the batch application. The profiles required on queue manager QM2 are the same as for the batch application.

The following profiles are required on queue manager QM1:

Table 78. Sample security profiles for the CICS application on queue manager QM1

Class	Profile	User ID	Access
MQCONN	QM1.CICS	CICSAD1	READ
MQADMIN	QM1.CONTEXT.**	CICSAD1 CICSTX1	CONTROL
MQQUEUE	QM1.LQ1	CICSAD1 CICSTX1	UPDATE
MQQUEUE	QM1.RQA	CICSAD1 CICSTX1	UPDATE
MQQUEUE	QM1.RQB	CICSAD1 CICSTX1	UPDATE

Security scenario: queue-sharing group on z/OS

In this scenario, an application uses the **MQPUT1** call to put messages to queues on queue manager QM1. Some of the messages are then forwarded to queues on QM2, using TCP and LU 6.2 channels. The application is a batch application, and the messages are put using the **MQPMO_SET_ALL_CONTEXT** option.

This is illustrated in Figure 86 on page 601.

The following assumptions are made about the queue managers:

- All the required WebSphere MQ definitions have been predefined or have been made through the CSQINP2 data set processed at queue manager startup.

If they have not, you need the appropriate access authority to the commands needed to define these objects.

- All the RACF profiles required have been defined and appropriate access authorities have been granted, before the queue manager and channel initiators started.

If they have not, you need the appropriate authority to issue the RACF commands required to define all the profiles needed and grant the appropriate access authorities to those profiles. You also need the appropriate authority to issue the MQSC security commands to start using the new security profiles.

Security switch settings for queue-sharing group scenario:

Switch settings and RACF profiles.

The following security switches are set for the queue-sharing group:

- Subsystem security on
- Queue-sharing group security on
- Queue manager security off
- Queue security on
- Alternate user security on
- Context security on
- Process security off
- Namelist security off
- Connection security on
- Command security on
- Command resource security on

The following profiles are defined in the MQADMIN class to turn process, namelist, and queue-manager level security off:

```
QSGA.NO.PROCESS.CHECKS
QSGA.NO.NLIST.CHECKS
QSGA.NO.QMGR.CHECKS
```

Queue manager QM1 in queue-sharing group scenario:

Queues and channels for QM1.

The following queues are defined on queue manager QM1:

LQ1 A local queue.

RQA A remote queue definition, with the following attributes:

- RNAME(LQA)
- RQMNAME(QM2)
- XMITQ(QM1.TO.QM2.TCP)

RQB A remote queue definition, with the following attributes:

- RNAME(LQB)
- RQMNAME(QM2)
- XMITQ(QM1.TO.QM2.LU62)

QM1.TO.QM2.TCP

A transmission queue.

QM1.TO.QM2.LU62

A transmission queue.

The following channels are defined on QM1:

QM1.TO.QM2.TCP

A sender channel definition, with the following attributes:

- CHLTYPE(SDR)
- TRPTYPE(TCP)
- XMITQ(QM1.TO.QM2.TCP)
- CONNAME(QM2TCP)

QM1.TO.QM2.LU62

A sender channel definition, with the following attributes:

- CHLTYPE(SDR)
- TRPTYPE(LU62)
- XMITQ(QM1.TO.QM2.LU62)
- CONNAME(QM2LU62)

(See “Security considerations for the channel initiator on z/OS” on page 590 for information about setting up APPC security.)

Queue manager QM2 in queue-sharing group scenario:

Queues and channels for QM2.

The following queues have been defined on queue manager QM2:

LQA A local queue.

LQB A local queue.

DLQ A local queue that is used as the dead-letter queue.

The following channels have been defined on QM2:

QM1.TO.QM2.TCP

A receiver channel definition, with the following attributes:

- CHLTYPE(RCVR)
- TRPTYPE(TCP)
- PUTAUT(CTX)
- MCAUSER(MCATCP)

QM1.TO.QM2.LU62

A receiver channel definition, with the following attributes:

- CHLTYPE(RCVR)
- TRPTYPE(LU62)
- PUTAUT(CTX)
- MCAUSER(MCALU62)

(See “Security considerations for the channel initiator on z/OS” on page 590 for information about setting up APPC security.)

User IDs used in queue-sharing group scenario:

Explanation of the user IDs in the scenario.

The following user IDs are used:

BATCHID

Batch application (Job or TSO ID)

MSGUSR

UserIdentifier in MQMD (context user ID)

MOVER1

QM1 channel initiator address space user ID

MOVER2

QM2 channel initiator address space user ID

MCATCP

MCAUSER specified on the TCP/IP receiver channel definition

MCALU62

MCAUSER specified on the LU 6.2 receiver channel definition

Security profiles and accesses required for queue-sharing group scenario:

Security profiles and accesses for either a batch or CICS implementation of the queue-sharing group scenario.

Table 79 through Table 82 on page 612 show the security profiles that are required to enable the scenario to work:

Table 79. Security profiles for the example scenario

Class	Profile	User ID	Access
MQCONN	QSGA.CHIN	MOVER1 MOVER2	READ
MQADMIN	QSGA.RESLEVEL	BATCHID MOVER1 MOVER2	NONE
MQADMIN	QSGA.CONTEXT.**	MOVER1 MOVER2	CONTROL

Table 79. Security profiles for the example scenario (continued)

Class	Profile	User ID	Access
MQQUEUE	QSGA.SYSTEM.COMMAND.INPUT	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.SYSTEM.CHANNEL.SYNCQ	MOVER1 MOVER	UPDATE
MQQUEUE	QSGA.SYSTEM.CHANNEL.INITQ	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.SYSTEM.COMMAND.REPLY.MODEL	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.SYSTEM.ADMIN.CHANNEL.EVENT	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.SYSTEM.QSG.CHANNEL.SYNCQ	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.SYSTEM.QSG.TRANSMIT.QUEUE	MOVER1 MOVER2	UPDATE
MQQUEUE	QSGA.QM1.TO.QM2.TCP	MOVER1	ALTER
MQQUEUE	QSGA.QM1.TO.QM2.LU62	MOVER1	ALTER
MQQUEUE	QSGA.DLQ	MOVER2	UPDATE

Security profiles required for a batch application:

Additional security profiles required for a batch implementation of the queue-sharing group scenario.

The batch application runs under user ID BATCHID on QM1. It connects to queue manager QM1 and puts messages to the following queues:

- LQ1
- RQA
- RQB

It uses the MQPMO_SET_ALL_CONTEXT option. The user ID found in the *UserIdentifier* field of the message descriptor (MQMD) is MSGUSR.

The following profiles are required on queue manager QM1:

Table 80. Sample security profiles for the batch application on queue manager QM1

Class	Profile	User ID	Access
MQCONN	QSGA.BATCH	BATCHID	READ
MQADMIN	QSGA.CONTEXT.**	BATCHID	CONTROL
MQQUEUE	QSGA.LQ1	BATCHID	UPDATE
MQQUEUE	QSGA.RQA	BATCHID	UPDATE
MQQUEUE	QSGA.RQB	BATCHID	UPDATE

The following profiles are required on queue manager QM2 for messages put to queue RQA on queue manager QM1 (for the TCP/IP channel):

Table 81. Sample security profiles for queue manager QM2 using TCP/IP

Class	Profile	User ID	Access
MQADMIN	QSGA.ALTERNATE.USER.MSGUSR	MCATCP MOVER2	UPDATE
MQADMIN	QSGA.CONTEXT.**	MCATCP MOVER2	CONTROL
MQQUEUE	QSGA.LQA	MOVER2 MSGUSR	UPDATE
MQQUEUE	QSGA.DLQ	MOVER2 MSGUSR	UPDATE

Notes:

1. The user ID passed in the MQMD of the message is used as the user ID for the **MQPUT1** on queue manager QM2 because the receiver channel was defined with PUTAUT(CTX) and MCAUSER(MCATCP).
2. The MCAUSER field of the receiver channel definition is set to MCATCP; this user ID is used in addition to the channel initiator address space user ID for the checks carried out against the alternate user ID and context profile.
3. The MOVER2 user ID and the *UserIdentifier* in the message descriptor (MQMD) are used for the resource checks against the queue.
4. The MOVER2 and MSGUSR user IDs both need access to the dead-letter queue so that messages that cannot be put to the destination queue can be sent there.
5. Two user IDs are checked on all three checks performed because RESLEVEL is set to NONE.

The following profiles are required on queue manager QM2 for messages put to queue RQB on queue manager QM1 (for the LU 6.2 channel):

Table 82. Sample security profiles for queue manager QM2 using LU 6.2

Class	Profile	User ID	Access
MQADMIN	QSGA.ALTERNATE.USER.MSGUSR	MCALU62 MOVER1	UPDATE
MQADMIN	QSGA.CONTEXT.**	MCALU62 MOVER1	CONTROL
MQQUEUE	QSGA.LQB	MOVER1 MSGUSR	UPDATE
MQQUEUE	QSGA.DLQ	MOVER1 MSGUSR	UPDATE

Notes:

1. The user ID passed in the MQMD of the message is used as the user ID for the **MQPUT1** on queue manager QM2 because the receiver channel was defined with PUTAUT(CTX) and MCAUSER(MCALU62).
2. The MCA user ID is set to the value of the MCAUSER field of the receiver channel definition (MCALU62).
3. Because LU 6.2 supports security on the communications system for the channel, the user ID received from the network is used as the channel user ID (MOVER1).
4. Two user IDs are checked on all three checks performed because RESLEVEL is set to NONE.
5. MCALU62 and MOVER1 are used for the checks performed against the alternate user ID and Context profiles, and MSGUSR and MOVER1 are used for the checks against the queue profile.
6. The MOVER1 and MSGUSR user IDs both need access to the dead-letter queue so that messages that cannot be put to the destination queue can be sent there.

Setting up WebSphere MQ MQI client security

You must consider WebSphere MQ MQI client security, so that the client applications do not have unrestricted access to resources on the server.

When running a client application, do not run the application using a user ID that has more access rights than necessary; for example, a user in the mqm group or even the mqm user itself.

Note: WebSphere MQ provides Secure Sockets Layer (SSL) support for WebSphere MQ MQI clients on the following platforms:

- IBM i
- UNIX and Linux systems
- Windows

By running an application as a user with too many access rights, you run the risk of the application accessing and changing parts of the queue manager, either by accident or maliciously.

There are two aspects to security between a client application and its queue manager server: authentication and access control.

- Authentication can be used to ensure that the client application, running as a specific user, is who they say they are. By using authentication you can prevent an attacker from gaining access to your queue manager by impersonating one of your applications.

You should use mutual authentication within SSL or TLS. For more information, see “Working with SSL or TLS” on page 616

- Access control can be used to give or remove access rights for a specific user or group of users. By running a client application with a specifically created user (or user in a specific group) you can then use access controls to ensure the application cannot access parts of your queue manager that the application is not supposed to.

When setting up access control you must consider channel authentication rules and the MCAUSER field on a channel. Both of these features have the ability to change which user id is being used for verifying access control rights.

For more information on access control, see “Authorizing access to objects” on page 704.

If you have set up a client application to connect to a specific channel with a restricted ID, but the channel has an administrator ID set in its MCAUSER field then, provided the client application connects successfully, the administrator ID is used for access control checks. Therefore, the client application will have full access rights to your queue manager.

If you have set up a client application to connect to a specific channel with a restricted ID, but the channel has an administrator ID set in its MCAUSER field then, provided the client application connects successfully, the administrator ID is used for access control checks. Therefore, the client application will have full access rights to your queue manager.

For more information on the MCAUSER attribute, see “Mapping a client asserted user ID to an MCAUSER user ID” on page 736.

Channel authentication rules can also be used as a method for controlling access to a queue manager, by setting up specific rules and criteria for a connection to be accepted.

For more information on channel authentication rules see: “Channel authentication records” on page 408.

For information on running SSL or TLS client applications on AIX with multiple GSKit V8.0 installations, see Running SSL or TLS client applications with multiple installations of GSKit v8.0 on AIX.

Related concepts:

“Planning authentication for a client application” on page 416

Related reference:

“Access control for clients” on page 432

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client

Create your key repositories using FIPS-compliant software, then specify that the channel must use FIPS-certified CipherSpecs.

In order to be FIPS-compliant at run time, the key repositories must have been created and managed using only FIPS-compliant software such as `runmqakm` with the `-fips` option.

You can specify that an SSL or TLS channel must use only FIPS-certified CipherSpecs in three ways, listed in order of precedence:

1. Set the `FipsRequired` field in the MQSCO structure to `MQSSL_FIPS_YES`.
2. Set the environment variable `MQSSLFIPS` to `YES`.
3. Set the `SSLFipsRequired` attribute in the client configuration file to `YES`.

By default, FIPS-certified CipherSpecs is not required.

These values have the same meanings as the equivalent parameter values on `ALTER QMGR SSLFIPS` (see  `ALTER QMGR`). If the client process currently has no active SSL or TLS connections, and a `FipsRequired` value is validly specified on an `SSL MQCONN`, all subsequent SSL connections associated with this process must use only the CipherSpecs associated with this value. This applies until this and all other SSL or TLS connections have stopped, at which stage a subsequent `MQCONN` can provide a new value for `FipsRequired`.

If cryptographic hardware is present, the cryptographic modules used by WebSphere MQ can be configured to be those modules provided by the hardware product, and these might be FIPS-certified to a particular level. The configurable modules and whether they are FIPS-certified depends on the hardware product in use.

Where possible, if FIPS-only CipherSpecs is configured then the MQI client rejects connections which specify a non-FIPS CipherSpec with `MQRC_SSL_INITIALIZATION_ERROR`. WebSphere MQ does not guarantee to reject all such connections and it is your responsibility to determine whether your WebSphere MQ configuration is FIPS-compliant.

Related concepts:

“Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows” on page 394

“Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows” on page 394

Related reference:

 `FipsRequired` (MQLONG) (*WebSphere MQ V7.1 Reference*)

 `MQSSLFIPS` (*WebSphere MQ V7.1 Installing Guide*)

 SSL stanza of the client configuration file (*WebSphere MQ V7.1 Installing Guide*)

 `MQSSLFIPS` (*WebSphere MQ V7.1 Installing Guide*)

 SSL stanza of the client configuration file (*WebSphere MQ V7.1 Installing Guide*)

 `FipsRequired` (MQLONG) (*WebSphere MQ V7.1 Reference*)

Running SSL or TLS client applications with multiple installations of GSKit V8.0 on AIX

SSL or TLS client applications on AIX might experience MQRC_CHANNEL_CONFIG_ERROR and error AMQ6175 when running on AIX systems with multiple GSKit V8.0 installations.

When running client applications on an AIX system with multiple GSKit V8.0 installations, the client connect calls can return MQRC_CHANNEL_CONFIG_ERROR when using SSL or TLS. The /var/mqm/errors logs record error AMQ6175 and AMQ9220 for the failing client application, for example:

```
09/08/11 11:16:13 - Process(24412.1) User(user) Program(example)
                    Host(machine.example.ibm.com) Installation(Installation1)
                    VRMF(7.1.0.0)
AMQ6175: The system could not dynamically load the shared library
'/usr/mqm/gskit8/lib64/libgsk8ssl_64.so'. The system returned
error number '8' and error message 'Symbol resolution failed
for /usr/mqm/gskit8/lib64/libgsk8ssl_64.so because:
    Symbol VALUE_EC_NamedCurve_secp256r1_9GSKASN0ID (number 16) is not
exported from dependent module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.
    Symbol VALUE_EC_NamedCurve_secp384r1_9GSKASN0ID (number 17) is not exported
from dependent module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.
    Symbol VALUE_EC_NamedCurve_secp521r1_9GSKASN0ID (number 18) is not exported
from dependent module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.
    Symbol VALUE_EC_ecPublicKey_9GSKASN0ID (number 19) is not exported from
dependent module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.
    Symbol VALUE_EC_ecdsa_with_SHA1_9GSKASN0ID (number 20) is not exported from
dependent module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.
    Symbol VALUE_EC_ecdsa_9GSKASN0ID (number 21) is not exported from dependent
module /db2data/db2inst1/sqllib/lib64/libgsk8cms_64.so.'
```

EXPLANATION:

This message applies to AIX systems. The shared library
'/usr/mqm/gskit8/lib64/libgsk8ssl_64.so' failed
to load correctly due to a problem with the library.

ACTION:

Check the file access permissions and that the file has not been corrupted.

```
----- amqxufnx.c : 1284 -----
09/08/11 11:16:13 - Process(24412.1) User(user) Program(example)
                    Host(machine.example.ibm.com) Installation(Installation1)
                    VRMF(7.1.0.0)
```

AMQ9220: The GSKit communications program could not be loaded.

EXPLANATION:

The attempt to load the GSKit library or procedure
'/usr/mqm/gskit8/lib64/libgsk8ssl_64.so' failed with error code
536895861.

ACTION:

Either the library must be installed on the system or the environment changed
to allow the program to locate it.

```
----- amqcgksa.c : 836 -----
```

A common cause of this error is that the setting of the LIBPATH or LD_LIBRARY_PATH environment variable has caused the IBM WebSphere MQ client to load a mixed set of libraries from two different GSKit V8.0 installations. Executing an IBM WebSphere MQ client application in a Db2 environment can cause this error.

To avoid this error, include the IBM WebSphere MQ library directories at the front of the library path so that the IBM WebSphere MQ libraries take precedence. This can be achieved using the **setmqenv** command with the **-k** parameter, for example:

```
./usr/mqm/bin/setmqenv -s -k
```

For more information about the use of the **setmqenv** command, refer to  **setmqenv** (set WebSphere MQ environment) (*WebSphere MQ V7.1 Reference*)

Related reference:

 **GSKit: Changes from GSKit V7.0 to GSKit V8.0** (*WebSphere MQ V7.1 Installing Guide*)

 **GSKit: Commands renamed** (*WebSphere MQ V7.1 Installing Guide*)

Working with SSL or TLS

These topics give instructions for performing single tasks related to using SSL or TLS with IBM WebSphere MQ.

Many of them are used as steps in the higher-level tasks described in the following sections:

- “Identifying and authenticating users” on page 689
- “Authorizing access to objects” on page 704
- “Confidentiality of messages” on page 761
- “Data integrity of messages” on page 781
- “Keeping clusters secure” on page 801

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

“Security protocols in WebSphere MQ” on page 386

Working with SSL or TLS on HP Integrity NonStop Server

Describes the IBM WebSphere MQ client for HP Integrity NonStop Server OpenSSL security implementation, including security services, components, supported protocol versions, supported CipherSpecs, and unsupported security functionality.

IBM WebSphere MQ SSL & TLS support provides the following security services for client channels:

- Authentication of the server and, optionally, authentication of the client.
- Encryption and decryption of the data that is flowing across a channel.
- Integrity checks on the data that is flowing across a channel.

The SSL and TLS support supplied with the IBM WebSphere MQ client for HP Integrity NonStop Server comprises the following components:

- OpenSSL libraries and the **openssl** command.
- IBM WebSphere MQ password stash command, **amqrs1c**.

The following required components for SSL or TLS client channel operation are not provided with the IBM WebSphere MQ client for HP Integrity NonStop Server:

- An entropy daemon to provide a source of random data for OpenSSL cryptography.

Supported protocol versions

The IBM WebSphere MQ client for HP Integrity NonStop Server supports the following protocol versions:

- SSL 3.0
- TLS 1.0
- TLS 1.2

Supported CipherSpecs

The IBM WebSphere MQ client for HP Integrity NonStop Server supports the following CipherSpecs versions:

- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- RC4_SHA_US
- RC4_MD5_US
- TRIPLE_DES_SHA_US
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- DES_SHA_EXPORT1024
- RC4_56_SHA_EXPORT1024
- RC4_MD5_EXPORT
- RC2_MD5_EXPORT
- DES_SHA_EXPORT
- TLS_RSA_WITH_DES_CBC_SHA
- NULL_SHA
- NULL_MD5
- FIPS_WITH_DES_CBC_SHA
- FIPS_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_NULL_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- ECDHE_ECDSA_AES_128_CBC_SHA256
- ECDHE_ECDSA_AES_256_CBC_SHA384
- ECDHE_RSA_AES_128_CBC_SHA256
- ECDHE_RSA_AES_256_CBC_SHA384
- ECDHE_ECDSA_AES_128_GCM_SHA256
- ECDHE_ECDSA_AES_256_GCM_SHA384
- ECDHE_RSA_AES_128_GCM_SHA256
- ECDHE_RSA_AES_256_GCM_SHA384

Unsupported security functionality

The IBM WebSphere MQ client for HP Integrity NonStop Server does not currently support:

- PKCS#11 Cryptographic hardware support
- LDAP Certificate Revocation List checking
- OCSP Online Certificate Status Protocol checking
- FIPS 140-2, NSA SUITE B cipher suite controls

Certificate management:

Use a set of files to store digital certificate and certificate revocation information.

IBM WebSphere MQ SSL and TLS support uses a set of files to store digital certificate and certificate revocation information. These files are located in a directory specified either programmatically by way of the KeyRepository field in the MQSCO structure passed on the MQCONN call, by the *MQSSLKEYR* environment variable, or, in the SSL stanza of the *mqclient.ini* using the SSLKeyRepository attribute.

The MQSCO structure takes precedence over the MQSSLKEYR environment variable which takes precedence over the ini file stanza value.

Important: The key repository location specifies a directory location and not a filename on the HP Integrity NonStop Server platform.

The IBM WebSphere MQ client for HP Integrity NonStop Server uses the following, case sensitive, named files in the key repository location:

- "Personal certificate store"
- "Certificate trust store"
- "Pass phrase stash file" on page 619
- "Certificate revocation list file" on page 619

Personal certificate store:

The personal certificate store file, *cert.pem*.

This file contains the personal certificate and the encrypted private key for the client to use, in PEM format. The existence of this file is optional when you are using SSL or TLS channels that do not require client authentication. Where client authentication is required by the channel, and SSLCAUTH(REQUIRED) is specified on the channel definition, this file must exist and contain both the certificate and encrypted private key.

File permissions must be set on this file to allow read access to the owner of the certificate store.

A correctly formatted *cert.pem* file must contain exactly two sections with the following headers and footers:

```
-----BEGIN PRIVATE KEY-----  
Base 64 ASCII encoded private key data here  
-----END PRIVATE KEY-----  
-----BEGIN CERTIFICATE-----  
Base 64 ASCII encoded certificate data here  
-----END CERTIFICATE-----
```

The pass phrase for the encrypted private key is stored in the pass phrase stash file, *Stash.sth*.

Certificate trust store:

The certificate truststore file, *trust.pem*.

This file contains the certificates that are needed to validate the personal certificates that are used by queue managers that the client connects to, in PEM format. The certificate truststore is mandatory for all SSL or TLS client channels.

File permissions must be set to limit write access to this file.

A correctly formatted `trust.pem` file must contain one or more sections with the following headers and footers:

```
-----BEGIN CERTIFICATE-----  
Base 64 ASCII encoded certificate data here  
-----END CERTIFICATE-----
```

Pass phrase stash file:

The pass phrase stash file, `Stash.sth`.

This file is a binary format private to IBM WebSphere MQ and contains the encrypted pass phrase for use when you are accessing the private key that is held in the `cert.pem` file. The private key itself is stored in the `cert.pem` certificate store.

This file is created or altered by using the IBM WebSphere MQ `amqrsslc` command-line tool with the `-s` parameter. For example, where the directory `/home/alice` contains a `cert.pem` file:

```
amqrsslc -s /home/alice/cert
```

```
Enter password for Keystore /home/alice/cert.pem :  
password
```

```
Stashed the password in file /home/alice/Stash.sth
```

File permissions must be set on this file to allow read access to the owner of the associated personal certificate store.

Certificate revocation list file:

The certificate revocation list file, `cr1.pem`.

This file contains the certificate revocation lists (CRLs) that the client uses to validate digital certificates, in PEM format. The existence of this file is optional. If this file is not present, no certificate revocation checks are done when you are validating certificates.

File permissions must be set to limit write access to this file.

A correctly formatted `cr1.pem` file must contain one or more sections with the following headers and footers:

```
-----BEGIN X509 CRL-----  
Base 64 ASCII encoded CRL data here  
-----END X509 CRL-----
```

Working with SSL or TLS on IBM i

This collection of topics gives instructions for individual tasks working with the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) in IBM WebSphere MQ for IBM i.

For IBM i the SSL support is integral to the operating system. Ensure that you have installed the prerequisites listed in  *Hardware and software requirements on IBM i (WebSphere MQ V7.1 Installing Guide)*.

On IBM i, you manage keys and digital certificates with the Digital Certificate Manager (DCM) tool.

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

Accessing DCM:

Follow these instructions to access the DCM interface.

About this task

Perform the following steps in a web browser that supports frames.

Procedure

1. Go to either `http://machine.domain:2001` or `https://machine.domain:2010`, where *machine* is the name of your computer.
2. Type a valid user profile and password when requested to. Ensure that your user profile has *ALLOBJ and *SECADM special authorities to enable you to create new certificate stores. If you do not have the special authorities, you can only manage your personal certificates or view the object signatures for the objects for which you are authorized. If you are authorized to use an object signing application, you can also sign objects from DCM.
3. On the Internet Configurations page, click **Digital Certificate Manager**. The Digital Certificate Manager page is displayed.

Assigning a certificate to a queue manager on IBM i:

Use DCM to assign a certificate to a queue manager.

Use traditional IBM i digital certificate management to assign a certificate to a queue manager. This means that you can specify that a queue manager uses the system certificate store, and that the queue manager is registered for use as an application with Digital Certificate Manager. To do this you change the value of the queue manager's SSLKEYR attribute to *SYSTEM.

When the SSLKEYR parameter is changed to *SYSTEM, WebSphere MQ registers the queue manager as a server application with a unique application label of QIBM_WEBSHERE_MQ_QMGRNAME and a label with a description of Qmgrname (WMQ). The queue manager then appears as a server application in Digital Certificate Manager, and you can assign to this application any server or client certificate in the system store.

Because the queue manager is registered as an application, advanced features of DCM such as defining CA trust lists can be carried out.

If the SSLKEYR parameter is changed to a value other than *SYSTEM, WebSphere MQ deregisters the queue manager as an application with Digital Certificate Manager. If a queue manager is deleted, it is also deregistered from DCM. A user with sufficient *SECADM authority can also manually add or remove applications from DCM.

Setting up a key repository on IBM i:

A key repository must be set up at both ends of the connection. The default certificate stores can be used or you can create your own.

An SSL or TLS connection requires a *key repository* at each end of the connection. Each queue manager and IBM WebSphere MQ MQI client must have access to a key repository. If you want to access the key repository using a file name and password (that is, not using the *SYSTEM option) ensure:

- the QMQM user profile has execute authority for the directory containing the key repository
- the QMQM user profile has read authority for the file containing the key repository

See “The SSL or TLS key repository” on page 390 for more information.

On IBM i, digital certificates are stored in a certificate store that is managed with DCM. These digital certificates have labels, which associate a certificate with a queue manager or a IBM WebSphere MQ MQI client. SSL and TLS use the certificates for authentication purposes.

WebSphere MQ uses the `ibmwebspheremq` prefix on a label to avoid confusion with certificates for other products. The prefix is followed by the name of the queue manager or IBM WebSphere MQ MQI client user logon ID, changed to lower case. Ensure that the entire certificate label is in lowercase.

The queue manager or IBM WebSphere MQ MQI client certificate store name comprises a path and stem name. The default path is `/QIBM/UserData/ICSS/Cert/Server/` and the default stem name is `Default`. On IBM i, the default certificate store, `/QIBM/UserData/ICSS/Cert/Server/Default.kdb`, is also known as `*SYSTEM`. Optionally, you can define your own path and stem name.

If you define your own path or file name, set the permissions to the file to tightly control access to it.

“Changing the key repository location for a queue manager on IBM i” on page 622 tells you about specifying the certificate store name. You can specify the certificate store name either before or after creating the certificate store.

Note: The operations you can perform with DCM might be limited by the authority of your user profile. For example, you require `*ALLOBJ` and `*SECADM` authorities to create a CA certificate.

Creating a certificate store on IBM i:

If you do not want to use the default certificate store, follow this procedure to create your own.

About this task

Create a new certificate store only if you do not want to use the IBM i default certificate store.

To specify that the IBM i system certificate store is to be used, change the value of the queue manager's `SSLKEYR` attribute to `*SYSTEM`. This value indicates that the queue manager uses the system certificate store, and the queue manager is registered for use as an application with Digital Certificate Manager (DCM).

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620
2. In the navigation panel, click **Create New Certificate Store**. The Create New Certificate Store page is displayed in the task frame.
3. In the task frame, select **Other System Certificate Store** and click **Continue**. The Create a Certificate in New Certificate Store page is displayed in the task frame.
4. Select **No - Do not create a certificate in the certificate store** and click **Continue**. The Certificate Store Name and Password page is displayed in the task frame.
5. In the **Certificate store path and filename** field, type an IFS path and file name, for example `/QIBM/UserData/mqm/qmgrs/qm1/key.kdb`
6. Type a password in the **Password** field and type it again in the **Confirm Password** field. Click **Continue**. Make a note of the password (which is case sensitive) because you need it when you stash the repository key.
7. To exit from DCM, close your browser window.

What to do next

When you have created the certificate store using DCM, ensure you stash the password, as described in “Stashing the certificate store password on IBM i”

Stashing the certificate store password on IBM i:

Stash the certificate store password using CL commands.

The following instructions apply to stashing the certificate store password on IBM i for a queue manager. Alternatively, for a IBM WebSphere MQ MQI client, if you are not using the *SYSTEM certificate store (that is, the MQSSLKEYR environment is set to a value other than *SYSTEM), follow the procedure described in the "Stash the certificate store password" section of WebSphere MQ SSL Client utility (amqrsslc) for IBM i.

If you have specified that the *SYSTEM certificate store is to be used (by changing the value of the SSLKEYR attribute of the queue manager to *SYSTEM) you must not follow these steps.

When you have created the certificate store using DCM, use the following commands to stash the password:

```
STRMQM MQMNAME('queue manager name')
```

```
CHGMQM MQMNAME('queue manager name') SSLKEYRPWD('password')
```

The password must be entered as a literal (in single quotation marks) exactly as you entered it in step 6 of “Creating a certificate store on IBM i” on page 621 (it is case-sensitive).

Note: If you are not using the default system certificate store, and you do not stash the password, attempts to start SSL channels fail because they cannot obtain the password required to access the certificate store.

Locating the key repository for a queue manager on IBM i:

Use this procedure to obtain the location of your queue manager's certificate store.

Procedure

1. Display your queue manager's attributes, using the following command:
DSPMQM MQMNAME('queue manager name')
2. Examine the command output for the path and stem name of the certificate store. For example: /QIBM/UserData/ICSS/Cert/Server/Default, where /QIBM/UserData/ICSS/Cert/Server is the path and Default is the stem name.

Changing the key repository location for a queue manager on IBM i:

Change the location of your queue manager's certificate store using either CHGMQM or ALTER QMGR.

Procedure

Use either the CHGMQM command or the ALTER QMGR MQSC command to set your queue manager's key repository attribute.

1. Using CHGMQM: CHGMQM MQMNAME('qm1') SSLKEYR('/QIBM/UserData/ICSS/Cert/Server/MyKey')
2. Using ALTER QMGR: ALTER QMGR SSLKEYR('/QIBM/UserData/ICSS/Cert/Server/MyKey')

In either case, the certificate store has the fully qualified file name: /QIBM/UserData/ICSS/Cert/Server/MyKey.kdb

What to do next

When you change the location of a queue manager's certificate store, certificates are not transferred from the old location. If the CA certificates preinstalled when you create the certificate store are insufficient, you must populate the new certificate store with certificates, as described in "Importing a certificate into a key repository on IBM i" on page 627. You must also stash the password for the new location, as described in "Stashing the certificate store password on IBM i" on page 622.

Creating a certificate authority and certificate for testing on IBM i:

Use this procedure to create a local CA certificate to sign certificate requests, and to create and install the CA certificate.

Before you begin

The instructions in this topic assume that a local certificate authority (CA) does not exist. If a local CA does exist, go to "Requesting a server certificate on IBM i" on page 624.

About this task

The CA certificates that are provided when you install SSL are signed by the issuing CA. On IBM i, you can generate a local certificate authority that can sign server certificates for testing SSL communications on your system. Follow these steps in a Web browser to create a local CA certificate:

Procedure

1. Access the DCM interface, as described in "Accessing DCM" on page 620.
2. In the navigation panel, click **Create a Certificate Authority**. The Create a Certificate Authority page is displayed in the task frame.
3. Type a password in the **Certificate store password** field and type it again in the **Confirm password** field.
4. Type a name in the **Certificate Authority (CA) name** field, for example SSL Test Certificate Authority.
5. Type appropriate values in the **Common Name** and **Organization** fields, and select a country. For the remaining optional fields, type the values you require.
6. Type a validity period for the local CA in the **Validity period** field. The default value is 1095 days.
7. Click **Continue**. The CA is created, and DCM creates a certificate store and a CA certificate for your local CA.
8. Click **Install certificate**. The download manager dialog box is displayed.
9. Type the full path name for the temporary file in which you want to store the CA certificate and click **Save**.
10. When download is complete, click **Open**. The Certificate window is displayed.
11. Click **Install certificate**. The Certificate Import wizard is displayed.
12. Click **Next**.
13. Select **Automatically select the certificate store based on the type of certificate** and click **Next**.
14. Click **Finish**. A confirmation window is displayed.
15. Click **OK**.
16. In the Certificate window, click **OK**.
17. Click **Continue**. The Certificate Authority Policy page is displayed in the task frame.
18. In the **Allow creation of user certificates** field, select **Yes**.
19. In the **Validity period** field, type the validity period of certificates that are issued by your local CA. The default value is 365 days.

20. Click **Continue**. The Create a Certificate in New Certificate Store page is displayed in the task frame.
21. Check that none of the applications are selected.
22. Click **Continue** to complete the setup of the local CA.

Requesting a server certificate on IBM i:

Digital certificates protect against impersonation, certifying that a public key belongs to a specified entity. A new server certificate can be requested from a certificate authority using the Digital Certificate Manager (DCM).

About this task

Perform the following steps in a Web browser:

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
3. Select the certificate store you want to use and click **Continue**.
4. Optional: If you selected ***SYSTEM** in step 3, enter the system store password and click **Continue**.
5. Optional: If you selected **Other System Certificate Store** in step 3, in the **Certificate store path and filename** field, type the IFS path and file name you set when you created your certificate store. Also type a password in the **Certificate Store Password** field. Then click **Continue**.
6. In the navigation panel, click **Create Certificate**.
7. In the task frame, select the **Server or client certificate** radio button and click **Continue**. The Select a Certificate Authority (CA) page is displayed in the task frame.
8. If you have a local CA on your workstation you choose either the local CA or a commercial CA to sign the certificate. Select the radio button for the CA you want and click **Continue**. The Create a Certificate page is displayed in the task frame.
9. Optional: For a queue manager, in the **Certificate label** field, type `ibmwebsphermq` followed by the name of your queue manager folded to lowercase. For example, for queue manager QM1, type `ibmwebsphermqm1`.
10. Optional: For a IBM WebSphere MQ MQI client, in the **Certificate label** field, type `ibmwebsphermq` followed by your logon user ID folded to lowercase. For example, type `ibmwebsphermqmyuserID`.
11. Type appropriate values in the **Common Name** and **Organization** fields, and select a country. For the remaining optional fields, type the values you require.

Results

If you selected a commercial CA to sign your certificate, DCM creates a certificate request in PEM (Privacy-Enhanced Mail) format. Forward the request to your chosen CA.

If you selected the local CA to sign your certificate, DCM informs you that the certificate has been created in the certificate store and can be used.

Related tasks:

“Creating a certificate store on IBM i” on page 621

Requesting a server certificate on IBM i for IBM Key Manager:

Follow this procedure to create a certificate signed by your local certificate authority (CA), or to apply for a server certificate signed by a commercial CA for import into the IBM Key Management (iKeyman) utility.

About this task

A user certificate must be used when the Digital Certificate Manager (DCM) serves as the certificate manager for WebSphere MQ on multiple platforms. For personal certificates distributed to other platforms and for import into the iKeyman utility, perform the following steps in a Web browser:

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the navigation pane, click **Create Certificate**. The Create Certificate page is displayed in the task frame.
3. On the Create Certificate panel, select the **User certificate** radio button and click **Continue**. The Create User Certificate page is displayed.
4. On the Create User Certificate panel, complete the required fields under Certificate Information for **Organization name**, **State or province**, **Country** or **region**. Optionally, put values in the **Organization unit** and **Locality** or **city** fields. Click **Continue**. The **Common name** is automatically set to the user ID with which you are logged on to the iSeries system.
5. On the next Create User Certificate panel, click **Install certificate** and click **Continue**. A message is displayed stating, Your personal certificate has been installed. You should keep a backup copy of this certificate.
6. Click **OK**.
7. Depending on the internet browser you used to access DCM, do the following steps:
 - a. For Internet Explorer choose: **Tools>Internet Options>Content tab>Certificates button>Personal tab>**. Select the certificate and click **Export**.
 - b. For Mozilla Firefox choose: **Tools>Options>Advanced>Encryption tab>View Certificates button>Your Certificates tab>**. Select the certificate and click **Backup**. Select the path and filename and click **OK**.
8. Transfer the exported certificate to the remote system using FTP in binary format.
9. Add the exported certificate from step 7 to the iKeyman utility in the key database.
 - a. If the certificate was saved using Internet Explorer, use the instructions described in Importing from a Microsoft .pfx.
 - b. If the certificate was saved using Mozilla Firefox, use the instructions described in Importing a personal certificate into a key repository.

During the import, ensure that the label name of the personal certificate and the signer certificate are changed to reflect the naming convention which WebSphere MQ is expecting. For example, the personal certificate looks like *userid's CN id*. It must be changed to *ibmwebspheremqXXXXXX* where XXXXXX is the name of your queue manager in lowercase letters.

Adding server certificates to a key repository on IBM i:

Follow this procedure to add a requested certificate to the key repository.

About this task

After the CA sends you a new server certificate, you add it to the certificate store from which you generated the request. If the CA sends the certificate as part of an email message, copy the certificate into a separate file.

Note:

- You do not need to perform this procedure if the server certificate is signed by your local CA.
- Before you import a server certificate in PKCS #12 format into DCM, you must first import the corresponding CA certificate.

Use the following procedure to receive a server certificate into the queue manager certificate store:

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the **Manage Certificates** task category in the navigation panel, click **Import Certificate**. The Import Certificate page displays in the task frame.
3. Select the radio button for your certificate type and click **Continue**. Either the Import Server or Client Certificate page or the Import Certificate Authority (CA) Certificate page displays in the task frame.
4. In the **Import File** field, type the file name of the certificate you want to import and click **Continue**. DCM automatically determines the format of the file.
5. If the certificate is a **Server or client** certificate, type the password in the task frame and click **Continue**. DCM informs you that the certificate has been imported.

Exporting a certificate from a key repository on IBM i:

Exporting a certificate exports both the public and private key. This action should be taken with extreme caution, since passing on a private key would completely compromise your security.

Before you begin

When you share a user's certificate with another user, you exchange public keys. When you export a certificate as described here, you export both the public and private key. This action should be taken with extreme caution, since passing on a private key would completely compromise your security.

About this task

Perform the following steps on the computer from which you want to export the certificate:

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
3. Select the certificate store you want to use and click **Continue**.
4. Optional: If you selected ***SYSTEM** in step 3, enter the system store password and click **Continue**.
5. Optional: If you selected **Other System Certificate Store** in step 3, in the **Certificate store path and filename** field, type the IFS path and file name you set when you created your certificate store and type a password in the **Certificate Store Password** field. Then click **Continue**

6. In the **Manage Certificates** task category in the navigation panel, click **Export Certificate**. The Export a Certificate page is displayed in the task frame.
7. Select the radio button for your certificate type and click **Continue**. Either the Export Server or Client Certificate page or the Export Certificate Authority (CA) Certificate page is displayed in the task frame.
8. Select the certificate you want to export.
9. Select the radio button to specify whether you want to export the certificate to a file or directly into another certificate store.
10. If you selected to export a server or client certificate to a file, provide the following information:
 - The path and file name of the location where you want to store the exported certificate.
 - For a personal certificate, the password that is used to encrypt the exported certificate and the target release. For CA certificates, you do not need to specify the password.
11. If you selected to export a certificate directly into another certificate store, specify the target certificate store and its password.
12. Click **Continue**.

Related tasks:

“Creating a certificate store on IBM i” on page 621

Importing a certificate into a key repository on IBM i:

Follow this procedure to import a certificate.

Before you begin

Before you import a personal certificate in PKCS #12 format into DCM, you must first import the corresponding CA certificate.

About this task

Perform these steps on the machine to which you want to import the certificate.

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
3. Select the certificate store you want to use and click **Continue**.
4. Optional: If you selected ***SYSTEM** in step 3, enter the system store password and click **Continue**.
5. Optional: If you selected **Other System Certificate Store** in step 3, in the **Certificate store path and filename** field, type the IFS path and file name you set when you created your certificate store and type a password in the **Certificate Store Password** field. Then click **Continue**
6. In the **Manage Certificates** task category in the navigation panel, click **Import Certificate**. The Import Certificate page is displayed in the task frame.
7. Select the radio button for your certificate type and click **Continue**. Either the Import Server or Client Certificate page or the Import Certificate Authority (CA) Certificate page is displayed in the task frame.
8. In the **Import File** field, type the file name of the certificate you want to import and click **Continue**. DCM automatically determines the format of the file.
9. If the certificate is a **Server or client** certificate, type the password in the task frame and click **Continue**. DCM informs you that the certificate has been imported.

Related tasks:

“Creating a certificate store on IBM i” on page 621

Removing certificates in IBM i:

Use this procedure to remove personal certificates.

Procedure

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
3. Select the **Other System Certificate Store** check box and click **Continue**. The Certificate Store and Password page is displayed.
4. In the **Certificate store path and filename** field, type the IFS path and file name you set when you created the certificate store.
5. Type a password in the **Certificate Store Password** field. Click **Continue**. The Current Certificate Store page is displayed in the task frame.
6. In the **Manage Certificates** task category in the navigation panel, click **Delete Certificate**. The Confirm Delete Certificate page is displayed in the task frame.
7. Select the certificate you want to delete. Click **Delete**.
8. Click **Yes** to confirm that you want to delete the certificate. Otherwise, click **No**. DCM informs you if it has deleted the certificate.

Using the *SYSTEM certificate store for one-way authentication on IBM i:

Follow these instructions to set up one-way authentication.

Before you begin

- You have created a queue manager, channels, and transmission queues.
- You have created a server or client certificate on the server queue manager.
- You have transferred the CA certificate to the client queue manager and imported it into the key repository.
- You have started a listener on the server and client queue managers.

About this task

To use one-way authentication, using a computer running IBM i as the SSL server, set the SSL Key Repository (SSLKEYR) parameter to *SYSTEM. This setting registers the WebSphere MQ queue manager as an application. You can then assign a certificate to the queue manager to enable one-way authentication.

You can also use private keystores to implement one-way authentication by creating a dummy certificate for the client queue manager in the key repository.

Procedure

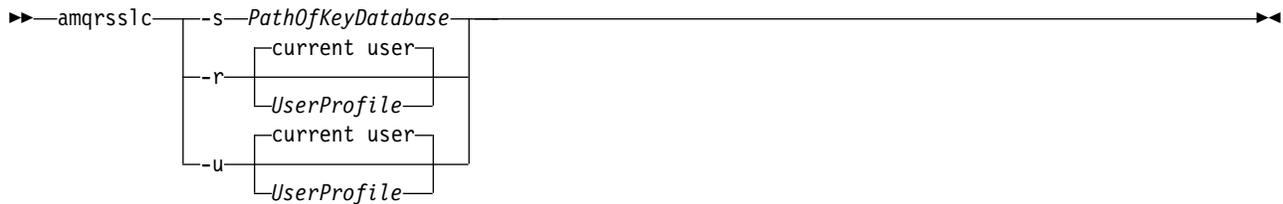
1. Perform the following steps on the server and client queue managers:
 - a. Alter the queue manager to set the SSLKEYR parameter by issuing the command `CHGMQM QMNAME(SSL) SSLKEYR(*SYSTEM)`.
 - b. Stash the password for the default key repository by issuing the command `CHGMQM QMNAME(SSL) SSLKEYRPWD('xxxxxxx')`. The password must be in single quotation marks.
 - c. Alter the channels to have the correct CipherSpec in the SSLCIPHER parameter.
 - d. Refresh SSL security by issuing the command `RFRMQMAUT QMNAME(QMGRNAME) TYPE(*SSL)`.

2. Assign the certificate to the server queue manager using DCM, as follows:
 - a. Access the DCM interface, as described in “Accessing DCM” on page 620.
 - b. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
 - c. Select the *SYSTEM certificate store and click **Continue**.
 - d. In the left panel, expand **Manage Applications**.
 - e. Select the **View Application** definition to check that the queue manager has been registered as an application. SSL (WMQ) is listed in the table.
 - f. Select **Update Certificate Assignment**.
 - g. Select **Server** and click **Continue**.
 - h. Select QMGRNAME (WMQ) and click **Update certificate assignment**.
 - i. Select the certificate and click **Assign New Certificate**. A window opens stating that the certificate has been assigned to the application.

WebSphere MQ SSL Client utility (amqrssl) for IBM i:

The WebSphere MQ SSL Client utility (amqrssl) for IBM i is used by the IBM WebSphere MQ MQI client on IBM i systems to register or unregister the client user profile, or stash the certificate store password. The utility can only be run by a user with a profile with *ALLOBJ special authority or a member of QMQMADM that has options to create or delete application registrations in the Digital Certificate Manager (DCM).

Syntax diagram



Register the client user profile

If the IBM WebSphere MQ MQI client is using the *SYSTEM certificate store, you must register the client user profile (logon user) for use as an application with Digital Certificate Manager (DCM).

If you want to register the client user profile, run the **amqrssl** program with the **-r** option with *UserProfile*. The user profile used when calling **amqrssl** must have *USE authority. Providing *UserProfile* with the **-r** option registers the *UserProfile* as a server application with a unique application label of QIBM_WEBSPPHERE_MQ_<UserProfile> and a label with a description of <UserProfile> (WMQ). This server application then is displayed in the DCM, and you can assign to this application any server or client certificate in the system store.

Note: If a user profile is not specified with **-r** option, then the user profile of the user running the **amqrssl** tool is registered.

The following code uses **amqrssl** to register a user profile. In the first example, the specified user profile is registered; in the second it is the profile of the logged in user:

```
CALL PGM(QMQM/AMQRSSLC) PARM('-r' UserProfile)
CALL PGM(QMQM/AMQRSSLC) PARM('-r')
```

Unregister the client user profile

To unregister the client profile, run the **amqrssl** program with the **-u** option with *UserProfile*. The user profile used when calling **amqrssl** must have *USE authority. Providing the *UserProfile* with the **-u** option unregisters *UserProfile* with label QIBM_WEBSPHERE_MQ_<UserProfile> from the DCM.

Note: If a user profile is not specified with **-u** option, then the user profile of the user running the **amqrssl** tool is unregistered.

The following code uses **amqrssl** to unregister a user profile. In the first example, the specified user profile is unregistered; in the second it is the profile of the logged in user:

```
CALL PGM(QMQM/AMQRSSLC) PARM('-u' UserProfile)
CALL PGM(QMQM/AMQRSSLC) PARM('-u')
```

Stash the certificate store password

If the IBM WebSphere MQ MQI client is not using the *SYSTEM certificate store and using another certificate store (that is, MQSSLKEYR is set to value other than *SYSTEM), then the password of the key database must be stashed. Use **-s** option for stashing the password of key database.

In the following code, the fully qualified file name of the certificate store is /Path/Of/KeyDatabase/MyKey.kdb:

```
CALL PGM(QMQM/AMQRSSLC) PARM('-s' '/Path/Of/KeyDatabase/MyKey')
```

Running this code results in a request for the password of this key database. This password is stashed in a file with the same name as key database with a .sth extension. This file is stored on the same path as the key database. The code example generates a stash file of /Path/Of/KeyDatabase/MyKey.sth. QMQM is the user owner and QMQMADM the group owner for this file. QMQM and QMQMADM have read, write permission, and other profiles have only read permission.

When changes to certificates or the certificate store become effective on IBM i:

When you change the certificates in a certificate store, or the location of the certificate store, the changes take effect depending on the type of channel and how the channel is running.

Changes to the certificates in the certificate store and to the key repository attribute become effective in the following situations:

- When a new outbound single channel process first runs an SSL channel.
- When a new inbound TCP/IP single channel process first receives a request to start an SSL channel.
- When the MQSC command REFRESH SECURITY TYPE(SSL) is issued to refresh the WebSphere MQ SSL environment.
- For client application processes, when the last SSL connection in the process is closed. The next SSL connection picks up the certificate changes.
- For channels that run as threads of a process pooling process (amqrmppa), when the process pooling process is started or restarted and first runs an SSL channel. If the process pooling process has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).
- For channels that run as threads of the channel initiator, when the channel initiator is started or restarted and first runs an SSL channel. If the channel initiator process has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).

- For channels that run as threads of a TCP/IP listener, when the listener is started or restarted and first receives a request to start an SSL channel. If the listener has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).

Configuring cryptographic hardware on IBM i:

Use this procedure to configure the 4758 PCI Cryptographic Coprocessor on IBM i

Before you begin

Ensure your user profile has *ALLOBJ and *SECADM special authorities to enable you to configure the coprocessor hardware.

Procedure

1. Go to either `http://machine.domain:2001` or `https://machine.domain:2010`, where *machine* is the name of your computer. A dialog box is displayed, requesting a user name and a password.
2. Type a valid IBM i user profile and password.
3. On the AS/400 Tasks page, click **4758 PCI Cryptographic Coprocessor**.

What to do next

For more information about configuring the 4758 PCI Cryptographic Coprocessor, refer to the IBM i product documentation here: [IBM i](#).

Working with SSL or TLS on UNIX, Linux and Windows systems

This information describes how you set up and work with the Secure Sockets Layer (SSL) on UNIX, Linux and Windows systems.

This collection of topics applies to the following:

- WebSphere MQ for AIX
- WebSphere MQ for HP-UX
- WebSphere MQ for Linux
- WebSphere MQ for Solaris
- WebSphere MQ for Windows

For these platforms, the SSL support is installed with WebSphere MQ.

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

Using iKeyman, iKeycmd, runmqakm, and runmqckm:

On UNIX, Linux and Windows systems, manage keys and digital certificates with the iKeyman GUI or from the command line using iKeycmd or runmqakm.

- For **UNIX and Linux** systems:
 - Use the **strmqikm** command to start the iKeyman GUI.
 - Use the **runmqckm** command to perform tasks with the iKeycmd command line interface.
 - Use the **runmqakm** command to perform tasks with the runmqakm command line interface. The command syntax for **runmqakm** is the same as the syntax for **runmqckm**.

If you need to manage SSL certificates in a way that is FIPS compliant, use the **runmqakm** command instead of the **runmqckm** or **strmqikm** commands.

See  Managing keys and certificates (*WebSphere MQ V7.1 Reference*) for a full description of the command line interfaces for the **runmqckm** and **runmqakm** commands.

If you are using certificates or keys stored on PKCS #11 cryptographic hardware, note that iKeycmd and iKeyman are 64-bit programs. External modules required for PKCS #11 support will be loaded into a 64-bit process, therefore you must have a 64-bit PKCS #11 library installed for the administration of cryptographic hardware. The Windows and Linux x86 32-bit platforms are the only exceptions, as the iKeyman and iKeycmd programs are 32-bit on those platforms.

See  GSKit: PKCS#11 and JRE addressing mode (*WebSphere MQ V7.1 Installing Guide*) for further information.

Before you run the **strmqikm** command to start the iKeyman GUI, ensure you are working on a machine that is able to run the X Window System and that you do the following:

- Set the DISPLAY environment variable, for example:
`export DISPLAY=myipc:0`
 - Ensure that your PATH environment variable contains **/usr/bin** and **/bin**. This is also required for the **runmqckm** and **runmqakm** commands. For example:
`export PATH=$PATH:/usr/bin:/bin`
 - For **Windows** systems:
 - Use the **strmqikm** command to start the iKeyman GUI.
 - Use the **runmqckm** command to perform tasks with the iKeycmd command line interface.
- If you need to manage SSL certificates in a way that is FIPS compliant, use the **runmqakm** command instead of the **runmqckm** or **strmqikm** commands.

To request SSL tracing on UNIX, Linux or Windows systems, see  strmqtrc (*WebSphere MQ V7.1 Reference*).

Related reference:

 runmqckm, and runmqakm commands (*WebSphere MQ V7.1 Reference*)

Setting up a key repository on UNIX, Linux, and Windows systems:

You can set up a key repository by using the iKeyman user interface, or by using the **iKeycmd** or **runmqakm** commands.

An SSL or TLS connection requires a *key repository* at each end of the connection. Each IBM WebSphere MQ queue manager and IBM WebSphere MQ MQI client must have access to a key repository. For more information, see “The SSL or TLS key repository” on page 390.

On UNIX, Linux, and Windows systems, digital certificates are stored in a key database file that is managed by using the **iKeyman** user interface, or by using the **iKeycmd** or **runmqakm** commands. These digital certificates have labels. A specific label associates a personal certificate with a queue manager or IBM WebSphere MQ MQI client. SSL and TLS use that certificate for authentication purposes. On UNIX, Linux, and Windows systems, IBM WebSphere MQ uses `ibmwebspheremq` as a label prefix to avoid confusion with certificates for other products. The prefix is followed by the name of the queue manager or IBM WebSphere MQ MQI client user logon ID, changed to lowercase. Ensure that you specify the entire certificate label in lowercase.

The key database file name comprises a path and stem name:

- On UNIX and Linux systems, the default path for a queue manager (set when you created the queue manager) is `/var/mqm/qmgrs/<queue_manager_name>/ssl`.

On Windows systems, the default path is *MQ_INSTALLATION_PATH\Qmgrs\queue_manager_name\ssl*, where *MQ_INSTALLATION_PATH* is the directory in which IBM WebSphere MQ is installed. For example, *C:\program files\IBM\WebSphere MQ\Qmgrs\QM1\ssl*.

The default stem name is *key*. Optionally, you can choose your own path and stem name, but the extension must be *.kdb*.

If you choose your own path or file name, set the permissions to the file to tightly control access to it.

- For a WebSphere MQ client, there is no default path or stem name. Tightly control access to this file. The extension must be *.kdb*.

Do not create key repositories on a file system that does not support file level locks, for example NFS version 2 on Linux systems.

See “Changing the key repository location for a queue manager on UNIX, Linux or Windows systems” on page 637 for information about checking and specifying the key database file name. You can specify the key database file name either before or after creating the key database file.

The user ID from which you run the **iKeyman** or **iKeycmd** commands must have write permission for the directory in which the key database file is created or updated. For a queue manager using the default *ssl* directory, the user ID from which you run **iKeyman** or **iKeycmd** must be a member of the *mqm* group. For a IBM WebSphere MQ MQI client, if you run **iKeyman** or **iKeycmd** from a user ID different from that under which the client runs, you must alter the file permissions to enable the IBM WebSphere MQ MQI client to access the key database file at run time. For more information, see “Accessing and securing your key database files on Windows” on page 635 or “Accessing and securing your key database files on UNIX and Linux systems” on page 635.

In **iKeyman** or **iKeycmd** version 7.0, new key databases are automatically populated with a set of pre-defined certificate authority (CA) certificates. In **iKeyman** or **iKeycmd** version 8.0, key databases are not automatically populated, making the initial setup more secure because you include only the CA certificates that you want, in your key database file.

Note: Because of this change in behavior for GSKit version 8.0 that results in CA certificates no longer being automatically added to the repository, you must manually add your preferred CA certificates. This change of behavior provides you with more granular control over the CA certificates used. See “Adding default CA certificates into an empty key repository, on UNIX, Linux or Windows systems with GSKit version 8.0” on page 635.

Using the **iKeyman** user interface

Note: If you must manage SSL or TLS certificates in a way that is FIPS-compliant, use the **runmqakm** command. The **iKeyman** user interface does not provide a FIPS-compliant option.

Complete the following steps to create a new CMS key database file for either a queue manager or a IBM WebSphere MQ MQI client:

1. On UNIX and Linux systems, log in as the root user. On Windows systems, log in as Administrator or as a member of the MQM group.
2. Start the **iKeyman** user interface by running the **strmqikm** command.
3. From the **Key Database File** menu, click **New**. The New window opens.
4. Click **Key database type** and select **CMS** (Certificate Management System).
5. In the **File Name** field, type a file name. This field already contains the text *key.kdb*. If your stem name is *key*, leave this field unchanged. If you have specified a different stem name, replace *key* with your stem name. However, you must not change the *.kdb* extension.
6. In the **Location** field, type the path. For example:

- For a queue manager: /var/mqm/qmgrs/QM1/ssl (on UNIX and Linux systems) or C:\Program Files\IBM\WebSphere MQ\qmgrs\QM1\ssl (on Windows systems).
The path must match the value of the **SSLKeyRepository** attribute of the queue manager.
 - For a IBM WebSphere MQ client: /var/mqm/ssl (on UNIX and Linux systems) or C:\mqm\ssl (on Windows systems).
7. Click **Open**. The Password Prompt window opens.
 8. Type a password in the **Password** field, and type it again in the **Confirm Password** field.
 9. Select the **Stash the password to a file** check box.

Note: If you do not stash the password, attempts to start SSL or TLS channels fail because they cannot obtain the password required to access the key database file.

10. Click **OK**.
11. Click **OK**. The Personal Certificates window opens.
12. Set the access permissions as described in “Accessing and securing your key database files on Windows” on page 635 or “Accessing and securing your key database files on UNIX and Linux systems” on page 635.

Using the command line

1. Use the following commands to create a new CMS key database file for either a queue manager or a IBM WebSphere MQ MQI client by using either the **iKeycmd** or **runmqakm** command:

- On UNIX, Linux, and Windows systems:
runmqckm -keydb -create -db *filename* -pw *password* -type *cms* -stash

- Using runmqakm:
runmqakm -keydb -create -db *filename* -pw *password* -type *cms*
-stash -fips -strong

where:

-db *filename*

Specifies the fully qualified file name of a CMS key database, and must have a file extension of .kdb.

-pw *password*

Specifies the password for the CMS key database.

-type *cms*

Specifies the type of database. (For IBM WebSphere MQ, it must be *cms*.)

-stash

Saves the key database password to a file.

-fips

Disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that are FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the **runmqakm** command fails.

-strong

Checks that the password entered satisfies the minimum requirements for password strength. The minimum requirements for a password are as follows:

- The password must be a minimum length of 14 characters.
- The password must contain a minimum of one lowercase character, one uppercase character, and one digit or special character. Special characters include the asterisk (*), the dollar sign (\$), the number sign (#), and the percent sign (%). A space is classified as a special character.
- Each character can occur a maximum of three times in a password.

- A maximum of two consecutive characters in the password can be identical.
- All characters are in the standard ASCII printable character set within the range 0x20 - 0x7E.

Accessing and securing your key database files on Windows:

The key database files might not have appropriate access permissions. You must set appropriate access to these files.

Set access control to the files *key.kdb*, *key.sth*, *key.crl*, and *key.rdb*, where *key* is the stem name of your key database, to grant authority to a restricted set of users.

Consider granting access as follows:

full authority

BUILTIN\Administrators, NT AUTHORITY\SYSTEM, and the user who created the database files.

read authority

For a queue manager, the local mqm group only. This assumes that the MCA is running under a user ID in the mqm group.

For a client, the user ID under which the client process is running.

Accessing and securing your key database files on UNIX and Linux systems:

The key database files might not have appropriate access permissions. You must set appropriate access to these files.

For a queue manager, set permissions on the key database files so that queue manager and channel processes can read them when necessary, but other users cannot read or modify them. Normally, the mqm user needs read permissions. If you have created the key database file by logging in as the mqm user, then the permissions are probably sufficient; if you were not the mqm user, but another user in the mqm group, you probably need to grant read permissions to other users in the mqm group.

Similarly for a client, set permissions on the key database files so that client application processes can read them when necessary, but other users cannot read or modify them. Normally, the user under which the client process runs needs read permissions. If you have created the key database file by logging in as that user, then the permissions are probably sufficient; if you were not the client process user, but another user in that group, you probably need to grant read permissions to other users in the group.

Set the permissions on the files *key.kdb*, *key.sth*, *key.crl*, and *key.rdb*, where *key* is the stem name of your key database, to read and write for the file owner, and to read for the mqm or client user group (-rw-r-----).

Adding default CA certificates into an empty key repository, on UNIX, Linux or Windows systems with GSKit version 8.0:

Follow this procedure to add one or more of the default CA certificates to an empty key repository with GSKit version 8.

In GSKit version 7.0, the behaviour when creating a new key repository was to automatically add in a set of default CA certificates for commonly-used Certificate Authorities. For GSKit version 8, this behaviour has changed so that CA certificates are no longer automatically added to the repository. The user is now required to manually add CA certificates into the key repository.

Using iKeyman

Perform the following steps on the machine on which you want to add the CA certificate:

1. Start the iKeyman GUI using the `strmqikm` command (on UNIX, Linux and Windows systems).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file to which you want to add the certificate, for example `key.kdb`.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file displays in the **File Name** field.
8. In the **Key database content** field, select **Signer Certificates**.
9. Click **Populate**. The Add CA's Certificate window opens.
10. The CA certificates that are available to be added to the repository are displayed in a hierarchical tree structure. Select the top level entry for the organization whose CA certificates you wish to trust to view the complete list of valid CA certificates.
11. Select the CA certificates you wish to trust from the list and click **OK**. The certificates are added to the key repository.

Using the command line

Use the following commands to list, then add CA certificates using `iKeycmd`:

- Issue the following command to list the default CA certificates along with the organizations which issue them:

```
runmqckm -cert -listsigners
```
- Issue the following command to add all of the CA certificates for the organization specified in the *label* field:

```
runmqckm -cert -populate -db filename -pw password -label label
```

where:

<code>-db <i>filename</i></code>	is the fully qualified path name of the key database.
<code>-pw <i>password</i></code>	is the password for the key database.
<code>-label <i>label</i></code>	is the label attached to the certificate.

Note: Adding a CA certificate to a key repository results in WebSphere MQ trusting all personal certificates signed by that CA certificate. Consider carefully which Certificate Authorities you wish to trust and only add the set of CA certificates needed to authenticate your clients and managers. It is not recommended to add the full set of default CA certificates unless this is a definitive requirement for your security policy.

Locating the key repository for a queue manager on UNIX, Linux or Windows systems:

Use this procedure to obtain the location of your queue manager's key database file

Procedure

1. Display your queue manager's attributes, using either of the following MQSC commands:

```
DISPLAY QMGR ALL  
DISPLAY QMGR SSLKEYR
```

You can also display your queue manager's attributes using the WebSphere MQ Explorer or PCF commands.

2. Examine the command output for the path and stem name of the key database file. For example,
 - a. on UNIX and Linux systems: `/var/mqm/qmgrs/QM1/ssl/key`, where `/var/mqm/qmgrs/QM1/ssl` is the path and `key` is the stem name
 - b. on Windows: `MQ_INSTALLATION_PATH\qmgrs\QM1\ssl\key`, where `MQ_INSTALLATION_PATH\qmgrs\QM1\ssl` is the path and `key` is the stem name. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Changing the key repository location for a queue manager on UNIX, Linux or Windows systems:

You can change the location of your queue manager's key database file by various means including the MQSC command `ALTER QMGR`.

You can change the location of your queue manager's key database file by using the MQSC command `ALTER QMGR` to set your queue manager's key repository attribute. For example, on UNIX and Linux systems:

```
ALTER QMGR SSLKEYR('/var/mqm/qmgrs/QM1/ssl/MyKey')
```

The key database file has the fully qualified file name: `/var/mqm/qmgrs/QM1/ssl/MyKey.kdb`

On Windows:

```
ALTER QMGR SSLKEYR('C:\Program Files\IBM\WebSphere MQ\Qmgrs\QM1\ssl\Mykey')
```

The key database file has the fully qualified file name: `C:\Program Files\IBM\WebSphere MQ\Qmgrs\QM1\ssl\Mykey.kdb`

Attention: Ensure that you do not include the `.kdb` extension in the file name on the `SSLKEYR` keyword, as the queue manager appends this extension automatically.

You can also alter your queue manager's attributes using the WebSphere MQ Explorer or PCF commands.

When you change the location of a queue manager's key database file, certificates are not transferred from the old location. If the key database file you are now accessing is a new key database file, you must populate it with the CA and personal certificates you need, as described in "Importing a personal certificate into a key repository on UNIX, Linux or Windows systems" on page 650.

Locating the key repository for a WebSphere MQ MQI client on UNIX, Linux and Windows systems.:

The location of the key repository is given by the `MQSSLKEYR` variable, or specified in the `MQCONN` call.

Examine the `MQSSLKEYR` environment variable to obtain the location of your WebSphere MQ MQI client's key database file. For example:

```
echo $MQSSLKEYR
```

Also check your application, because the key database file name can also be set in an `MQCONN` call, as described in "Specifying the key repository location for a WebSphere MQ MQI client on UNIX, Linux or Windows systems" on page 638. The value set in an `MQCONN` call overrides the value of `MQSSLKEYR`.

Specifying the key repository location for a WebSphere MQ MQI client on UNIX, Linux or Windows systems:

There is no default key repository for a WebSphere MQ MQI client. You can specify its location in either of two ways. Ensure that the key database file can be accessed only by intended users or administrators to prevent unauthorized copying to other systems.

You can specify the location of your WebSphere MQ MQI client's key database file in either of two ways:

- Setting the MQSSLKEYR environment variable. For example, on UNIX and Linux systems:
`export MQSSLKEYR=/var/mqm/ssl/key`

The key database file has the fully-qualified file name:

`/var/mqm/ssl/key.kdb`

On Windows:

`set MQSSLKEYR=C:\Program Files\IBM\WebSphere MQ\ssl\key`

The key database file has the fully-qualified file name:

`C:\Program Files\IBM\WebSphere MQ\ssl\key.kdb`

Note: The .kdb extension is a mandatory part of the file name, but is not included as part of the value of the environment variable.

- Providing the path and stem name of the key database file in the *KeyRepository* field of the MQSCO structure when an application makes an MQCONN call. For more information about using the MQSCO structure in MQCONN, see  Overview for MQSCO (*WebSphere MQ V7.1 Reference*).

When changes to certificates or the certificate store become effective on UNIX, Linux or Windows systems.:

When you change the certificates in a certificate store, or the location of the certificate store, the changes take effect depending on the type of channel and how the channel is running.

Changes to the certificates in the key database file and to the key repository attribute become effective in the following situations:

- When a new outbound single channel process first runs an SSL channel.
- When a new inbound TCP/IP single channel process first receives a request to start an SSL channel.
- When the MQSC command REFRESH SECURITY TYPE(SSL) is issued to refresh the Websphere MQ SSL environment.
- For client application processes, when the last SSL connection in the process is closed. The next SSL connection will pick up the certificate changes.
- For channels that run as threads of a process pooling process (amqrmppa), when the process pooling process is started or restarted and first runs an SSL channel. If the process pooling process has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).
- For channels that run as threads of the channel initiator, when the channel initiator is started or restarted and first runs an SSL channel. If the channel initiator process has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).
- For channels that run as threads of a TCP/IP listener, when the listener is started or restarted and first receives a request to start an SSL channel. If the listener has already run an SSL channel, and you want the change to become effective immediately, run the MQSC command REFRESH SECURITY TYPE(SSL).

You can also refresh the WebSphere MQ SSL environment using the WebSphere MQ Explorer or PCF commands.

Creating a self-signed personal certificate on UNIX, Linux, and Windows systems:

You can create a self-signed certificate by using iKeyman, iKeycmd, or runmqakm.

Note: WebSphere MQ does not support SHA-3 or SHA-5 algorithms. You can use the digital signature algorithm names SHA384WithRSA and SHA512WithRSA because both algorithms are members of the SHA-2 family.

The digital signature algorithm names SHA3WithRSA and SHA5WithRSA are deprecated because they are an abbreviated form of SHA384WithRSA and SHA512WithRSA respectively.

For more information about why you might want to use self-signed certificates, see “Using self-signed certificates for mutual authentication of two queue managers” on page 762.

Not all digital certificates can be used with all CipherSpecs. Ensure that you create a certificate that is compatible with the CipherSpecs you need to use. WebSphere MQ supports three different types of CipherSpec. For details, see “Interoperability of Elliptic Curve and RSA CipherSpecs” on page 405 in the “Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403 topic. To use the Type 1 CipherSpecs (those with names beginning ECDHE_ECDSA_) you must use the **runmqakm** command to create the certificate and you must specify an Elliptic Curve ECDSA signature algorithm parameter; for example, **-sig_alg** EC_ecdsa_with_SHA384.

Using iKeyman

iKeyman does not provide a FIPS-compliant option. If you need to manage SSL or TLS certificates in a way that is FIPS-compliant, use the **runmqakm** command.

Use the following procedure to obtain a self-signed certificate for your queue manager or WebSphere MQ MQI client:

1. Start the iKeyman GUI by using the **strmqikm** command .
2. From the **Key Database File** menu, click **Open**. The Open window displays.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file in which you want to save the certificate, for example key.kdb.
6. Click **Open**. The Password Prompt window displays.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field.
8. From the **Create** menu, click **New Self-Signed Certificate**. The Create New Self-Signed Certificate window is displayed.
9. In the **Key Label** field, type:
 - For a queue manager, `ibmwebspheremq` followed by the name of your queue manager folded to lowercase. For example, for QM1, `ibmwebspheremqqm1`, or,
 - For a WebSphere MQ client, `ibmwebspheremq` followed by your logon user ID folded to lowercase, for example `ibmwebspheremqmyuserid`.
10. Type or select a value for any field in the **Distinguished name**, or any of the **Subject alternative name** fields.
11. For the remaining fields, either accept the default values, or type or select new values. For more information about Distinguished Names, see “Distinguished Names” on page 373.
12. Click **OK**. The **Personal Certificates** list shows the label of the self-signed personal certificate you created.

Using the command line

Use the following commands to create a self-signed personal certificate by using iKeycmd or runmqkm:

- Using iKeycmd on UNIX, Linux and Windows systems:

```
runmqckm -cert -create -db filename -pw password -label label
          -dn distinguished_name -size key_size -x509version version
          -expire days -sig_alg algorithm
```

Instead of `-dn distinguished_name`, you can use `-san_dnsname DNS_names`, `-san_emailaddr email_addresses`, or `-san_ipaddr IP_addresses`.

- Using runmqakm:

```
runmqakm -cert -create -db filename -pw password -label label
          -dn distinguished_name -size key_size -x509version version -expire days
          -fips -sig_alg algorithm
```

<code>-db filename</code>	The fully qualified file name of a CMS key database.
<code>-pw password</code>	The password for the CMS key database.
<code>-label label</code>	The key label attached to the certificate.
<code>-dn distinguished_name</code>	The X.500 distinguished name enclosed in double quotation marks. At least one attribute is required. You can supply multiple OU or DC attributes.
<code>-size key_size</code>	The key size. For iKeycmd, the value can be 512 or 1024. For runmqakm, the value can be 512, 1024, 2048 or 4096.
<code>-x509version version</code>	The version of X.509 certificate to create. The value can be 1, 2, or 3. The default is 3.
<code>-expire days</code>	The expiration time in days of the certificate. The default is 365 days for a certificate.
<code>-fips</code>	Specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the runmqakm command fails.
<code>-sig_alg</code>	For runmqakm, the hashing algorithm used during the creation of a self-signed certificate. This hashing algorithm is used to create the signature associated with the newly created self-signed certificate. The value can be md5, MD5_WITH_RSA, MD5WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, sha1, SHA1WithDSA, SHA1WithECDSA, SHA1WithRSA, sha224, SHA224_WITH_RSA, SHA224WithDSA, SHA224WithECDSA, SHA224WithRSA, sha256, SHA256_WITH_RSA, SHA256WithDSA, SHA256WithECDSA, SHA256WithRSA, SHA2WithRSA, sha384, SHA384_WITH_RSA, SHA384WithECDSA, SHA384WithRSA, sha512, SHA512_WITH_RSA, SHA512WithECDSA, SHA512WithRSA, SHAWithDSA, SHAWithRSA, EC_ecdsa_with_SHA1, EC_ecdsa_with_SHA224, EC_ecdsa_with_SHA256, EC_ecdsa_with_SHA384, or EC_ecdsa_with_SHA512. The default value is SHA1WithRSA.
<code>-sig_alg</code>	For iKeycmd, the asymmetric signature algorithm used for the creation of the entry's key pair. The value can be MD2_WITH_RSA, MD2WithRSA, MD5_WITH_RSA, MD5WithRSA, SHA1WithDSA, SHA1WithRSA, SHA256_WITH_RSA, SHA256WithRSA, SHA2WithRSA, SHA384_WITH_RSA, SHA384WithRSA, SHA512_WITH_RSA, SHA512WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, SHAWithDSA, or SHAWithRSA. The default value is SHA1WithRSA.
<code>-san_dnsname DNS_names</code>	A comma- or space-delimited list of DNS names for the entry being created.
<code>-san_emailaddr email_addresses</code>	A comma- or space-delimited list of email addresses for the entry being created.
<code>-san_ipaddr IP_addresses</code>	A comma- or space-delimited list of IP addresses for the entry being created.

Requesting a personal certificate on UNIX, Linux, and Windows systems:

You can request a personal certificate by using the iKeyman user interface, or by using the **iKeycmd** or **runmqakm** commands.

About this task

Note: WebSphere MQ does not support SHA-3 or SHA-5 algorithms. You can use the digital signature algorithm names SHA384WithRSA and SHA512WithRSA because both algorithms are members of the SHA-2 family.

The digital signature algorithm names SHA3WithRSA and SHA5WithRSA are deprecated because they are an abbreviated form of SHA384WithRSA and SHA512WithRSA respectively.

Not all digital certificates can be used with all CipherSpecs. Ensure that you request a certificate that is compatible with the CipherSpecs you need to use. WebSphere MQ supports three different types of CipherSpec. For details, see “Interoperability of Elliptic Curve and RSA CipherSpecs” on page 405 in the “Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403 topic. To use the Type 1 CipherSpecs (with names beginning EC_{DHE}_ECDSA_) you must use the **runmqakm** command to request the certificate and you must specify an Elliptic Curve ECDSA signature algorithm parameter; for example, **-sig_alg** EC_ecdsa_with_SHA384.

Using the iKeyman user interface:

About this task

iKeyman does not provide a FIPS-compliant option. If you need to manage SSL or TLS certificates in a way that is FIPS-compliant, use the **runmqakm** command.

Procedure

Complete the following steps to apply for a personal certificate, by using the iKeyman user interface:

1. Start the iKeyman user interface by using the **strmqikm** command on UNIX, Linux, and Windows systems.
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to generate the request; for example, key.kdb.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is shown in the **File Name** field.
8. From the **Create** menu, click **New Certificate Request**. The Create New Key and Certificate Request window opens.
9. In the **Key Label** field, enter the following labels:
 - For a queue manager, enter `ibmwebsphermq` followed by the name of your queue manager changed to lowercase. For example, for a queue manager called QM1, enter `ibmwebsphermqmq1`.
 - For a IBM WebSphere MQ MQI client, enter `ibmwebsphermq` followed by your logon user ID, all in lowercase; for example, `ibmwebsphermqmyuserid`.
10. Type or select a value for any field in the **Distinguished name** field, or any of the **Subject alternative name** fields. For the remaining fields, either accept the default values, or type or select new values. For more information about Distinguished Names, see “Distinguished Names” on page 373.
11. In the **Enter the name of a file in which to store the certificate request** field, either accept the default `certreq.arm`, or type a new value with a full path.

12. Click **OK**. A confirmation window is displayed.
13. Click **OK**. The **Personal Certificate Requests** list shows the label of the new personal certificate request you created. The certificate request is stored in the file you chose in step 11 on page 641.
14. Request the new personal certificate either by sending the file to a certificate authority (CA), or by copying the file into the request form on the website for the CA.

Using the command line:

Procedure

Use the following commands to request a personal certificate by using either the **iKeycmd** or **runmqakm** command:

- Using **iKeycmd** on UNIX, Linux, and Windows systems:

```
runmqckm -certreq -create -db filename -pw password -label label
        -dn distinguished_name -size key_size -file filename -sig_alg algorithm
```

Instead of `-dn distinguished_name`, you can use `-san_dsname DNS_names`, `-san_emailaddr email_addresses`, or `-san_ipaddr IP_addresses`.

- Using **runmqakm**:

```
runmqakm -certreq -create -db filename -pw password -label label
        -dn distinguished_name -size key_size -file filename -fips
        -sig_alg algorithm
```

where:

-db filename

Specifies the fully qualified file name of a CMS key database.

-pw password

Specifies the password for the CMS key database.

-label label

Specifies the key label attached to the certificate.

-dn distinguished_name

Specifies the X.500 distinguished name enclosed in double quotation marks. At least one attribute is required. You can supply multiple OU and DC attributes.

-size key_size

Specifies the key size. If you are using **iKeycmd**, the value can be 512 or 1024. If you are using **runmqakm**, the value can be 512, 1024, or 2048.

-file filename

Specifies the file name for the certificate request.

-fips

Specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that are FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the **runmqakm** command fails.

-sig_alg

For **iKeycmd**, specifies the asymmetric signature algorithm used for the creation of the entry's key pair. The value can be MD2_WITH_RSA, MD2WithRSA, MD5_WITH_RSA, MD5WithRSA, SHA1WithDSA, SHA1WithRSA, SHA256_WITH_RSA, SHA256WithRSA, SHA2WithRSA, SHA384_WITH_RSA, SHA384WithRSA, SHA512_WITH_RSA, SHA512WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, SHAWithDSA, or SHAWithRSA. The default value is SHA1WithRSA

-sig_alg

For **runmqakm**, specifies the hashing algorithm used during the creation of a certificate request. This hashing algorithm is used to create the signature associated with the newly created certificate request.

The value can be md5, MD5_WITH_RSA, MD5WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, sha1, SHA1WithDSA, SHA1WithECDSA, SHA1WithRSA, sha224, SHA224_WITH_RSA, SHA224WithDSA, SHA224WithECDSA, SHA224WithRSA, sha256, SHA256_WITH_RSA, SHA256WithDSA, SHA256WithECDSA, SHA256WithRSA, SHA2WithRSA, sha384, SHA384WithECDSA, SHA384WithRSA, sha512, SHA512_WITH_RSA, SHA512WithECDSA, SHA512WithRSA, SHAWithDSA, SHAWithRSA, EC_ecdsa_with_SHA1, EC_ecdsa_with_SHA224, EC_ecdsa_with_SHA256, EC_ecdsa_with_SHA384, or EC_ecdsa_with_SHA512. The default value is SHA1WithRSA.

-san_dnsname *DNS_names*

Specifies a comma-delimited or space-delimited list of DNS names for the entry being created.

-san_emailaddr *email_addresses*

Specifies a comma-delimited or space-delimited list of email addresses for the entry being created.

-san_ipaddr *IP_addresses*

Specifies a comma-delimited or space-delimited list of IP addresses for the entry being created.

If you are using cryptographic hardware, see “Requesting a personal certificate for your PKCS #11 hardware” on page 657.

Renewing an existing personal certificate on UNIX, Linux, and Windows systems:

You can renew a personal certificate by using the iKeyman user interface, or by using the **iKeycmd** or **runmqkm** commands.

Before you begin

If you have a requirement to use larger key sizes for your personal certificates, the renewal steps described below do not work, because the recreated certificate request is generated from an existing key.

Follow the steps described in “Requesting a personal certificate on UNIX, Linux, and Windows systems” on page 641 to create a new certificate request, using the key sizes you require. This process replaces your existing key.

About this task

A personal certificate has an expiry date, after which the certificate can no longer be used. This task explains how to renew an existing personal certificate before it expires.

Using the iKeyman user interface:

About this task

iKeyman does not provide a FIPS-compliant option. If you need to manage SSL or TLS certificates in a way that is FIPS-compliant, use the **runmqkm** command.

Procedure

Complete the following steps to apply for a personal certificate, by using the iKeyman user interface:

1. Start the iKeyman user interface by using the **strmqikm** command on UNIX, Linux, and Windows systems.
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to generate the request; for example, key.kdb.
6. Click **Open**. The Password Prompt window opens.

7. Type the password you set when you created the key database and click **OK**. The name of your key database file is shown in the **File Name** field.
8. Select **Personal Certificates** from the drop down selection menu, and select the certificate from the list that you want to renew.
9. Click the **Recreate Request...** button. A window opens for you to enter the file name and file location information.
10. In the **file name** field, either accept the default `certreq.arm`, or type a new value, including the full file path.
11. Click **OK**. The certificate request is stored in the file you selected in step 9.
12. Request the new personal certificate either by sending the file to a certificate authority (CA), or by copying the file into the request form on the website for the CA.

Using the command line:

Procedure

Use the following commands to request a personal certificate by using either the **iKeycmd** or **runmqakm** command:

- Using **iKeycmd** on UNIX, Linux, and Windows systems:
`runmqckm -certreq -recreate -db filename -pw password -label label
-target filename`
- Using **runmqakm**:
`runmqakm -certreq -recreate -db filename -pw password -label label
-target filename`

where:

- db filename**
Specifies the fully qualified file name of a CMS key database.
- pw password**
Specifies the password for the CMS key database.
- target filename**
Specifies the file name for the certificate request.

What to do next

Once you have received the signed personal certificate from the certificate authority, you can add it to your key database using the steps described in "Receiving personal certificates into a key repository on UNIX, Linux and Windows systems."

Receiving personal certificates into a key repository on UNIX, Linux and Windows systems:

Use this procedure to receive a personal certificate into the key database file. The key repository must be the same repository where you created the certificate request.

After the CA sends you a new personal certificate, you add it to the key database file from which you generated the new certificate request. If the CA sends the certificate as part of an email message, copy the certificate into a separate file.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the `runmqakm` command. **iKeyman** does not provide a FIPS-compliant option.

Ensure that the certificate file to be imported has write permission for the current user, and then use the following procedure for either a queue manager or a WebSphere MQ MQI client to receive a personal certificate into the key database file:

1. Start the iKeyman GUI using the **strmqikm** command (on Windows UNIX and Linux).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file to which you want to add the certificate, for example key.kdb.
6. Click **Open**, and then click **OK**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field. Select the **Personal Certificates** view.
8. Click **Receive**. The Receive Certificate from a File window opens.
9. Type the certificate file name and location for the new personal certificate, or click **Browse** to select the name and location.
10. Click **OK**. If you already have a personal certificate in your key database, a window opens, asking if you want to set the key you are adding as the default key in the database.
11. Click **Yes** or **No**. The Enter a Label window opens.
12. Click **OK**. The **Personal Certificates** field shows the label of the new personal certificate you added.

Using the command line

Use the following commands to add a personal certificate to a key database file using iKeycmd or runmqakm.:

- On UNIX, Linux and Windows, issue the following command:

```
runmqckm -cert -receive -file filename -db filename -pw password
        -format ascii
```

- On UNIX, Linux or Windows systems, issue the following command:

```
runmqakm -cert -receive -file filename -db filename -pw password -fips
```

where:

<i>-file filename</i>	is the fully qualified file name of the file containing the personal certificate.
<i>-db filename</i>	is the fully qualified file name of a CMS key database.
<i>-pw password</i>	is the password for the CMS key database.
<i>-format ascii</i>	is the format of the certificate. The value can be <i>ascii</i> for Base64-encoded ASCII or <i>binary</i> for Binary DER data. The default is <i>ascii</i> .
<i>-fips</i>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the runmqakm command fails.

If you are using cryptographic hardware, refer to “Importing a personal certificate to your PKCS #11 hardware” on page 659.

Extracting a CA certificate from a key repository:

Follow this procedure to extract a CA certificate.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the `runmqakm` command. iKeyman does not provide a FIPS-compliant option.

Perform the following steps on the machine from which you want to extract the CA certificate:

1. Start the iKeyman GUI using the `strmqikm` command..
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to extract, for example `key.kdb`.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field.
8. In the **Key database content** field, select **Signer Certificates** and select the certificate you want to extract.
9. Click **Extract**. The Extract a Certificate to a File window opens.
10. Select the **Data type** of the certificate, for example **Base64-encoded ASCII data** for a file with the `.arm` extension.
11. Type the certificate file name and location where you want to store the certificate, or click **Browse** to select the name and location.
12. Click **OK**. The certificate is written to the file you specified.

Using the command line

Use the following commands to extract a CA certificate using `iKeycmd` or `runmqakm`:

- On UNIX, Linux and Windows:

```
runmqckm -cert -extract -db filename -pw password -label label -target filename  
-format ascii
```

- Using `runmqakm`:

```
runmqakm -cert -extract -db filename -pw password -label label  
-target filename -format ascii -fips
```

where:

<code>-db <i>filename</i></code>	is the fully qualified path name of a CMS key database.
<code>-pw <i>password</i></code>	is the password for the CMS key database.
<code>-label <i>label</i></code>	is the label attached to the certificate.
<code>-target <i>filename</i></code>	is the name of the destination file.
<code>-format <i>ascii</i></code>	is the format of the certificate. The value can be <code>ascii</code> for Base64-encoded ASCII or <code>binary</code> for Binary DER data. The default is <code>ascii</code> .
<code>-fips</code>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the <code>runmqakm</code> command fails.

Extracting the public part of a self-signed certificate from a key repository on UNIX, Linux and Windows systems:

Follow this procedure to extract the public part of a self-signed certificate.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the `runmqakm` command. iKeyman does not provide a FIPS-compliant option.

Perform the following steps on the machine from which you want to extract the public part of a self-signed certificate:

1. Start the iKeyman GUI using the `strmqikm` command (on UNIX, Linux and Windows).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to extract the certificate, for example `key.kdb`.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field.
8. In the **Key database content** field, select **Personal Certificates** and select the certificate.
9. Click **Extract certificate**. The Extract a Certificate to a File window opens.
10. Select the **Data type** of the certificate, for example **Base64-encoded ASCII data** for a file with the `.arm` extension.
11. Type the certificate file name and location where you want to store the certificate, or click **Browse** to select the name and location.
12. Click **OK**. The certificate is written to the file you specified. Note that when you extract (rather than export) a certificate, only the public part of the certificate is included, so a password is not required.

Using the command line

Use the following commands to extract the public part of a self-signed certificate using `iKeycmd` or `runmqakm`:

- On UNIX, Linux and Windows:

```
runmqckm -cert -extract -db filename -pw password -label label -target filename
        -format ascii
```
- Using `runmqakm`:

```
runmqakm -cert -extract -db filename -pw password -label label
        -target filename -format ascii -fips
```

where:

<code>-db <i>filename</i></code>	is the fully qualified path name of a CMS key database.
<code>-pw <i>password</i></code>	is the password for the CMS key database.
<code>-label <i>label</i></code>	is the label attached to the certificate.
<code>-target <i>filename</i></code>	is the name of the destination file.
<code>-format <i>ascii</i></code>	is the format of the certificate. The value can be <code>ascii</code> for Base64-encoded ASCII or <code>binary</code> for Binary DER data. The default is <code>ascii</code> .
<code>-fips</code>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the <code>runmqakm</code> command fails.

Adding a CA certificate (or the public part of a self-signed certificate) into a key repository, on UNIX, Linux or Windows systems:

Follow this procedure to add a CA certificate or the public part of a self-signed certificate to the key repository.

If the certificate that you want to add is in a certificate chain, you must also add all the certificates that are above it in the chain. You must add the certificates in strictly descending order starting from the root, followed by the CA certificate immediately below it in the chain, and so on.

Where the following instructions refer to a CA certificate, they also apply to the public part of a self-signed certificate.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the `runmqakm` command. iKeyman does not provide a FIPS-compliant option.

Perform the following steps on the machine on which you want to add the CA certificate:

1. Start the iKeyman GUI using the `strmqikm` command (on UNIX, Linux and Windows systems).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file to which you want to add the certificate, for example `key.kdb`.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file displays in the **File Name** field.
8. In the **Key database content** field, select **Signer Certificates**.
9. Click **Add**. The Add CA's Certificate from a File window opens.
10. Type the certificate file name and location where the certificate is stored, or click **Browse** to select the name and location.
11. Click **OK**. The Enter a Label window opens.
12. In the Enter a Label window, type the name of the certificate.
13. Click **OK**. The certificate is added to the key database.

Using the command line

Use the following commands to add a CA certificate using `iKeycmd` or `runmqakm`:

- On UNIX, Linux and Windows, issue the following command:

```
runmqckm -cert -add -db filename -pw password -label label -file filename  
-format ascii
```
- On UNIX, Linux or Windows systems, issue the following command:

```
runmqakm -cert -add -db filename -pw password -label label -file filename  
-format ascii -fips
```

where:

<code>-db filename</code>	is the fully qualified path name of the CMS key database.
<code>-pw password</code>	is the password for the CMS key database.
<code>-label label</code>	is the label attached to the certificate.
<code>-file filename</code>	is the name of the file containing the certificate.
<code>-format ascii</code>	is the format of the certificate. The value can be <code>ascii</code> for Base64-encoded ASCII or <code>binary</code> for Binary DER data. The default is <code>ascii</code> .
<code>-fips</code>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the <code>runmqakm</code> command fails.

Exporting a personal certificate from a key repository:

Follow this procedure to exporting a personal certificate.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the `runmqakm` command. iKeyman does not provide a FIPS-compliant option.

Perform the following steps on the machine from which you want to export the personal certificate:

1. Start the iKeyman GUI using the `strmqikm` command (on Windows UNIX and Linux).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to export the certificate, for example `key.kdb`.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field.
8. In the **Key database content** field, select **Personal Certificates** and select the certificate you want to export.
9. Click **Export/Import**. The Export/Import key window opens.
10. Select **Export Key**.
11. Select the **Key file type** of the certificate you want to export, for example **PKCS12**.
12. Type the file name and location to which you want to export the certificate, or click **Browse** to select the name and location.
13. Click **OK**. The Password Prompt window opens. Note that when you export (rather than extract) a certificate, both the public and private parts of the certificate are included. This is why the exported file is protected by a password. When you extract a certificate, only the public part of the certificate is included, so a password is not required.
14. Type a password in the **Password** field, and type it again in the **Confirm Password** field.
15. Click **OK**. The certificate is exported to the file you specified.

Using the command line

Use the following commands to export a personal certificate using iKeycmd:

- On UNIX, Linux and Windows:


```
runmqckm -cert -export -db filename -pw password -label label -type cms
          -target filename -target_pw password -target_type pkcs12
```

To export a personal certificate using runmqakm, use the following command:

```
runmqakm -cert -export -db filename -pw password -label label -type cms  
-target filename -target_pw password -target_type pkcs12  
-encryption strong | weak -fips
```

where:

-db <i>filename</i>	is the fully qualified path name of the CMS key database.
-encryption	is the strength of encryption used in certificate export command. The value can be strong or weak. The default is strong.
-fips	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the runmqakm command fails.
-pw <i>password</i>	is the password for the CMS key database.
-label <i>label</i>	is the label attached to the certificate.
-type <i>cms</i>	is the type of the database.
-target <i>filename</i>	is the fully qualified path name of the destination file.
-target_pw <i>password</i>	is the password for encrypting the certificate.
-target_type <i>pkcs12</i>	is the type of the certificate.

Importing a personal certificate into a key repository on UNIX, Linux or Windows systems:

Follow this procedure to import a personal certificate

Before importing a personal certificate in PKCS #12 format into the key database file, you must first add the full valid chain of issuing CA certificates to the key database file (see “Adding a CA certificate (or the public part of a self-signed certificate) into a key repository, on UNIX, Linux or Windows systems” on page 648).

PKCS #12 files should be considered temporary and deleted after use.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS-compliant, use the runmqakm command. iKeyman does not provide a FIPS-compliant option.

Perform the following steps on the machine to which you want to import the personal certificate:

1. Start the iKeyman GUI using the **strmqikm** command .
2. From the **Key Database File** menu, click **Open**. The Open window displays.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file to which you want to add the certificate, for example key.kdb.
6. Click **Open**. The Password Prompt window displays.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file displays in the **File Name** field.
8. In the **Key database content** field, select **Personal Certificates**.
9. If there are certificates in the Personal Certificates view, follow these steps:
 - a. Click **Export/Import**. The Export/Import key window is displayed.
 - b. Select **Import Key**.
10. If there are no certificates in the Personal Certificates view, click **Import**.
11. Select the **Key file type** of the certificate you want to import, for example PKCS12.

12. Type the certificate file name and location where the certificate is stored, or click **Browse** to select the name and location.
13. Click **OK**. The Password Prompt window displays.
14. In the **Password** field, type the password used when the certificate was exported.
15. Click **OK**. The Change Labels window is displayed. This window allows the labels of certificates being imported to be changed if, for example, a certificate with the same label already exists in the target key database. Changing certificate labels has no effect on certificate chain validation. This can be used to change the personal certificate label to that required by WebSphere MQ in order to associate the certificate with the particular queue manager or client (*ibmwebsphermqmqm1* for example).
16. To change a label, select the required label from the **Select a label to change** list. The label is copied into the **Enter a new label** entry field. Replace the label text with that of the new label and click **Apply**.
17. The text in the **Enter a new label** entry field is copied back into the **Select a label to change** field, replacing the originally selected label and so relabelling the corresponding certificate.
18. When you have changed all the labels that needed to be changed, click **OK**. The Change Labels window closes, and the original IBM Key Management window reappears with the **Personal Certificates** and **Signer Certificates** fields updated with the correctly labeled certificates.
19. The certificate is imported to the target key database.

Using the command line

To import a personal certificate using `iKeycmd`, use the following commands:

- On UNIX, Linux and Windows:

```
runmqckm -cert -import -file filename -pw password -type pkcs12 -target filename
-target_pw password -target_type cms -label label
```

To import a personal certificate using `runmqakm`, use the following command:

```
runmqakm -cert -import -file filename -pw password -type pkcs12 -target filename
-target_pw password -target_type cms -label label -fips
```

where:

<code>-file filename</code>	is the fully qualified file name of the file containing the PKCS #12 certificate.
<code>-pw password</code>	is the password for the PKCS #12 certificate.
<code>-type pkcs12</code>	is the type of the file.
<code>-target filename</code>	is the name of the destination CMS key database.
<code>-target_pw password</code>	is the password for the CMS key database.
<code>-target_type cms</code>	is the type of the database specified by <code>-target</code>
<code>-label label</code>	is the label of the certificate to import from the source key database.
<code>-new_label label</code>	is the label that the certificate will be assigned in the target database. If you omit <code>-new_label</code> option, the default is to use the same as the <code>-label</code> option.
<code>-fips</code>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the <code>runmqakm</code> command fails.

`iKeycmd` does not provide a command to change certificate labels directly. Use the following steps to change a certificate label:

1. Export the certificate to a PKCS #12 file using the **-cert -export** command. Specify the existing certificate label for the `-label` option.
2. Remove the existing copy of the certificate from the original key database using the **-cert -delete** command.

3. Import the certificate from the PKCS #12 file using the **-cert -import** command. Specify the old label for the **-label** option and the required new label for the **-new_label** option. The certificate will be imported back into the key database with the required label.

Importing from a Microsoft .pfx file:

Follow this procedure to import from a Microsoft .pfx file using iKeyman. You cannot use runmqakm to import a .pfx file.

A .pfx file can contain two certificates relating to the same key. One is a personal or site certificate (containing both a public and private key). The other is a CA (signer) certificate (containing only a public key). These certificates cannot coexist in the same CMS key database file, so only one of them can be imported. Also, the “friendly name” or label is attached to only the signer certificate.

The personal certificate is identified by a system generated Unique User Identifier (UUID). This section shows the import of a personal certificate from a pfx file while labeling it with the friendly name previously assigned to the CA (signer) certificate. The issuing CA (signer) certificates should already be added to the target key database. Note that PKCS#12 files should be considered temporary and deleted after use.

Follow these steps to import a personal certificate from a source pfx key database:

1. Start the iKeyman GUI using the **strmqikm** command (on Linux, UNIX or Windows). The IBM Key Management window is displayed.
2. From the **Key Database File** menu, click **Open**. The Open window is displayed.
3. Select a key database type of **PKCS12**.
4. **You are recommended to take a backup of the pfx database before performing this step.** Select the pfx key database that you want to import. Click **Open**. The Password Prompt window is displayed.
5. Enter the key database password and click **OK**. The IBM Key Management window is displayed. The title bar shows the name of the selected pfx key database file, indicating that the file is open and ready.
6. Select **Signer Certificates** from the list. The “friendly name” of the required certificate is displayed as a label in the Signer Certificates panel.
7. Select the label entry and click **Delete** to remove the signer certificate. The Confirm window is displayed.
8. Click **Yes**. The selected label is no longer displayed in the Signer Certificates panel.
9. Repeat steps 6, 7, and 8 for all the signer certificates.
10. From the **Key Database File** menu, click **Open**. The Open window is displayed.
11. Select the target key CMS database which the pfx file is being imported into. Click **Open**. The Password Prompt window is displayed.
12. Enter the key database password and click **OK**. The IBM Key Management window is displayed. The title bar shows the name of the selected key database file, indicating that the file is open and ready.
13. Select **Personal Certificates** from the list.
14. If there are certificates in the Personal Certificates view, follow these steps:
 - a. Click **Export/Import key**. The Export/Import key window is displayed.
 - b. Select **Import** from Choose Action Type.
15. If there are no certificates in the Personal Certificates view, click **Import**.
16. Select the PKCS12 file.
17. Enter the name of the pfx file as used in Step 4. Click **OK**. The Password Prompt window is displayed.
18. Specify the same password that you specified when you deleted the signer certificate. Click **OK**.

19. The Change Labels window is displayed (as there should be only a single certificate available for import). The label of the certificate should be a UUID which has a format xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.
20. To change the label select the UUID from the **Select a label to change:** panel. The label will be replicated into the **Enter a new label:** field. Replace the label text with that of the friendly name that was deleted in Step 7 and click **Apply**. The friendly name must be in the form `ibmwebspheremq`, followed by the queue manager name or the WebSphere MQ MQI client user logon ID in lowercase.
21. Click **OK**. The Change Labels window is now removed and the original IBM Key Management window reappears with the Personal Certificates and Signer Certificates panels updated with the correctly labeled personal certificate.
22. The pfx personal certificate is now imported to the (target) database.

It is not possible to change a certificate label using `iKeycmd` or `runmqakm`.

Importing from a PKCS #7 file:

The `iKeyman` and `iKeycmd` tools do not support PKCS #7 (.p7b) files. Use the `runmqckm` tool to import certificates from a PKCS #7 file.

Use the following command to add a CA certificate from a PKCS #7 file:

```
runmqckm -cert -add -db filename -pw password -type cms -file filename
-label label
```

<code>-db <i>filename</i></code>	is the fully qualified file name of the CMS key database.
<code>-pw <i>password</i></code>	is the password for the key database.
<code>-type <i>cms</i></code>	is the type of the key database.
<code>-file <i>filename</i></code>	is the name of the PKCS #7 file.
<code>-label <i>label</i></code>	is the label that the certificate is assigned in the target database. The first certificate takes the label given. All other certificates, if present, are labeled with their subject name.

Use the following command to import a personal certificate from a PKCS #7 file:

```
runmqckm -cert -import -db filename -pw password -type pkcs7 -target filename
-target_pw password -target_type cms -label label -new_label label
```

<code>-db <i>filename</i></code>	is the fully qualified file name of the file containing the PKCS #7 certificate.
<code>-pw <i>password</i></code>	is the password for the PKCS #7 certificate.
<code>-type <i>pkcs7</i></code>	is the type of the file.
<code>-target <i>filename</i></code>	is the name of the destination key database.
<code>-target_pw <i>password</i></code>	is the password for the destination key database.
<code>-target_type <i>cms</i></code>	is the type of the database specified by <code>-target</code>
<code>-label <i>label</i></code>	is the label of the certificate that is to be imported.
<code>-new_label <i>label</i></code>	is the label that the certificate will be assigned in the target database. If you omit the <code>-new_label</code> option, the default is to use the same as the <code>-label</code> option.

Generating strong passwords for key repository protection:

You can generate strong passwords for key repository protection using the **runmqakm** command.

You can use the **runmqakm** command with the following parameters to generate a strong password:

```
runmqakm -random -create -length 14 -strong -fips
```

When using the generated password on the **-pw** parameter of subsequent certificate administration commands, always place double quotation marks around the password. On UNIX and Linux systems, you must also use a backslash character to escape the following characters if they appear in the password string:

```
! \ " ' & * ( )
```

When entering the password in response to a prompt from **runmqckm**, **runmqakm** or the iKeyman GUI then it is not necessary to quote or escape the password. It is not necessary because the operating system shell does not affect data entry in these cases.

Deleting a certificate from a key repository on UNIX, Linux or Windows systems:

Use this procedure to remove personal or CA certificates.

Using iKeyman

If you need to manage SSL certificates in a way that is FIPS compliant, use the **runmqakm** command. iKeyman does not provide a FIPS-compliant option.

1. Start the iKeyman GUI using the **strmqikm** command (on UNIX, Linux and Windows systems).
2. From the **Key Database File** menu, click **Open**. The Open window opens.
3. Click **Key database type** and select **CMS** (Certificate Management System).
4. Click **Browse** to navigate to the directory that contains the key database files.
5. Select the key database file from which you want to delete the certificate, for example **key.kdb**.
6. Click **Open**. The Password Prompt window opens.
7. Type the password you set when you created the key database and click **OK**. The name of your key database file is displayed in the **File Name** field.
8. From the drop down list, select **Personal Certificates** or **Signer Certificates**
9. Select the certificate you want to delete.
10. If you do not already have a copy of the certificate and you want to save it, click **Export/Import** and export it (see "Exporting a personal certificate from a key repository" on page 649).
11. With the certificate selected, click **Delete**. The Confirm window opens.
12. Click **Yes**. The **Personal Certificates** field no longer shows the label of the certificate you deleted.

Using the command line

Use the following commands to delete a certificate using iKeycmd or **runmqakm**:

- On UNIX, Linux and Windows:

```
runmqckm -cert -delete -db filename -pw password -label label
```
- Using **runmqakm**:

```
runmqakm -cert -delete -db filename -pw password -label label -fips
```

where:

<code>-db filename</code>	is the fully qualified file name of a CMS key database.
<code>-pw password</code>	is the password for the CMS key database.
<code>-label label</code>	is the label attached to the personal certificate.
<code>-fips</code>	specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that have been FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the <code>runmqakm</code> command fails.

Configuring for cryptographic hardware on UNIX, Linux or Windows systems:

You can configure cryptographic hardware for a queue manager or client in a number of ways.

You can configure cryptographic hardware for a queue manager on UNIX, Linux or Windows systems using either of the following methods:

- Use the ALTER QMGR MQSC command with the SSLCRYP parameter, as described in  ALTER QMGR (*WebSphere MQ V7.1 Reference*).
- Use WebSphere MQ Explorer to configure the cryptographic hardware on your UNIX, Linux or Windows system. For more information, refer to the online help.

You can configure cryptographic hardware for a WebSphere MQ client on UNIX, Linux or Windows systems using either of the following methods:

- Set the MQSSLCRYP environment variable. The permitted values for MQSSLCRYP are the same as for the SSLCRYP parameter, as described in  ALTER QMGR (*WebSphere MQ V7.1 Reference*). If you use the GSK_PCS11 version of the SSLCRYP parameter, the PKCS #11 token label must be specified entirely in lower-case.
- Set the **CryptoHardware** field of the SSL configuration options structure, MQSCO, on an MQCONN call. For more information, see  Overview for MQSCO (*WebSphere MQ V7.1 Reference*).

If you have configured cryptographic hardware which uses the PKCS #11 interface using any of these methods, you must store the personal certificate for use on your channels in the key database file for the cryptographic token you have configured. This is described in “Managing certificates on PKCS #11 hardware.”

Managing certificates on PKCS #11 hardware:

You can manage digital certificates on cryptographic hardware that supports the PKCS #11 interface.

You must create a key database to prepare the IBM WebSphere MQ environment, even if you do not intend to store certificate authority (CA) certificates in it, but will store all your certificates on your cryptographic hardware. A key database is necessary for the queue manager to reference in its SSLKEYR field, or for the client application to reference in the MQSSLKEYR environment variable. This key database is also required if you are creating a certificate request.

Using the iKeyman user interface

1. On UNIX and Linux systems, log in as the root user. On Windows systems, log in as Administrator or as a member of the MQM group.
2. Start the iKeyman user interface by running the `strmqikm` command.
3. Click **Key Database File > Open**.
4. Click **Key database type** and select **PKCS11Direct**

5. In the **File Name** field, type the name of the module for managing your cryptographic hardware; for example, PKCS11_API.so.

If you are using certificates or keys stored on PKCS #11 cryptographic hardware, note that iKeycmd and iKeyman are 64-bit programs. External modules required for PKCS #11 support will be loaded into a 64-bit process, therefore you must have a 64-bit PKCS #11 library installed for the administration of cryptographic hardware. The Windows and Linux x86 32-bit platforms are the only exceptions, as the iKeyman and iKeycmd programs are 32-bit on those platforms.

6. In the **Location** field, enter the path:
 - On UNIX and Linux systems, this might be /usr/lib/pksc11, for example.
 - On Windows systems, you can type the library name; for example, cryptoki.

Click **OK**. The Open Cryptographic Token window opens.

7. In the **Cryptographic Token Password** field, type the password that you set when you configured the cryptographic hardware.
8. If your cryptographic hardware has the capacity to hold the signer certificates required to receive or import a personal certificate, clear both secondary key database check boxes and continue from step 12. If you require a secondary CMS key database to hold the signer certificates, select either **Open existing secondary key database file** or **Create new secondary key database file**.
9. In the **File Name** field, type a file name. This field already contains the text key.kdb. If your stem name is key, leave this field unchanged. If you specified a different stem name, replace key with your stem name. You must not change the .kdb suffix.
10. In the **Location** field, type the path, for example:
 - For a queue manager: /var/mqm/qmgrs/QM1/ssl
 - For a IBM WebSphere MQ MQI client: /var/mqm/ssl

Click **OK**. The Password Prompt window opens.

11. Enter a password.
 - If you selected **Open existing secondary key database file** in step 8, type a password in the **Password** field.
 - If you selected **Create new secondary key database file** in step 8:
 - a. Type a password in the **Password** field, and type it again in the **Confirm Password** field.
 - b. Select **Stash the password to a file**. Note that if you do not stash the password, attempts to start SSL channels fail because they cannot obtain the password required to access the key database file.
 - c. Click **OK**. A window opens, confirming that the password is in file key.sth (unless you specified a different stem name).
12. Click **OK**. The Key database content frame displays.

Using a command line interface

1. Use the following commands to create a new CMS key database file for either a queue manager or a IBM WebSphere MQ MQI client by using either the **iKeycmd** or **runmqakm** command:
 - On UNIX, Linux, and Windows systems:

```
runmqckm -keydb -create -db filename -pw password -type cms -stash
```
 - Using runmqakm:

```
runmqakm -keydb -create -db filename -pw password -type cms  
-stash -fips -strong
```

where:

-db *filename*

Specifies the fully qualified file name of a CMS key database, and must have a file extension of .kdb.

-pw *password*

Specifies the password for the CMS key database.

-type *cms*

Specifies the type of database. (For IBM WebSphere MQ, it must be *cms*.)

-stash

Saves the key database password to a file.

-fips

Disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that are FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the **runmqakm** command fails.

-strong

Checks that the password entered satisfies the minimum requirements for password strength. The minimum requirements for a password are as follows:

- The password must be a minimum length of 14 characters.
- The password must contain a minimum of one lowercase character, one uppercase character, and one digit or special character. Special characters include the asterisk (*), the dollar sign (\$), the number sign (#), and the percent sign (%). A space is classified as a special character.
- Each character can occur a maximum of three times in a password.
- A maximum of two consecutive characters in the password can be identical.
- All characters are in the standard ASCII printable character set within the range 0x20 - 0x7E.

Requesting a personal certificate for your PKCS #11 hardware:

Use this procedure for either a queue manager or a IBM WebSphere MQ MQI client to request a personal certificate for your cryptographic hardware.

Using the iKeyman user interface:

About this task

Note: WebSphere MQ does not support SHA-3 or SHA-5 algorithms. You can use the digital signature algorithm names SHA384WithRSA and SHA512WithRSA because both algorithms are members of the SHA-2 family.

The digital signature algorithm names SHA3WithRSA and SHA5WithRSA are deprecated because they are an abbreviated form of SHA384WithRSA and SHA512WithRSA respectively.

Procedure

To request a personal certificate from the iKeyman user interface, complete the following steps:

1. Complete the steps to work with your cryptographic hardware. See "Managing certificates on PKCS #11 hardware" on page 655.
2. From the **Create** menu, click **New Certificate Request**. The Create New Key and Certificate Request window opens.
3. In the **Key Label** field, enter the following labels:
 - For a queue manager, enter `ibmwebsphermq` followed by the name of your queue manager changed to lowercase. For example, for a queue manager called QM1, enter `ibmwebsphermqm1`.
 - For a IBM WebSphere MQ MQI client, enter `ibmwebsphermq` followed by your logon user ID, all in lowercase; for example, `ibmwebsphermqmyuserid`.

4. Enter values for **Common Name** and **Organization**, and select a **Country**. For the remaining optional fields, either accept the default values, or type or select new values. Note that you can supply only one name in the **Organizational Unit** field. For more information about these fields, see “Distinguished Names” on page 373.
5. In the **Enter the name of a file in which to store the certificate request** field, either accept the default `certreq.arm`, or type a new value with a full path.
6. Click **OK**. A confirmation window opens.
7. Click **OK**. The **Personal Certificate Requests** list shows the label of the new personal certificate request you created. The certificate request is stored in the file you chose in step 5.
8. Request the new personal certificate either by sending the file to a certificate authority (CA), or by copying the file into the request form on the website for the CA.

Using the command line:

Procedure

Use the following commands to request a personal certificate by using either the **iKeycmd** or **runmqkm** command:

- Using **iKeycmd** on UNIX, Linux, and Windows systems:

```
runmqckm -certreq -create -db filename -pw password -label label
        -dn distinguished_name -size key_size -file filename -sig_alg algorithm
```

Instead of `-dn distinguished_name`, you can use `-san_dsname DNS_names`, `-san_emailaddr email_addresses`, or `-san_ipaddr IP_addresses`.

- Using **runmqakm**:

```
runmqakm -certreq -create -db filename -pw password -label label
        -dn distinguished_name -size key_size -file filename -fips
        -sig_alg algorithm
```

where:

-db filename

Specifies the fully qualified file name of a CMS key database.

-pw password

Specifies the password for the CMS key database.

-label label

Specifies the key label attached to the certificate.

-dn distinguished_name

Specifies the X.500 distinguished name enclosed in double quotation marks. At least one attribute is required. You can supply multiple OU and DC attributes.

-size key_size

Specifies the key size. If you are using **iKeycmd**, the value can be 512 or 1024. If you are using **runmqakm**, the value can be 512, 1024, or 2048.

-file filename

Specifies the file name for the certificate request.

-fips

Specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that are FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the **runmqakm** command fails.

-sig_alg

For **iKeycmd**, specifies the asymmetric signature algorithm used for the creation of the entry's key pair. The value can be `MD2_WITH_RSA`, `MD2withRSA`, `MD5_WITH_RSA`, `MD5withRSA`, `SHA1withDSA`, `SHA1withRSA`,

SHA256_WITH_RSA, SHA256WithRSA, SHA2WithRSA, SHA384_WITH_RSA, SHA384WithRSA, SHA512_WITH_RSA, SHA512WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, SHAWithDSA, or SHAWithRSA. The default value is SHA1WithRSA

-sig_alg

For **runmqakm**, specifies the hashing algorithm used during the creation of a certificate request. This hashing algorithm is used to create the signature associated with the newly created certificate request. The value can be md5, MD5_WITH_RSA, MD5WithRSA, SHA_WITH_DSA, SHA_WITH_RSA, sha1, SHA1WithDSA, SHA1WithECDSA, SHA1WithRSA, sha224, SHA224_WITH_RSA, SHA224WithDSA, SHA224WithECDSA, SHA224WithRSA, sha256, SHA256_WITH_RSA, SHA256WithDSA, SHA256WithECDSA, SHA256WithRSA, SHA2WithRSA, sha384, SHA384WithECDSA, SHA384WithRSA, sha512, SHA512_WITH_RSA, SHA512WithECDSA, SHA512WithRSA, SHAWithDSA, SHAWithRSA, EC_ecdsa_with_SHA1, EC_ecdsa_with_SHA224, EC_ecdsa_with_SHA256, EC_ecdsa_with_SHA384, or EC_ecdsa_with_SHA512. The default value is SHA1WithRSA.

-san_dnsname *DNS_names*

Specifies a comma-delimited or space-delimited list of DNS names for the entry being created.

-san_emailaddr *email_addresses*

Specifies a comma-delimited or space-delimited list of email addresses for the entry being created.

-san_ipaddr *IP_addresses*

Specifies a comma-delimited or space-delimited list of IP addresses for the entry being created.

If you are using cryptographic hardware, see “Requesting a personal certificate for your PKCS #11 hardware” on page 657.

Importing a personal certificate to your PKCS #11 hardware:

Use this procedure for either a queue manager or a IBM WebSphere MQ MQI client to import a personal certificate to your cryptographic hardware.

Using iKeyman:

Procedure

To request a personal certificate from the iKeyman user interface, complete the following steps:

1. Complete the steps to work with your cryptographic hardware. See “Managing certificates on PKCS #11 hardware” on page 655.
2. Click **Receive**. The Receive Certificate from a File window opens.
3. Select the **Data type** of the new personal certificate; for example, Base64–encoded ASCII data for a file with the .arm extension.
4. Type the certificate file name and location for the new personal certificate, or click **Browse** to select the name and location.
5. Click **OK**. If you already have a personal certificate in your key database a window opens, asking if you want to set the key you are adding as the default key in the database.
6. Click **Yes** or **No**. The Enter a Label window opens.
7. Type a label. For example, you might use the same label as when you requested the personal certificate. Note that the label must be in the correct IBM WebSphere MQ format:
 - For a queue manager, `ibmwebspheremq` followed by the name of your queue manager in lowercase. For example, for a queue manager called QM1, the label would be: `ibmwebspheremqqm1`.
 - For a IBM WebSphere MQ MQI client, `ibmwebspheremq` followed by your logon user ID in lowercase. For example, for a user ID MyUserID, the label would be: `ibmwebspheremqmyuserid`.
8. Click **OK**. The **Personal Certificates** list shows the label of the new personal certificate you added. This label is formed by adding the cryptographic token label before the label you supplied.

Using the command line:

Procedure

To request a personal certificate from a command line, complete the following steps:

1. Open a command window that is configured for your environment.
2. Enter the appropriate command for your operating system and configuration:
 - On Windows, UNIX and Linux systems, use one of the following commands:

```
runmqckm -cert -receive -file filename -crypto path  
-tokenlabel hardware_token -pw hardware_password -format cert_format  
runmqakm -cert -receive -file filename -crypto path  
-tokenlabel hardware_token -pw hardware_password -format cert_format -fips
```

where:

-file filename

Specifies the fully qualified file name of the file containing the personal certificate.

-crypto path

Specifies the fully qualified path to the PKCS #11 library supplied with the hardware.

-tokenlabel hardware_token

Specifies the label given to the storage part of the cryptographic hardware during installation.

-pw hardware_password

Specifies the password for access to the hardware.

-format cert_format

Specifies the format of the certificate. The value can be `asci i` for Base64-encoded ASCII or `binary` for binary DER data. The default is ASCII.

-fips

Specifies that the command is run in FIPS mode. This mode disables the use of the BSafe cryptographic library. Only the ICC component is used and this component must be successfully initialized in FIPS mode. When in FIPS mode, the ICC component uses algorithms that are FIPS 140-2 validated. If the ICC component does not initialize in FIPS mode, the `runmqakm` command fails.

Certificate validation and trust policy design on UNIX, Linux, and Windows systems:

WebSphere MQ validates SSL or TLS certificates according to two types of policy, basic, and standard. Standard policy checking conforms to RFC 5280.

The information in these topics applies to the following systems:

- WebSphere MQ for UNIX and Linux systems
- WebSphere MQ for Windows systems

The following terms are used in this section:

Certificate policy

Determines which fields in a certificate are understood and processed.

OCSP policy

Determines which fields in an OCSP request or response are understood and processed.

CRL policy

Determines which fields in a certificate revocation list are understood and processed.

Path validation policy

Determines how the certificate, OCSP, and CRL policy types interact with each other to determine whether a certificate chain (a trust point "RootCA" to an end-entry "EE") is valid.

The basic and standard path validation policies are described separately because it reflects the implementation within WebSphere MQ for UNIX, Linux, and Windows systems. However, the standard OCSP and CRL policies are the same as the basic policies, and the standard certificate policy is an extended version of the basic policy, so these policies are not described separately.

By default, WebSphere MQ applies basic policy validation first. If basic policy validation fails, WebSphere MQ applies standard policy (RFC 5280) validation. If basic policy validation succeeds, standard policy validation is not applied. Thus, a validation failure means that both basic and standard policy validation failed, possibly for different reasons. A validation success means that either basic policy validation succeeded and standard policy validation was therefore not applied, or basic policy validation failed and standard policy validation succeeded.

Enforcing strict RFC 5280 compliance

To enforce strict RFC 5280 compliance, use the certificate validation policy configuration setting. This setting allows you to disable the basic policy, so that only the standard RFC 5280 policy is used. For more information about the certificate validation policy configuration setting, see “Certificate validation policies in WebSphere MQ” on page 402.

The following examples are digital certificates which are accepted by the basic certificate validation policy, but which are rejected by the RFC 5280 compliant standard policy. In order for a digital certificate chain to be trusted, the entire chain must satisfy the configured validation policy.

To view the full details of a digital certificate, use the **runmqakm** command:

```
runmqakm -cert -details -db key.kdb -pw password -label certificate_label
```

A certificate which has trust status enabled in the **runmqakm** output is not necessarily trusted for use in an SSL or TLS handshake. Trust status enabled means that the certificate is eligible to be used as a CA certificate to verify other certificates, if the certificate also satisfies the rules of the certificate validation policy. For more information about the RFC 5280 compliant standard certificate validation policy, see “Standard path validation policy” on page 670 (*WebSphere MQ V7.1 Reference*).

Example certificate 1 - incorrect key usage

This example shows a certificate where the key usage field does not comply with the standard certificate validation policy rules for a CA certificate. One of the requirements for a certificate to be valid for use as a CA certificate is that the key usage field must indicate that it is permitted to sign other certificates using the keyCertSign flag. A certificate without this flag cannot be used as a CA certificate.

```
Label : root
Key Size : 1024
Version : X509 V3
Serial : 54cb6f740c7ee410
Issuer : CN=Example Root CA,O=Example,C=GB
Subject : CN=Example Root CA,O=Example,C=GB
Not Before : 9 February 2012 17:19:00 GMT
Not After : 1 October 2019 18:19:00 GMT+01:00
Public Key
30 81 9F 30 0D 06 09 2A 86 48 86 F7 0D 01 01 01
05 00 03 81 8D 00 30 81 89 02 81 81 00 CC 44 D9
25 6D 26 1C 9D B9 FF DE B8 AC 44 AB E3 64 80 44
AF BE E0 00 93 53 92 33 F8 7E BD D7 71 ED 21 52
24 75 DF D6 EE 3C 54 97 84 29 EA 93 4C 4A D1 19
5D C1 A0 82 F5 74 E1 AD D9 87 10 D5 6A 2B 6F 90
04 0F 7E 6E 85 6D 32 99 33 9C D9 BB 57 86 DE 68
23 C9 F2 6D 53 E3 F5 FF D1 0B E7 23 19 3A F6 70
6B C8 C7 EB DB 78 8E 8C 9E 55 58 66 B6 31 DB 40
5F 6A 97 AB 12 D7 E2 3E 2E 79 EE 78 7B 02 03 01
```

```

00 01
Public Key Type : RSA (1.2.840.113549.1.1.1)
Fingerprint : SHA1 :
    EE 68 D4 4F 73 4F F4 21 DE 1A 01 11 5E DE B1 B8
    DF 40 AA D8
Fingerprint : MD5 :
    50 B5 E9 B2 D7 35 05 6A DC 6D 4B 1E B2 F2 DF A4
Fingerprint : SHA256 :
    B4 D7 6E C4 47 26 24 C7 4F 41 C3 83 03 6F 5C C7
    07 11 61 E0 0E 36 59 1F 1C E6 69 39 2D 18 05 D2
Extensions
    basicConstraints
        ca = true
        pathLen = 1239876
        critical
    key usage: encipherOnly
Signature Algorithm : SHA256WithRSASignature (1.2.840.113549.1.1.11)
Value
    9D AE 54 A9 9D 68 01 68 15 B5 53 9F 96 C9 5B D1
    52 40 DB CB 33 AF FD B9 26 D5 90 3F 1E 0B FC A6
    D9 8C 04 90 EB AA FD A8 7A 3C AB 60 5F 20 4F 0D
    7B 73 41 27 6A 2B BF 8C 99 91 B6 49 96 82 6A 24
    0A E8 B9 A5 AF 69 3D 2C A3 3C C8 12 39 FB 56 58
    4E 2A FE AC AC 10 89 53 B1 8F 0F C0 50 BF 5E 00
    91 64 B4 A1 4C 9A 4E D5 1F 38 7C AD 32 A9 8A E1
    91 16 2C 6D 1E 4A CA 99 8D CC 22 CD BF 90 49 FC
Trust Status : Enabled

```

In this example, the key usage field contains only the encipherOnly flag. The keyCertSign flag is not set, so this certificate is not permitted to sign other certificates. It therefore cannot be used as a CA certificate.

Example certificate 2 - missing basic constraints extension

This example shows a certificate which lacks the basic constraints extension. The basic constraints extension is used to indicate whether this certificate is permitted for use as a CA. It is also used to indicate the maximum length of any certificate chain which can be signed by the certificate. The standard certificate validation policy requires that the certificate has a basic constraints extension with the isCA flag set in order to be used as a CA.

```

Label : root
Key Size : 1024
Version : X509 V3
Serial : 1c7dfea316570bf6
Issuer : CN=Second Example Root CA,O=Example,C=GB
Subject : CN=Second Example Root CA,O=Example,C=GB
Not Before : 9 February 2012 17:18:22 GMT
Not After : 1 October 2019 18:18:22 GMT+01:00
Public Key
    30 81 9F 30 0D 06 09 2A 86 48 86 F7 0D 01 01 01
    05 00 03 81 8D 00 30 81 89 02 81 81 00 B2 70 49
    7C AE 1B A7 B3 06 49 6C 99 19 BC A8 77 BE 86 33
    21 6B C9 26 CC A6 28 52 9F 7B CF 03 A4 37 A7 4D
    6B 06 AA ED 7D 58 E3 70 F3 F7 C1 06 DA E8 27 C6
    3D 1B AC FA EF AA 59 7A 9A AB C1 14 4E AF 13 14
    4B 71 CA 8D FE C3 F5 2F E8 AC AD EF 21 80 6D 12
    89 4A 2A 84 AA 9D E0 4F C1 93 B1 3E 16 E8 3C 75
    39 2A 74 1E 90 CC B1 C3 2B 1D 55 26 76 D2 65 C1
    06 47 2A BF 79 96 42 76 A9 6E 65 88 5F 02 03 01
    00 01
Public Key Type : RSA (1.2.840.113549.1.1.1)
Fingerprint : SHA1 :

```

```

33 9F A1 81 43 F1 43 95 48 A5 66 B4 CD 98 E8 15
9C B3 CA 90
Fingerprint : MD5 :
91 EA D9 C0 2C 05 5B E2 CD 0B F6 DD 8A 11 44 23
Fingerprint : SHA256 :
62 46 35 0B 0E A1 A7 2A D5 74 70 0F AA 47 9A 9C
6B 80 1B F1 0B 4C 81 05 85 0E 91 11 A4 21 D2 34
Extensions
key usage: digitalSignature, keyCertSign
Signature Algorithm : SHA256WithRSASignature (1.2.840.113549.1.1.11)
Value
79 34 BA 5B 6F DC 06 A3 99 24 4E 8A 2B 27 05 47
0D 4D BE 6A 77 D1 1D 5F 54 82 9D CC F6 92 D4 9A
AB 4D B6 DD 6E AD 86 C3 6A A3 32 E3 B3 ED E0 62
4A EB 51 08 AC BE 49 9E 9C D7 FE AE C8 9D 17 16
68 31 6B F4 BA 74 1E 4F 5F 05 48 9F E7 46 BA DC
17 7A 60 88 F8 5B DB 3C 51 D4 98 97 28 82 CF 36
47 DA D2 0F 47 FF 70 EA 45 3A 49 66 E6 E2 F9 67
2C C8 3E 24 A2 3B EC 76 1F D6 31 2B BD A9 B5 08
Trust Status : Enabled

```

In this example, the certificate lacks the basic constraints field entirely. Therefore this certificate cannot be used as a CA certificate.

Example certificate 3 - intermediate CA with old version of X.509

This example shows an intermediate CA certificate which is at X.509 version 1. The standard certificate validation policy requires that all intermediate CA certificates must be at least X.509 version 3. Root CA certificates are exempt from this requirement as there are still some commonly used version 1 root CA certificates in existence. However, this exemption might change in future.

```

Label : intermediate
Key Size : 1024
Version : X509 V1
Serial : 02
Issuer : CN=Test Root CA,O=Example,C=GB
Subject : CN=Test Intermediate CA,O=Example,C=GB
Not Before : 10 February 2012 17:33:45 GMT
Not After : 11 April 2018 18:33:45 GMT+01:00
Public Key
30 81 9F 30 0D 06 09 2A 86 48 86 F7 0D 01 01 01
05 00 03 81 8D 00 30 81 89 02 81 81 00 C0 07 C2
D0 9F 84 DB 7C 20 8F 51 F9 C2 1A 3F CF E2 D7 F2
F1 56 F2 A4 8F 8F 06 B7 3B 01 31 DE 7C CC 03 63
AA D3 2F 1C 50 15 E3 56 80 40 7D FF 75 87 D3 F3
00 89 9A 26 F5 57 05 FA 4F ED 3B DD 93 FA F2 DF
38 26 D4 3A 92 51 CC F3 70 27 42 7A 9F AD 51 45
67 B7 AE 11 AD 4F 2D AB D2 CF 73 E6 F0 45 92 F0
47 16 66 7E 01 C7 76 A3 7B EC D2 76 3F E5 15 EC
D7 72 2C FE 14 F5 78 83 AA C4 20 AB F7 02 03 01
00 01
Public Key Type : RSA (1.2.840.113549.1.1.1)
Fingerprint : SHA1 :
DE BB 75 4B 14 E1 44 B9 B6 44 33 97 49 D0 82 6D
81 F2 2F DE
Fingerprint : MD5 :
72 49 44 42 E2 E6 89 F1 CC 37 C9 F6 B5 8F F3 AE
Fingerprint : SHA256 :
83 A4 52 AF 49 34 F1 DC 49 E6 95 AE 93 67 80 13
C2 64 D9 26 22 A0 E8 0A 5A A9 71 EC E8 33 E1 D1
Signature Algorithm : SHA256WithRSASignature (1.2.840.113549.1.1.11)
Value

```

```
40 4A 09 94 A0 18 07 5E 96 D7 A6 52 6B 8D 20 50
E8 91 F7 7E EA 76 B4 08 DF 76 66 1F FA FF 91 79
2E E0 66 8B 9F 40 FA 14 13 79 81 DB 31 A5 55 1D
44 67 41 F4 EA 1A F7 83 4F 21 F4 43 78 4E F8 5E
6F B2 B8 3A F7 6B B4 F5 C6 F8 EB 4C BF 62 6F 3E
C7 20 EC 53 B3 40 51 36 C1 0A 4E 73 ED 74 D1 93
02 C5 FB 61 F7 87 64 A5 94 06 7D 25 7C E3 73 DD
08 D4 07 D0 A4 3F 77 88 12 59 DB A4 DB 68 8F C1
Trust Status : Enabled
```

In this example, the version field is X.509 V1. This certificate is an X.509 version 1 certificate and therefore cannot be used as an intermediate CA.

Basic and standard certificate policies:

The basic and standard certificate policies support the same fields: the standard policy supports additional certificate extensions.

The supported fields for both the basic and standard policies are as follows:

- OuterSigAlgID²
- Signature³
- Version
- SerialNumber
- InnerSigAlgID⁴
- Issuer
- Validity
- SubjectName
- SubjectPublicKeyInfo
- IssuerUniqueID
- SubjectUniqueID

The supported extensions for the basic policy are as follows. Where an entry is marked as "not supported", WebSphere MQ does not attempt to process extensions containing a field of that specific type, but does process other types of the same extension.

- AuthorityKeyID
- AuthorityInfoAccess
- SubjectKeyID
- IssuerAltName
- SubjectAltName
- KeyUsage
- BasicConstraints
- PrivateKeyUsage
- CRLDistributionPoints
 - DistributionPoint
 - DistributionPointName (X.500 Name and LDAP Format URI only)
 - NameRelativeToCRLIssuer (not supported)

2. This field is called *signatureAlgorithm* in RFC 5280.

3. This field is called *signatureValue* in RFC 5280.

4. This field is called *signature* in RFC 5280.

- Reasons (ignored)
- CRLIssuer fields (not supported)

The supported extensions for the standard policy are all those listed for the basic policy and those in the following list. Where an entry is marked as "not supported", WebSphere MQ does not attempt to process extensions containing a field of that specific type, but does process other types of the same extension.

- NameConstraints
- ExtendedKeyUsage
- CertificatePolicies
 - PolicyInformation
 - PolicyIdentifier
 - PolicyQualifiers (not supported)
- PolicyMappings
- PolicyConstraints

Basic and standard OCSP policies:

The basic and standard OCSP policies support the same fields.

The supported fields for a request are as follows. Where an entry is marked as "not supported", WebSphere MQ does not attempt to process a request containing a field of that specific type, but does process other requests containing the same higher-level field.

- Signature (Optional)
- Version (Version 1 Only)
- RequesterName (Optional)
- RequestList (single request only)
 - CertID ⁵
 - singleRequestExtensions (not supported)
- RequestExtensions
 - Nonce (if enabled)

The supported fields for a response are as follows:

- ResponseStatus
- Response
 - responseType (id-pkix-ocsp-basic)
 - BasicOCSPResponse
 - Signature
 - Certs
 - Extensions
 - extendedKeyUsage
 - id-kp-OCSPSigning
 - id-pkix-ocsp-nocheck
 - ResponseData
 - Version (Version 1 Only)
 - ResponderID (by name or by hash)
 - ProducedAt (ignored)

5. This field is called reqCert in RFC 2560

- Responses (multiple responses supported)
 - SingleResponse
 - certID
 - certStatus
 - RevokedInfo (ignored)
 - thisUpdate (ignored)
 - nextUpdate
 - singleExtensions (ignored)
 - responseExtensions
 - Nonce (if enabled)

Basic and standard CRL policies:

The basic and standard CRL policies support the same fields and extensions.

The supported fields for these policies are as follows:

- OuterSigAlgID⁶
- Signature⁷
- Version
- InnerSigAlgID⁸
- Issuer
- ThisUpdate
- NextUpdate
- RevokedCertificate
 - UserCertificate
 - RevocationDate

There are no supported CRLEntry extensions.

The supported CRL extensions for these policies are as follows. Where an entry is marked as "not supported", WebSphere MQ does not attempt to process extensions containing a field of that specific type, but does process other types of the same extension.

- AuthorityKeyID
- IssuerAltName
- CRLNumber
- IssuingDistributionPoint
 - DistributionPoint
 - DistributionPointName
 - FullName (X.500 Name and LDAP Format URI only)
 - NameRelativeToCRLIssuer (not supported)
 - Reasons (ignored)
 - CRLIssuer
 - OnlyContainsUserCerts (not supported)
 - OnlyContainsCACerts (not supported)

6. This field is called *signatureAlgorithm* in RFC 5280.

7. This field is called *signatureValue* in RFC 5280.

8. This field is called *signature* in RFC 5280.

- OnlySomeReasons (not supported)
- IndirectCRL⁹ (rejected)

Related concepts:

“Basic and standard CRL policies” on page 666 (*WebSphere MQ V7.1 Reference*)

Related reference:

“Basic and standard OCSP policies” on page 665 (*WebSphere MQ V7.1 Reference*)

Basic path validation policy:

The basic path validation policy determines how the certificate, OCSP, and CRL policy types interact with each other to determine if a certificate chain is valid.

The validation of a chain is performed in the following manner (but not necessarily in the following order):

1. Ensure that the name of the certificate's issuer is equal to the subject name in the previous certificate, and that there is not an empty issuer name in this certificate or the previous certificate subject name. If no previous certificate exists in the path and this is the first certificate in the chain, ensure that the issuer and subject name are identical and that the trust status is set for the certificate¹⁰.

Note: WebSphere MQ for UNIX, Linux and Windows systems will fail path validation in situations where the previous certificate in a path has the same subject name as the current certificate.

2. Ensure that the signature algorithm used to actually sign the certificate matches the signature algorithm indicated within the certificate, by ensuring that the issuer signature algorithm identifier in the certificate matches the algorithm identifier in the signature data.
3. Ensure that the certificate was signed by the issuer, using the subject public key from the previous certificate in the path to verify the signature on the certificate. If no previous certificate exists and this is the first certificate, use the subject public key of the certificate to verify the signature on it. WebSphere MQ supports DSA and RSA signature algorithms; however it does not support DSA Parameter Inheritance.
4. Ensure that the certificate is a known X509 version, unique IDs are not present for version 1 certificates, and extensions are not present for version 1 and version 2 certificates.
5. Ensure that the certificate has not expired, or not been activated yet, and that its validity period is good¹¹.
6. Ensure that there are no unknown critical extensions or any duplicate extensions.
7. Ensure that the certificate has not been revoked. Here, the following operations apply:
 - a. If the OCSP connection is enabled and a Responder Address is configured or the Certificate has a valid AuthorityInfoAccess extension specifying a HTTP format GENERALNAME_uniformResourceID check revocation status with OCSP.
 - b. If revocation status from 7a above is undetermined the CRLDistributionPoints extension is checked for a list of X.500 distinguished name GENERALNAME_directoryname and URI GENERALNAME_uniformResourceID. Only LDAP, HTTP and FILE format URIs are supported. If the extension is not present, or use of the CRLDistributionPoints extension results in

9. IndirectCRL extensions will result in CRL validation failing. IndirectCRL extensions must not be used because they cause identified certificates to not be rejected.

10. Trust status is an administrative setting in the key database file. You can access and alter the trust status of a particular signer certificate in iKeyman. Select the required certificate from the signer list and click **View/Edit...** The **Set the certificate as a trusted root** check box on the resulting panel indicates the trust status. You can also set Trust status using iKeycmd or

runqmakm with the -trust flag on the -cert -modify command. For further information about this command, see  Managing keys and certificates (*WebSphere MQ V7.1 Reference*).

11. There are no checks to ensure the subject's validity is within bounds of the issuer's validity. This is not required, and it has been shown that certificates from some CAs do not pass such a check.

undetermined status and the extension is not Critical, the certificate's issuer's name is used to query revocation status. A CRL database (LDAP) is then queried for CRLs. If the certificate is not the last certificate, or if the last certificate has the basic constraint extension with the "isCA" flag turned on, the database is queried for ARLs and CRLs instead. If CRL checking is enabled, and no CRL database can be queried, the certificate is treated as revoked. Currently, the X500 directory name form and the LDAP/HTTP/FILE URI forms are the only supported name forms used to look up CRLs and ARLs¹².

Note: RelativeDistinguishedNames are not supported.

- c. If revocation status from both 7a on page 667 and 7b on page 667 is undetermined, WebSphere MQ checks the *OCSPAuthentication* configuration setting to decide whether to allow the connection.¹³
8. If the issuerAltName extension is marked critical, ensure that the name forms are recognized. The following general name forms are currently recognized:
 - rfc822
 - DNS
 - directory
 - URI
 - IPAddress(v4/v6)
9. If the subjectAltName extension is marked critical, ensure that the name forms are recognized. The following general name forms are currently recognized:
 - rfc822
 - DNS
 - directory
 - URI
 - IPAddress(v4/v6)
10. If the KeyUsage extension is critical on a non-EE certificate, ensure that the keyCertSign flag is on, and ensure that if the BasicConstraints extension is present, the "isCA" flag is true.
11. If the BasicConstraints extension is present, the following checks are made:
 - If the "isCA" flag is false, ensure the certificate is the last certificate in the chain and that the pathLength field is not present.
 - If the "isCA" flag is true and the certificate is NOT the last certificate in the chain, ensure that the number of certificates until the last certificate in the chain is not greater than the pathLength field.
12. The AuthorityKeyID extension is not used for path validation, but is used when building the certificate chain.
13. The SubjectKeyID extension is not used for path validation, but is used when building the certificate chain.
14. The PrivateKeyUsagePeriod extension is ignored by the validation engine, because it cannot determine when the CA actually signed the certificate. The extension is always non-critical and therefore can be safely ignored.

An OCSP Response is also validated to ensure that the response itself is valid. Validation is performed in the following manner (but not necessarily the following order):

1. Ensure that response status is Successful and the response type is PKIX_AD_OCSP_basic.r
2. Ensure that response version data is present and the response is the correct version (Version 1)

12. After they are retrieved from the database, ARLs are evaluated in exactly the same fashion as CRLs. Many CAs do not issue ARLs. However, WebSphere MQ will look for ARLs and CRLs if checking a CA certificate for revocation status.

13. If *OCSPAuthentication* is set to WARN, WebSphere MQ logs the unknown revocation status and allows the connection to continue.

3. Ensure that the response is correctly signed. The signature will be rejected if the signer does not meet at least one of the following criteria:
 - The signer matches a local configuration of OCSP signing authority¹⁴ for the certificate.
 - The signer is using the CA key for which the public key is contained in the CA certificate, that is, the CA itself is directly signing the response.
 - The signer is a direct sub-ordinate of the CA that signed the certificate for which revocation information is being checked and is authorized by the CA by including the value of `id-ad-ocspSigning` in an `ExtendedKeyUsage` extension.

Note: Revocation checking of the response signer certificate is not performed if the `id-pkix-ocsp-nocheck` extension is present.

4. Ensure that response hash algorithm, `serialNumber`, `issuerNameHash`, and `issuerKeyHash` match those of the request.
5. Ensure that the response has not expired, that is, that the `nextUpdate` time is greater than the current time.¹⁵
6. Ensure that the certificate has valid revocation status.

The validation of a CRL is also performed to ensure that the CRL itself is valid, and is performed in the following manner (but not necessarily the following order):

1. Ensure that the signature algorithm used to actually sign the CRL matches the signature algorithm indicated within the CRL, by ensuring that the issuer signature algorithm identifier in the CRL matches the algorithm identifier in the signature data.
2. Ensure that the CRL was signed by the issuer of certificate in question, verifying that the CRL has been signed with the key of the certificate issuer.
3. Ensure that the CRL has not expired¹⁶, or not been activated yet, and that its validity period is good.
4. Ensure that if the version field is present, it is version 2. Otherwise the CRL is version 1 and must not have any extensions. However, WebSphere MQ for UNIX, Linux and Windows systems only verifies that no critical extensions are present for a version 1 CRL.
5. Ensure that the certificate in question is on the `revokedCertificates` field list and that the revocation date is not in the future.
6. Ensure that there are no duplicate extensions.
7. If unknown critical extensions, including critical entry extensions, are detected in the CRL, this causes identified certificates to be treated as revoked¹⁷ (provided the CRL passes all other checks).

14. This is a Certificate in the KeyStore a user has installed and that has Trust Status set.

In the case of PKCS#11 Devices in Common Criteria mode of operation, the Certificate MUST have the `CKA_TRUSTED` Attribute set.

15. If no current OCSP responses are returned from the responder, WebSphere MQ will attempt to use out of date responses in determining the revocation status of a Certificate. WebSphere MQ attempts to use out of date Responses so that security will not be adversely reduced.
16. If no current CRLs are found, WebSphere MQ for UNIX, Linux and Windows systems will attempt to use out of date CRLs to determine the revocation status of a Certificate. It is not clearly specified in RFC 5280 what action to take in the event of no current CRLs. WebSphere MQ for UNIX, Linux and Windows systems attempt to use out of date CRLs so that security will not be adversely reduced.
17. ITU X.509 and RFC 5280 are in conflict in this case because the RFC mandates that CRLs with unknown critical extensions must fail validation. However, ITU X.509 requires that identified certificates must still be treated as revoked provided the CRL passes all other checks. WebSphere MQ for UNIX, Linux and Windows systems adopt the ITU X.509 guidance so that security will not be adversely reduced.

A potential scenario exists where the CA that issues a CRL might set an unknown critical extension to indicate that even though all other validation checks are successful, a certificate which is identified must not be considered revoked and thus not rejected by the application. In this scenario, following X.509, WebSphere MQ for UNIX, Linux and Windows systems will function in a fail-secure mode of operation. That is, they might reject certificates that the CA did not intend to be rejected and therefore might

8. If the authorityKeyID extension in the CRL and the subjectKeyID in the CA certificate are present and if the keyIdentifier field is present within the authorityKeyID of the CRL, match it with the CACertificate's subjectKeyID.
9. If the issuerAltName extension is marked critical, ensure that the name forms are recognized. The following general name forms are currently recognized:
 - rfc822
 - DNS
 - directory
 - URI
 - IPAddress(v4/v6)
10. If the issuingDistributionPoint extension is present in the CRL, process as follows:
 - If the issuingDistributionPoint specifies an InDirectCRL then fail the CRL validation.
 - If the issuingDistributionPoint indicates that a CRLDistributionPoint is present but no DistributionPointName is found, fail the CRL validation
 - If the issuingDistributionPoint indicates that a CRLDistributionPoint is present and specifies a DistributionPointName ensure that it is a GeneralName or LDAP format URI that matches the name given by the certificate's CRLDistributionPoint or the certificate's issuer's name. If the DistributionPointName is not a GeneralName then the CRL validation will fail.

Note: RelativeDistinguishedNames are not supported and will fail CRL validation if encountered.

Standard path validation policy:

The standard path validation policy determines how the certificate, OCSP, and CRL policy types interact with each other to determine if a certificate chain is valid. Standard policy checking conforms to RFC 5280.

Path validation uses the following concepts:

- A certification path of length n , where the trust point or root certificate is certificate 1, and the EE is n .
- A set of initial policy identifiers (each comprising a sequence of policy element identifiers), that identifies one or more certificate policies, any one of which is acceptable for the purposes of certification path processing, or the special value "any-policy". Currently this is always set to "any-policy".

Note: WebSphere MQ for UNIX, Linux and Windows systems only supports policy identifiers that are created by WebSphere MQ for UNIX, Linux and Windows systems.

- Acceptable policy set: a set of certificate policy identifiers comprising the policy or policies recognized by the public key user, together with policies deemed equivalent through policy mapping. The initial value of the acceptable policy set is the special value "any-policy".
- Constrained subtrees: a set of root names defining a set of subtrees within which all subject names in subsequent certificates in the certification path can fall. The initial value is "unbounded".
- Excluded subtrees: a set of root names defining a set of subtrees within which no subject name in subsequent certificates in the certification path can fall. The initial value is "empty".
- Explicit policy: an integer which indicates if an explicit policy identifier is required. The integer indicates the first certificate in the path where this requirement is imposed. When set, this variable can be decreased, but cannot be increased. (That is, if a certificate in the path requires explicit policy identifiers, a later certificate cannot remove this requirement.) The initial value is $n+1$.

deny service to some valid users. A fail-insecure mode ignores a CRL because it has an unknown critical extension and therefore certificates that the CA intended to be revoked are still accepted. The administrator of the system should then query this behavior with the issuing CA.

- Policy mapping: an integer which indicates if policy mapping is permitted. The integer indicates the last certificate on which policy mapping may be applied. When set, this variable can be decreased, but cannot be increased. (That is, if a certificate in the path specifies policy mapping is not permitted, it cannot be overridden by a later certificate.) The initial value is n+1.

The validation of a chain is performed in the following manner (but not necessarily the following order):

1. The information in the following paragraph is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that the name of the certificate's issuer is equal to the subject name in the previous certificate, and that there is not an empty issuer name in this certificate or the previous certificate subject name. If no previous certificate exists in the path and this is the first certificate in the chain, ensure that the issuer and subject name are identical and that the trust status is set for the certificate¹⁸.
 If the certificate does not have a subject name, the subjectAltName extension must be present and critical.
2. The information in the following paragraph is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that the signature algorithm used to actually sign the certificate matches the signature algorithm indicated within the certificate, by ensuring that the issuer signature algorithm identifier in the certificate matches the algorithm identifier in the signature data.
 If both the certificate's issuersUniqueID and the issuer's subjectUniqueID are present, ensure they match.
3.
 The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that the certificate was signed by the issuer, using the subject public key from the previous certificate in the path to verify the signature on the certificate. If no previous certificate exists and this is the first certificate, use the subject public key of the certificate to verify the signature on it.
4.
 The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that the certificate is a known X509 version, unique IDs are not present for version 1 certificates and extensions are not present for version 1 and version 2 certificates.
5.
 The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that the certificate has not expired, or not been activated yet, and that its validity period is good¹⁹
6.
 The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):
 Ensure that there are no unknown critical extensions, nor any duplicate extensions.
7. The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):

18. Trust status is an administrative setting in the key database file. You can access and alter the trust status of a particular signer certificate in iKeyman. Select the required certificate from the signer list and click **View/Edit...** The **Set the certificate as a trusted root** check box on the resulting panel indicates the trust status. You can also set Trust status using iKeycmd or

runmqakm with the **-trust** flag on the **-cert -modify** command. For further information about this command, see  Managing keys and certificates (*WebSphere MQ V7.1 Reference*).

19. There are no checks to ensure the subject's validity is within bounds of the issuer's validity. This is not required, and certificates from some CAs have been shown to not pass such a check.

Ensure that the certificate has not been revoked. Here, the following operations apply:

- a. If the OCSP connection is enabled and a Responder Address is configured or the Certificate has a valid AuthorityInfoAccess extension specifying an HTTP format GENERALNAME_uniformResourceID check revocation status with OCSP.
 - 1) WebSphere MQ for UNIX and Windows systems allows the OCSP Request to be optionally signed for preconfigured responders but this has otherwise no impact on OCSP Response processing.
- b. If revocation status from 7a is undetermined the CRLDistributionPoints extension is checked for a list of X.500 distinguished name GENERALNAME_directoryname and URI GENERALNAME_uniformResourceID. If the extension is not present, the certificate's issuer's name is used. A CRL database (LDAP) is then queried for CRLs. If the certificate is not the last certificate, or if the last certificate has the basic constraint extension with the "isCA" flag turned on, the database is queried for ARL's and CRL's instead. If CRL checking is enabled, and no CRL database can be queried, the certificate is treated as revoked. Currently, the X500 directory name form and the LDAP/HTTP/FILE URI forms are the only supported name forms used to look up CRLs and ARLs¹⁵.

Note: RelativeDistinguishedNames are not supported.

8.

The following information is consistent with the basic path validation policy described in "Basic path validation policy" on page 667 (*WebSphere MQ V7.1 Reference*):

If the subjectAltName extension is marked critical, ensure that the name forms are recognized. The following general name forms are currently recognized:

- rfc822
 - DNS
 - directory
 - URI
 - IPAddress(v4/v6)
9. Ensure that the subject name and subjectAltName extension (critical or noncritical) is consistent with the constrained and excluded subtrees state variables.
 10. If the EmailAddress OID is present in the subject name field as an IA5 string, and there is no subjectAltName extension, the EmailAddress must be consistent with the constrained and excluded subtrees state variable.
 11. Ensure that policy information is consistent with the initial policy set:
 - a. If the explicit policy state variable is less than or equal to the current certificate's numeric sequence value, a policy identifier in the certificate shall be in the initial policy set.
 - b. If the policy mapping variable is less than or equal to the current certificate's numeric sequence value, the policy identifier cannot be mapped.
 12. Ensure that policy information is consistent with the acceptable policy set:
 - a. If the certificate policies extension is marked critical²⁰, the intersection of the policies extension and the acceptable policy set is non-null.
 - b. The acceptable policy set is assigned the resulting intersection as its new value.
 13. Ensure that the intersection of the acceptable policy set and the initial policy set is non-null. If the special Policy of anyPolicy is present then allow it only if it has not been inhibited by the inhibitAnyPolicy extension at this chain position.

20. This is maintained as a legacy requirement from RFC2459 (6.1 (e)(1))

14. If an `inhibitAnyPolicy` extension is present ensure that it is marked Critical and, if so, set the `inhibitAnyPolicy` state and chain position to the value of the integer value of the extension provided it is not greater than the current value. This is the number of certificates to allow with an `anyPolicy` Policy before disallowing the `anyPolicy` Policy.
15. The following steps are performed for all certificates except the last one:
 - a. If the `issuerAltName` extension is marked critical, ensure that the name forms are recognized. The following general name forms are currently recognized:
 - `rfc822`
 - `DNS`
 - `directory`
 - `URI`
 - `IPAddress(v4/v6)`
 - b.
 - 1) If the `BasicConstraints` extension is not present, the certificate is only valid as an EE certificate.
 - 2) If the `BasicConstraints` extension is present, ensure that the "isCA" flag is true²¹. If the `pathLength` field is present, ensure the number of certificates until the last certificate is not greater than the `pathLength` field.
 - c. If the `KeyUsage` extension is critical, ensure that the `keyCertSign` flag is on, and ensure that if the `BasicConstraints` extension is present, that the "isCA" flag is true²².
 - d. If a policy constraints extension is included in the certificate, modify the explicit policy and policy mapping state variables as follows:
 - i. If `requireExplicitPolicy` is present and has value r , the explicit policy state variable is set to the minimum of its current value and the sum of r and i (the current certificate in the sequence).
 - ii. If `inhibitPolicyMapping` is present and has value q , the policy mapping state variable is set to the minimum of its current value and the sum of q and i (the current certificate in the sequence).
 - e. If the `policyMappings` extension is present (see 12(b)), ensure that it is not critical, and if policy mapping is allowed, these mappings are used to map between this certificate's policies and its signee's policies.
 - f. If the `nameConstraints` extension is present, ensure that it is critical, and that the permitted and excluded subtrees adhere to the following rules before updating the chain's subtree's state in accordance with the algorithm described in RFC 5280 section 6.1.4 part (g):
 - 1) The minimum field is set to zero.
 - 2) The maximum field is not present.
 - 3) The base field name forms are recognized. The following general name forms are currently recognized:
 - `rfc822`
 - `DNS`
 - `directory`
 - `URI`
 - `IPAddress(v4/v6)`
16. The `ExtendedKeyUsage` extension is not checked by WebSphere MQ.
- 17.

21. "isCA" is always checked to ensure it is true to be as part of the chain building itself, however this specific test is still made.

22. This check is in fact redundant because of step (b), but the check is still made.

The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):

The AuthorityKeyID extension is not used for path validation, but is used when building the certificate chain.

18.

The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):

The SubjectKeyID extension is not used for path validation, but is used when building the certificate chain.

19.

The following information is consistent with the basic path validation policy described in “Basic path validation policy” on page 667 (*WebSphere MQ V7.1 Reference*):

The PrivateKeyUsagePeriod extension is ignored by the validation engine, because it cannot determine when the CA actually signed the certificate. The extension is always non-critical and therefore can be safely ignored.

Working with SSL or TLS on z/OS

This information describes how you set up and work with the Secure Sockets Layer (SSL) on z/OS.

Each topic includes examples of performing each task using RACF. You can perform similar tasks using the other external security managers.

On z/OS, you must also set the number of server subtasks that each queue manager uses for processing SSL calls, as described in “Setting the SSLTASKS parameter” on page 675.

z/OS SSL support is integral to the operating system, and is known as *System SSL*. System SSL is part of the Cryptographic Services Base element of z/OS. The Cryptographic Services Base members are installed in the *pdsname.SIEALNKE* partitioned data set (PDS). When you install System SSL, ensure that you choose the appropriate options to provide the CipherSpecs that you require.

Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

Related reference:



ALTER QMGR (*WebSphere MQ V7.1 Reference*)

Additional user ID requirements:

This information describes the additional requirements your user ID needs to set up and work with the Secure Sockets Layer (SSL) on z/OS.

Ensure that you have all the appropriate High Impact or Pervasive (HIPER) updates on your system.

Ensure that you have set up the following prerequisites:

- The *ssidCHIN* user ID is defined correctly in RACF, and that the *ssidCHIN* user ID has READ access to the following profiles:
 - IRR.DIGTCERT.LIST
 - IRR.DIGTCERT.LISTRING

These variables are defined in the RACF FACILITY Class.

- The *CHINIT* is the owner of the key ring.
- The personal certificate of the queue manager, if created by the RACDCERT command, is created with a certificate type user ID that is also the same as the *ssidCHIN* user ID.
- The channel initiator is recycled to pick up any changes you make to the key ring.

- The WebSphere MQ Channel Initiator procedure has access to the system SSL runtime library *pdsname.SIEALNKE* through the link list, LPA, or a STEPLIB DD statement. This library must be APF-authorized.
- The user ID under whose authority the channel initiator is running is configured to use UNIX System Services (USS), as described in the z/OS UNIX System Services Planning documentation.
The user ID must have an OMVS segment in its RACF profile, with a valid home directory, and its default group must also have a valid OMVS group ID defined in RACF. The group ID must be set with UID=00.

Setting the SSLTASKS parameter:

Use the ALTER QMGR command to set the number of server subtasks for processing SSL calls

To use SSL channels, ensure that there are at least two server subtasks by setting the SSLTASKS parameter, using the ALTER QMGR command. For example:

```
ALTER QMGR SSLTASKS(5)
```

To avoid problems with storage allocation, do not set the SSLTASKS parameter to a value greater than 50. The value of the TCPCHL parameter on the queue manager must be set to a value greater than zero, otherwise the value set for SSLTASKS is ignored.

Setting up a key repository on z/OS:

Set up a key repository at both ends of the connection. Associate each key repository with its queue manager.

An SSL or TLS connection requires a *key repository* at each end of the connection. Each queue manager must have access to a key repository. Use the SSLKEYR parameter on the ALTER QMGR command to associate a key repository with a queue manager. See “The SSL or TLS key repository” on page 390 for more information.

On z/OS, digital certificates are stored in a *key ring* that is managed by your External Security Manager (ESM). These digital certificates have labels, which associate the certificate with a queue manager. SSL and TLS use these certificates for authentication purposes. All the examples that follow use RACF commands. Equivalent commands exist for other ESM programs.

On z/OS, WebSphere MQ uses the *ibmWebSphereMQ* prefix on a label to avoid confusion with certificates for other products. The prefix is followed by the name of the queue manager.

The key repository name for a queue manager is the name of a key ring in your RACF database. You can specify the key ring name either before or after creating the key ring.

Use the following procedure to create a new key ring for a queue manager:

1. Ensure that you have the appropriate authority to issue the RACDCERT command (see the *SecureWay Security Server RACF Command Language Reference* for more details).

2. Issue the following command:

```
RACDCERT ID(userid1) ADDRING(ring-name)
```

where:

- *userid1* is the user ID of the channel initiator address space, or the user ID that is going to own the key ring (if the key ring is shared).
- *ring-name* is the name you want to give to your key ring. The length of this name can be up to 237 characters. This name is case-sensitive. Specify *ring-name* in uppercase characters to avoid problems.

Making CA certificates available to a queue manager on z/OS:

After you have created your key ring, connect any relevant CA certificates to it.

If you have the CA certificate in a dataset, you must first add the certificate to the RACF® database by using the following command:

```
RACDCERT ID(userid1) ADD(input-data-set-name) WITHLABEL('My CA')
```

Then to connect a CA certificate for My CA to your key ring, use the following command:

```
RACDCERT ID(userid1)  
CONNECT(CERTAUTH LABEL('My CA') RING(ring-name) USAGE(CERTAUTH))
```

where *userid1* is either the channel initiator user ID or the owner of a shared key ring.

For more information about CA certificates, refer to “Digital certificates” on page 371.

Locating the key repository for a queue manager on z/OS:

Use this procedure to obtain the location of your queue manager's key ring.

1. Display your queue manager's attributes, using either of the following MQSC commands:
DISPLAY QMGR ALL
DISPLAY QMGR SSLKEYR
2. Examine the command output for the location of the key ring.

Specifying the key repository location for a queue manager:

To specify the location of your queue manager's key ring, use the ALTER QMGR MQSC command to set your queue manager's key repository attribute.

For example:

```
ALTER QMGR SSLKEYR(CSQ1RING)
```

if the key ring is owned by the channel initiator address space, or:

```
ALTER QMGR SSLKEYR(userid1/CSQ1RING)
```

if it is a shared key ring, where *userid1* is the user ID that owns the key ring.

Giving the channel initiator the correct access rights:

The channel initiator (CHINIT) needs access to the key repository and to certain security profiles.

Granting the CHINIT access to read the key repository

If the key repository is owned by the CHINIT user ID, this user ID needs read access to the IRR.DIGTCERT.LISTRING profile in the FACILITY class, and update access otherwise. Grant access by using the PERMIT command with ACCESS(UPDATE) or ACCESS(READ) as appropriate:

```
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(userid) ACCESS(UPDATE)
```

where *userid* is the user ID of the channel initiator address space.

Granting the CHINIT read access to the appropriate CSF* profiles

For hardware support provided through the Integrated Cryptographic Service Facility (ICSF) to be used, ensure your CHINIT user ID has read access to the appropriate CSF* profiles in the CSFSERV class by using the following command:

```
PERMIT csf-resource CLASS(CSFSERV) ID(userid) ACCESS(READ)
```

where *csf-resource* is the name of the CSF* profile and *userid* is the user ID of the channel initiator address space.

Repeat this command for each of the following CSF* profile names:

- CSFPKD
- CSFPKE
- CSFPKI
- CSFDSG
- CSFDSV

If your certificate keys are stored in ICSF and your installation has established access control over keys stored in ICSF, ensure your CHINIT user ID has read access to the profile in the CSFKEYS class by using the following command:

```
PERMIT IRR.DIGTCERT.userid.* CLASS(CSFKEYS) ID(userid) ACCESS(READ)
```

where *userid* is the user ID of the channel initiator address space.

When changes to certificates or the key repository become effective on z/OS:

Changes become effective when the channel initiator starts or the repository is refreshed.

Specifically, changes to the certificates in the key ring and to the key repository attribute become effective on either of the following occasions:

- When the channel initiator is started or restarted.
- When the REFRESH SECURITY TYPE(SSL) command is issued to refresh the contents of the SSL key repository.

Creating a self-signed personal certificate on z/OS:

Use this procedure to create a self-signed personal certificate.

1. Generate a certificate and a public and private key pair using the following command:

```
RACDCERT ID(userid2) GENCERT  
SUBJECTSDN(CN('common-name')  
            T('title')  
            OU('organizational-unit')  
            O('organization')  
            L('locality')  
            SP('state-or-province')  
            C('country'))  
WITHLABEL('label-name')
```

2. Connect the certificate to your key ring using the following command:

```
RACDCERT ID(userid1)  
CONNECT(ID(userid2) LABEL('label-name') RING(ring-name) USAGE(PERSONAL))
```

where:

- *userid1* is the user ID of the channel initiator address space or owner of the shared key ring.

- *userid2* is the user ID associated with the certificate and must be the user ID of the channel initiator address space.
userid1 and *userid2* can be the same ID.
- *ring-name* is the name you gave the key ring in “Setting up a key repository on z/OS” on page 675.
- *label-name* must be in the correct WebSphere MQ format for a queue manager or queue-sharing group: *ibmWebSphereMQ* followed by the name of your queue manager or queue-sharing group, for example, *ibmWebSphereMQCSQ1*.

Requesting a personal certificate on z/OS:

Apply for a personal certificate using RACF.

To apply for a personal certificate, use RACF as follows:

1. Create a self-signed personal certificate, as in “Creating a self-signed personal certificate on z/OS” on page 677. This certificate provides the request with the attribute values for the Distinguished Name.
2. Create a PKCS #10 Base64-encoded certificate request written to a data set, using the following command:

```
RACDCERT ID(userid2) GENREQ(LABEL('label-name')) DSN(output-data-set-name)
```

 where *label-name* is the label used when creating the self-signed certificate, and *userid2* is the user ID associated with the certificate and must be the user ID of the channel initiator address space.
3. Send the data set to a Certificate Authority (CA) to request a new personal certificate.
4. When the signed certificate is returned to you by the Certificate Authority, add the certificate back into the RACF database, using the original label, as described in “Adding personal certificates to a key repository on z/OS” on page 679.

Creating a RACF signed personal certificate:

RACF can function as a certificate authority and issue its own CA certificate.

This section uses the term *signer certificate* to denote a CA certificate issued by RACF.

The private key for the signer certificate must be in the RACF database before you carry out the following procedure:

1. Use the following command to generate a personal certificate signed by RACF, using the signer certificate contained in your RACF database:

```
RACDCERT ID(userid2) GENCERT
SUBJECTSDN(CN('common-name')
            T('title')
            OU('organizational-unit')
            O('organization')
            L('locality')
            SP('state-or-province')
            C('country'))
WITHLABEL('label-name')
SIGNWITH(CERTAUTH LABEL('signer-label'))
```

2. Connect the certificate to your key ring using the following command:

```
RACDCERT ID(userid1)
CONNECT(ID(userid2) LABEL('label-name') RING(ring-name) USAGE(PERSONAL))
```

where:

- *userid1* is the user ID of the channel initiator address space or owner of the shared key ring.
- *userid2* is the user ID associated with the certificate and must be the user ID of the channel initiator address space.

- *userid1* and *userid2* can be the same ID.
- *ring-name* is the name you gave the key ring in “Setting up a key repository on z/OS” on page 675.
- *label-name* must be in the correct WebSphere MQ format for a queue manager or queue-sharing group: *ibmWebSphereMQ* followed by the name of your queue manager or queue-sharing group, for example, *ibmWebSphereMQCSQ1*.
- *signer-label* is the label of your own signer certificate.

Adding personal certificates to a key repository on z/OS:

Use this procedure to add or import a personal certificate to a key ring.

After the certificate authority sends you a new personal certificate, add it to the key ring using the following procedure:

1. Add the certificate to the RACF database using the following command:

```
RACDCERT ID(userid2) ADD(input-data-set-name) WITHLABEL('label-name')
```
2. Connect the certificate to your key ring using the following command:

```
RACDCERT ID(userid1)  
CONNECT(ID(userid2) LABEL('label-name') RING(ring-name) USAGE(PERSONAL))
```

where:

- *userid1* is the user ID of the channel initiator address space or owner of the shared key ring.
- *userid2* is the user ID associated with the certificate and must be the user ID of the channel initiator address space.
- *ring-name* is the name you gave the key ring in “Setting up a key repository on z/OS” on page 675.
- *input-data-set-name* is the name of the data set containing the CA signed certificate. The data set must be cataloged and must not be a PDS or a member of a PDS. The record format (RECFM) expected by RACDCERT is VB. RACDCERT dynamically allocates and opens the data set, and reads the certificate from it as binary data.
- *label-name* is the label name that was used when you created the original request. It must be in the correct WebSphere MQ format for a queue manager or queue-sharing group: *ibmWebSphereMQ* followed by the name of your queue manager or queue-sharing group, for example, *ibmWebSphereMQCSQ1*.

Exporting a personal certificate from a key repository on z/OS:

Export the certificate using the RACDCERT command.

On the system from which you want to export the certificate, use the following command:

```
RACDCERT ID(userid2) EXPORT(LABEL('label-name'))  
DSN(output-data-set-name) FORMAT(CERTB64)
```

where:

- *userid2* is the user ID under which the certificate was added to the key ring.
- *label-name* is the label of the certificate you want to extract.
- *output-data-set-name* is the data set into which the certificate is placed.
- CERTB64 is a DER encoded X.509 certificate that is in Base64 format. You can choose an alternative format, for example:

CERTDER

DER encoded X.509 certificate in binary format

PKCS12B64

PKCS #12 certificate in Base64 format

PKCS12DER

PKCS #12 certificate in binary format

Note that **PKCS12DER** is supported only on OS/390® V2.10 and z/OS V1.1 and subsequent releases.

Deleting a personal certificate from a key repository on z/OS:

Delete a personal certificate using the RACDCERT command.

Before deleting a personal certificate, you might want to save a copy of it. To copy your personal certificate to a data set before deleting it, follow the procedure in “Exporting a personal certificate from a key repository on z/OS” on page 679. Then use the following command to delete your personal certificate:

```
RACDCERT ID(userid2) DELETE(LABEL('label-name'))
```

where:

- *userid2* is the user ID under which the certificate was added to the key ring.
- *label-name* is the name of the certificate you want to delete.

Renaming a personal certificate in a key repository on z/OS:

Rename a certificate using the RACDCERT command.

If you do not want a certificate with a specific label to be found, but do not want to delete it, you can rename it temporarily using the following command:

```
RACDCERT ID(userid2) LABEL('label-name') NEWLABEL('new-label-name')
```

where:

- *userid2* is the user ID under which the certificate was added to the key ring.
- *label-name* is the name of the certificate you want to rename.
- *new-label-name* is the new name of the certificate.

This can be useful when testing SSL client authentication.

Associating a user ID with a digital certificate on z/OS:

WebSphere MQ can use a user ID associated with a RACF certificate as a channel user ID. Associate a user ID with a certificate by installing it under that user ID, or using a Certificate Name Filter.

The method described in this topic is an alternative to the platform-independent method for associating a user ID with a digital certificate, which uses channel authentication records. For more information about channel authentication records, see “Channel authentication records” on page 408.

When an entity at one end of an SSL channel receives a certificate from a remote connection, the entity asks RACF if there is a user ID associated with that certificate. The entity uses that user ID as the channel user ID. If there is no user ID associated with the certificate, the entity uses the user ID under which the channel initiator is running.

Associate a user ID with a certificate in either of the following ways:

- Install that certificate into the RACF database under the user ID with which you want to associate it, as described in “Adding personal certificates to a key repository on z/OS” on page 679.
- Use a Certificate Name Filter (CNF) to map the Distinguished Name of the subject or issuer of the certificate to the user ID, as described in “Setting up a certificate name filter” on page 681.

Setting up a certificate name filter:

Use the RACDCERT command to define a certificate name filter (CNF), which maps a Distinguished Name to a user ID.

Perform the following steps to set up a CNF.

1. Enable CNF functions using the following command. You require update authority on the class DIGTNMAP to do this.

```
SETROPTS CLASSACT(DIGTNMAP) RACLIST(DIGTNMAP)
```

2. Define the CNF. For example:

```
RACDCERT ID(USER1) MAP WITHLABEL('filter1') TRUST  
SDNFILTER('O=IBM.C=UK') IDNFILTER('O=ExampleCA.L=Internet')
```

where USER1 is the user ID to be used when:

- The DN of the subject has an Organization of IBM and a Country of UK.
- The DN of the issuer has an Organization of ExampleCA and a Locality of Internet.

3. Refresh the CNF mappings:

```
SETROPTS RACLIST(DIGTNMAP) REFRESH
```

Note:

1. If the actual certificate is stored in the RACF database, the user ID under which it is installed is used in preference to the user ID associated with any CNF. If the certificate is not stored in the RACF database, the user ID associated with the most specific matching CNF is used. Matches of the subject DN are considered more specific than matches of the issuer DN.
2. Changes to CNFs do not apply until you refresh the CNF mappings.
3. A DN matches the DN filter in a CNF only if the DN filter is identical to the *least significant portion* of the DN. The least significant portion of the DN comprises the attributes that are usually listed at the right-most end of the DN, but which appear at the beginning of the certificate.
For example, consider the SDNFILTER 'O=IBM.C=UK'. A subject DN of 'CN=QM1.O=IBM.C=UK' matches that filter, but a subject DN of 'CN=QM1.O=IBM.L=Hursley.C=UK' does not match that filter.
The least significant portion of some certificates can contain fields that do not match the DN filter. Consider excluding these certificates by specifying a DN pattern in the SSLPEER pattern on the DEFINE CHANNEL command.
4. If the most specific matching CNF is defined to RACF as NOTRUST, the entity uses the user ID under which the channel initiator is running.
5. RACF uses the '.' character as a separator. WebSphere MQ uses either a comma or a semicolon.

You can define CNFs to ensure that the entity never sets the channel user ID to the default, which is the user ID under which the channel initiator is running. For each CA certificate in the key ring associated with the entity, define a CNF with an IDNFILTER that exactly matches the subject DN of that CA certificate. This ensures that all certificates that the entity might use match at least one of these CNFs. This is because all such certificates must either be connected to the key ring associated with the entity, or must be issued by a CA for which a certificate is connected to the key ring associated with the entity.

Refer to the *SecureWay Security Server RACF Security Administrator's Guide* for more information about the commands you use to manipulate CNFs.

Defining a sender channel and transmission queue on QMA:

Use the **DEFINE CHANNEL** and **DEFINE QLOCAL** commands to set up the required objects.

Procedure

On QMA, issue commands like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QMB.MACH.COM) XMITQ(QMB)
SSLCIPH(RC2_MD5_EXPORT) DESCR('Sender channel using SSL from QMA to QMB')

DEFINE QLOCAL(QMB) USAGE(XMITQ)
```

Results

A sender channel, TO.QMB, and a transmission queue, QMB, are created.

Defining a receiver channel on QMB:

Use the **DEFINE CHANNEL** command to set up the required object.

Procedure

On QMB, issue a command like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC2_MD5_EXPORT)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL to QMB')
```

Results

A receiver channel, TO.QMB, is created.

Starting the sender channel:

If necessary, start a listener program and refresh security. Then start the channel using the **START CHANNEL** command.

Procedure

1. Optional: If you have not already done so, start a listener program on QMB. The listener program listens for incoming network requests and starts the receiver channel when it is needed. For information about how to start a listener, see  Starting a channel listener (*WebSphere MQ V7.1 Installing Guide*).
2. Optional: If any SSL or TLS channels have run previously, issue the command `REFRESH SECURITY TYPE(SSL)`. This ensures that all the changes made to the key repository are available.
3. Start the channel on QMA, using the command `START CHANNEL(TO.QMB)`.

Results

The sender channel is started.

Exchanging self-signed certificates:

Exchange the certificates you previously extracted. If you use FTP, use the correct format.

Procedure

Transfer the CA part of the QM1 certificate to the QM2 system and vice versa, for example, by FTP.

If you transfer the certificates using FTP, you must do so in the correct format.

Transfer the following certificate types in *binary* format:

- DER encoded binary X.509
- PKCS #7 (CA certificates)
- PKCS #12 (personal certificates)

Transfer the following certificate types in ASCII format:

- PEM (privacy-enhanced mail)
- Base64 encoded X.509

Defining a sender channel and transmission queue on QM1:

Use the **DEFINE CHANNEL** and **DEFINE QLOCAL** commands to set up the required objects.

Procedure

On QM1, issue commands like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QM1.MACH.COM)
XMITQ(QM2) SSLCIPH(RC4_MD5_US) DESCR('Sender channel using SSL from QM1 to QM2')
DEFINE QLOCAL(QM2) USAGE(XMITQ)
```

This example uses CipherSpec RC4_MD5. The CipherSpecs at each end of the channel must be the same. Only the SSLCIPH parameter is mandatory if you want your channel to use SSL or TLS. See “CipherSpecs and CipherSuites in IBM WebSphere MQ” on page 399 for information about the permitted values for the SSLCIPH parameter.

Results

A sender channel, QM1.TO.QM2, and a transmission queue, QM2, are created.

Defining a receiver channel on QM2:

Use the **DEFINE CHANNEL** command to set up the required object.

Procedure

On QM2, issue a command like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC4_MD5_US)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from QM1 to QM2')
```

The channel must have the same name as the sender channel you defined in “Defining a sender channel and transmission queue on QM1,” and use the same CipherSpec.

Starting the sender channel:

If necessary, start a listener program and refresh security. Then start the channel using the **START CHANNEL** command.

Procedure

1. Optional: If you have not already done so, start a listener program on QM2. The listener program listens for incoming network requests and starts the receiver channel when it is needed. For information about how to start a listener, see  Starting a channel listener (*WebSphere MQ V7.1 Installing Guide*)
2. Optional: If any SSL or TLS channels have run previously, issue the command **REFRESH SECURITY TYPE(SSL)**. This ensures that all the changes made to the key repository are available.
3. On QM1, start the channel, using the command **START CHANNEL(QM1.TO.QM2)**.

Results

The sender channel is started.

Related concepts:

“Manipulating authentication information objects” on page 703

Related reference:

 **REFRESH SECURITY** (*WebSphere MQ V7.1 Reference*)

Refreshing the SSL or TLS environment:

Refresh the SSL or TLS environment on queue manager QMA using the **REFRESH SECURITY** command.

Procedure

On QMA, enter the following command:

```
REFRESH SECURITY TYPE(SSL)
```

This ensures that all the changes made to the key repository are available.

Allowing anonymous connections on a receiver channel:

Use the **ALTER CHANNEL** command to make SSL or TLS client authentication optional.

Procedure

On QMB, enter the following command:

```
ALTER CHANNEL(TO.QMB) CHLTYPE(RCVR) SSLCAUTH(OPTIONAL)
```

Starting the sender channel:

If necessary, start the channel initiator, start a listener program, and refresh security. Then start the channel using the **START CHANNEL** command.

Procedure

1. Optional: if you have not already done so, start the channel initiator.
2. Optional: If you have not already done so, start a listener program on QM2. The listener program listens for incoming network requests and starts the receiver channel when it is needed. For information about how to start a listener, see  Starting a channel listener (*WebSphere MQ V7.1 Installing Guide*)
3. Optional: If the channel initiator was already running or any SSL or TLS channels have run previously, issue the command `REFRESH SECURITY TYPE(SSL)`. This ensures that all the changes made to the key repository are available.
4. On QM1, start the channel, using the command `START CHANNEL(QM1.TO.QM2)`.

Results

The sender channel is started.

Starting the sender channel on z/OS:

If necessary, start the channel initiator, start a listener program, and refresh security. Then start the channel using the **START CHANNEL** command.

Procedure

1. Optional: If you have not already done so, start the channel initiator.
2. Optional: If you have not already done so, start a listener program on QMB. The listener program listens for incoming network requests and starts the receiver channel when it is needed. For information about how to start a listener, see  Starting a channel listener (*WebSphere MQ V7.1 Installing Guide*).
3. Optional: If the channel initiator was already running or if any SSL or TLS channels have run previously, issue the command `REFRESH SECURITY TYPE(SSL)`. This ensures that all the changes made to the key repository are available.
4. Start the channel on QMA, using the command `START CHANNEL(TO.QMB)`.

Results

The sender channel is started.

Configuring IBM WebSphere MQ for Common Criteria

When you use IBM WebSphere MQ in Common Criteria mode, you must comply with the following guidance.

- Use certificates that pass the PKIX validation specified in RFC 5280. You must also configure IBM WebSphere MQ to enforce RFC 5280 compliant certificate validation as described in this topic.
- Always use the `runmqakm` command to manage certificates and keys. Do not use `iKeyman` or `iKeycmd`.
- The IBM WebSphere MQ key database file is protected by a password. To allow unattended access to the key database file IBM WebSphere MQ provides a stash file to store the password. This stash file must be protected by an ACL or permission bits while it resides on the system and by encrypting the stash file when it is backed up.

- Ensure that key database files are protected by a randomly chosen strong password. At a minimum, ensure that the password adheres to the following rules:
 - The password must be a minimum length of 14 characters.
 - The password must contain a minimum of one lowercase character, one uppercase character, and one digit or special character.
 - Each character can occur a maximum of three times in a password.
 - A maximum of two consecutive characters in the password can be identical.

The **runmqakm** command can generate a strong random password that complies with these criteria. For more information, see “Generating strong passwords for key repository protection” on page 654.

- Use key database files that were created using strong encryption. That is, you must specify the parameters **strong** and **FIPS** when you use **runmqakm** command to create key database files.
- When you change the password for a key database file using the **runmqakm** command, you must specify the parameters **strong** and **FIPS**.
- When you convert an existing key database file in the old format to the newer secure mode using the **runmqakm** command, you must specify the parameters **strong** and **FIPS**.
- When you create a certificate request, a self-signed certificate, or sign a certificate using the **runmqakm** command, you must specify a **sig_alg** parameter of either **sha1** or a stronger digital signature algorithm. It is recommended to use one of the SHA-2 based digital signature algorithms such as **sha256**.
- If you use the **runmqakm** command to export a certificate, do not use the **target** parameter to automatically create a target key database file. You must have already created a target key database file by running the **runmqakm** command in **FIPS** mode with strong encryption.
- If you use the **runmqakm** command to import a certificate, the **FIPS** parameter must be specified on the import command.
- Do not configure IBM WebSphere MQ to use hardware key database files or hardware cipher acceleration. That is, you must not use the **crypto**, **tokenlabel**, **secondaryDB**, or **secondaryDBpw** parameters in the **runmqakm** tool.
- Do not use a value of 512 for the **size** parameter when using the **runmqakm** tool to generate certificates or certificate signing requests with RSA public keys. The value of the **size** parameter for RSA digital certificates must be a minimum of 1024, although 2048 is recommended.
- After making any changes to the certificates and keys in the queue manager's key repository, run the **REFRESH SECURITY TYPE(SSL) MQSC** command to ensure that the queue manager is using the latest key repository.
- Allow only certificates with a Basic Constraints extension present to participate in a certificate chain. By default, when validating a certificate chain, IBM WebSphere MQ operates such that certificates with no Basic Constraints extension present are allowed to sign other certificates in a certificate chain. In order to operate WebSphere MQ in Common Criteria mode, certificates with no Basic Constraints extension present must not be used in this way. This requirement is automatically enforced when IBM WebSphere MQ is configured to use the RFC 5280 compliant certificate validation policy.
- When tracing is switched on for problem determination, be aware that IBM WebSphere MQ is not in Common Criteria mode.
- If you have applications that use IBM WebSphere MQ and that also have other trusted and untrusted applications on the same system, run applications to use the maximum operating system protection and minimize the risk to privileged applications from unprivileged user IDs.
- An application running under a privileged user, like root or a Windows administrator, might have access to process environments containing TSF data. To prevent untrusted applications harming your system, avoid this privileged mode of operation.
- If you use certificate revocation lists (CRLs) for certificate validation, the IBM WebSphere MQ administrator is responsible for providing current and correct CRLs. The IBM WebSphere MQ administrator must ensure that the CRL provided to IBM WebSphere MQ includes all revoked certificates from any certificate authorities (CAs) recognized by IBM WebSphere MQ.

CRLs can be retrieved using an LDAP server in the IBM WebSphere MQ Common Criteria environment. This retrieval requires an LDAP server provided or accessed by the IBM WebSphere MQ administrator. In Common Criteria mode, the administrator must configure only one LDAP CRL server to IBM WebSphere MQ.

WebSphere MQ establishes a TCP/IP connection to the LDAP server and retrieves the CRL. Finally, this CRL is used for certificate validation. If no CRL can be obtained (for example due to LDAP failure), the certificate to be validated is considered invalid.

Ensure the connection between IBM WebSphere MQ and the LDAP server is an internal communication link within a trusted network, because this connection is made over an unprotected TCP/IP connection.

- If an OCSP responder server is used for certificate validation, the IBM WebSphere MQ administrator is responsible for providing the correct URL to a trustworthy OCSP responder. The administrator must also add the CA certificates necessary to verify the digital signature of any OCSP responses into the key repository, after checking that each CA is trustworthy.
- No environment variable whose name begins with `AMQ_SSL_ALLOW_DEFAULT` can be set.

Configuration requirements

To configure a queue manager for Common Criteria compliant operation, you must perform the configuration steps described in this section. Note that you must perform these steps on all queue managers where you require Common Criteria compliance.

Firstly, you must increase the command level of the queue manager to 711 or higher. IBM WebSphere MQ fixpack 7.1.0.2 is the minimum version that supports command level 711.

However, all queue managers created with WebSphere MQ 7.1 have a default command level of 710, even if the queue manager was created with fixpack 7.1.0.2 or later.

Note: Increasing the command level is an irreversible operation that prevents the queue manager running at IBM WebSphere MQ versions prior to 7.1.0.2; see  [Migrating queue managers to new-function fix packs \(WebSphere MQ V7.1 Installing Guide\)](#) for more information about queue manager command level migration.

For example, to increase the command level of a queue manager named `QMGR`, ensure that the queue manager is not running and then process the following commands:

```
strmqm -e CMDLEVEL=711 QMGR
strmqm QMGR
```

The first `strmqm` command increases the command level, and the second `strmqm` starts the queue manager with the new command level. Once you have increased the queue manager command level, you must run the MQSC commands described in this section.

For example, on a queue manager named `QMGR` you must first run the following command:

```
runmqsc QMGR
```

For more information about the `runmqsc` command, see  [runmqsc \(WebSphere MQ V7.1 Reference\)](#). You can then enter each MQSC command as shown below.

In order that auditing of authority, channel, command, and configuration events is implemented, execute the following MQSC command:

```
ALTER QMGR AUTHOREV(ENABLED) CHLEV(ENABLED) SSLEV(ENABLED) CMDEV(ENABLED) CONFIGEV(ENABLED)
```

If any of these settings are disabled, auditing is no longer performed, and IBM WebSphere MQ does not operate in accordance with the evaluated configuration.

You must configure the relevant event queues to have a maximum depth of 999999999, the maximum permitted by IBM WebSphere MQ. To do this, issue the following MQSC commands:

```
ALTER QLOCAL(SYSTEM.ADMIN.CHANNEL.EVENT) MAXDEPTH(999999999)
ALTER QLOCAL(SYSTEM.ADMIN.COMMAND.EVENT) MAXDEPTH(999999999)
ALTER QLOCAL(SYSTEM.ADMIN.CONFIG.EVENT) MAXDEPTH(999999999)
ALTER QLOCAL(SYSTEM.ADMIN.QMGR.EVENT) MAXDEPTH(999999999)
```

You must also ensure that an event monitoring program is running at all times to remove messages from each event queue, and process messages from each event queue. This is necessary to ensure that the event queues never become full, and so prevent the loss of audit messages.

In order that the queue manager enforces FIPS 140-2 and RFC 5280 compliance, execute the following MQSC command:

```
ALTER QMGR SSLFIPS(YES) CERTVPOL(RFC5280)
```

Note: In order to set the certificate validation policy (CERTVPOL) queue manager attribute, the queue manager must be at command level 711 or higher. If the CERTVPOL command causes error AMQ8894, you must end the queue manager and increase the command level to 711, as described previously in this section.

For more information about FIPS 140-2 configuration on the client, see “Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client” on page 614. For more information about RFC 5280 configuration on the client, see “Configuring certificate validation policies in WebSphere MQ” on page 403.

Enable the use of channel authentication records by executing the following MQSC command:

```
ALTER QMGR CHLAUTH(ENABLED)
```

You must remove the default channel authentication rule which allows use of the SYSTEM.ADMIN.SVRCONN channel and define a rule to block all inbound connections by executing the following MQSC commands:

```
SET CHLAUTH(SYSTEM.ADMIN.SVRCONN) TYPE(ADDRESSMAP) ACTION(REMOVEALL)
SET CHLAUTH(*) TYPE(ADDRESSMAP) ADDRESS(*) USERSRC(NOACCESS) ACTION(REPLACE)
```

Ensure that all channels use TLS secure sockets with mutual authentication by selecting appropriate FIPS 140-2 compliant CipherSpecs for the SSLCIPH channel attribute and setting the SSLCAUTH attribute to REQUIRED. You can check which channels are using TLS mutual authentication with the following command:

```
DISPLAY CHANNEL(*) SSLCIPH SSLCAUTH
```

Correctly configured channels show SSLCAUTH(REQUIRED) and an SSLCIPH value which is listed as FIPS compliant in “Specifying CipherSpecs” on page 774.

For each client application or remote queue manager which needs to connect, you must define an SSLPEERMAP channel authentication rule which maps the remote certificate Subject Distinguished Name onto an MCAUSER user ID specified in that channel authentication record. This user ID must have the WebSphere MQ authorities to access the queues, topics, and other objects it needs.

For example, to map the digital certificate whose Subject DN is "CN=Fred Bloggs, OU=Sales, O=Example, C=US" onto user ID fred, you would issue an MQSC command such as:

```
SET CHLAUTH(*) TYPE(SSLPEERMAP) SSLPEER('CN=Fred Bloggs, OU=Sales, O=Example, C=US')
USERSRC(MAP) MCAUSER('fred')
```

In a Common Criteria compliant configuration, all remote connections must send a personal certificate and the user ID used for WebSphere MQ object authority checks must be mapped onto the MCAUSER value in an SSLPEERMAP channel authentication record. This mapping prevents impersonation of authorized users.

To confirm whether the required settings are enabled, execute the following MQSC command:

```
DISPLAY QMGR
```

Obtaining the latest information

Always refer to the IBM website for the latest IBM WebSphere MQ support information and for the latest versions of the IBM WebSphere MQ documentation. For more information, see  Finding the latest information (*WebSphere MQ V7.1 Installing Guide*).

Related concepts:

“Common Criteria environmental considerations” on page 445

“Common Criteria in WebSphere MQ” on page 384

Identifying and authenticating users

You can identify and authenticate users by using the MQCSP structure or in several types of user exit program.

Using the MQCSP structure

You specify the MQCSP connection security parameters structure on an MQCONN call; this structure contains a user ID and password. If necessary, you can alter the MQCSP in a security exit.

Note: The object authority manager (OAM) does not use the password. However the OAM does some limited work with the user ID, that could be considered a trivial form of authentication. These checks stop you adopting another user ID, if you use those parameters in your applications.

Implementing identification and authentication in security exits

The primary purpose of a security exit is to enable the MCA at each end of a channel to authenticate its partner. At each end of a message channel, and at the server end of an MQI channel, an MCA typically acts on behalf of the queue manager to which it is connected. At the client end of an MQI channel, an MCA typically acts on behalf of the user of the WebSphere MQ client application. In this situation, mutual authentication actually takes place between two queue managers, or between a queue manager and the user of a WebSphere MQ MQI client application.

The supplied security exit (the SSPI channel exit) illustrates how mutual authentication can be implemented by exchanging authentication tokens that are generated, and then checked, by a trusted authentication server such as Kerberos. For more details, see “The SSPI channel exit program” on page 487.

Mutual authentication can also be implemented by using Public Key Infrastructure (PKI) technology. Each security exit generates some random data, signs it using the private key of the queue manager or user it is representing, and sends the signed data to its partner in a security message. The partner security exit performs the authentication by checking the digital signature using the public key of the queue manager or user. Before exchanging digital signatures, the security exits might need to agree the algorithm for generating a message digest, if more than one algorithm is available for use.

When a security exit sends the signed data to its partner, it also needs to send some means of identifying the queue manager or user it is representing. This might be a Distinguished Name, or even a digital

certificate. If a digital certificate is sent, the partner security exit can validate the certificate by working through the certificate chain to the root CA certificate. This provides assurance of the ownership of the public key that is used to check the digital signature.

The partner security exit can validate a digital certificate only if it has access to a key repository that contains the remaining certificates in the certificate chain. If a digital certificate for the queue manager or user is not sent, one must be available in the key repository to which the partner security exit has access. The partner security exit cannot check the digital signature unless it can find the signer's public key.

The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) use PKI techniques like the ones just described. For more information about how SSL and TLS perform authentication, see "Secure Sockets Layer (SSL) and Transport Layer Security (TLS) concepts" on page 376.

If a trusted authentication server or PKI support is not available, other techniques can be used. A common technique, which can be implemented in security exits, uses a symmetric key algorithm.

One of the security exits, exit A, generates a random number and sends it in a security message to its partner security exit, exit B. Exit B encrypts the number using its copy of a key which is known only to the two security exits. Exit B sends the encrypted number to exit A in a security message with a second random number that exit B has generated. Exit A verifies that the first random number has been encrypted correctly, encrypts the second random number using its copy of the key, and sends the encrypted number to exit B in a security message. Exit B then verifies that the second random number has been encrypted correctly. During this exchange, if either security exit is not satisfied with the authenticity of other, it can instruct the MCA to close the channel.

An advantage of this technique is that no key or password is sent over the communications connection during the exchange. A disadvantage is that it does not provide a solution to the problem of how to distribute the shared key in a secure way. One solution to this problem is described in "Implementing confidentiality in user exit programs" on page 780. A similar technique is used in SNA for the mutual authentication of two LUs when they bind to form a session. The technique is described in "Session level authentication" on page 454.

All the preceding techniques for mutual authentication can be adapted to provide one-way authentication.

Implementing identification and authentication in message exits

When an application puts a message on a queue, the *UserIdentifier* field in the message descriptor contains a user ID associated with the application. However, there is no data present that can be used to authenticate the user ID. This data can be added by a message exit at the sending end of a channel and checked by a message exit at the receiving end of the channel. The authenticating data can be an encrypted password or a digital signature, for example.

This service might be more effective if it is implemented at the application level. The basic requirement is for the user of the application that receives the message to be able to identify and authenticate the user of the application that sent the message. It is therefore natural to consider implementing this service at the application level. For more information, see "Identity mapping in the API exit and API-crossing exit" on page 694.

Implementing identification and authentication in the API exit and API-crossing exit

At the level of an individual message, identification and authentication is a service that involves two users, the sender and the receiver of the message. The basic requirement is for the user of the application that receives the message to be able to identify and authenticate the user of the application that sent the message. Note that the requirement is for one way, not two way, authentication.

Depending on how it is implemented, the users and their applications might need to interface, or even interact, with the service. In addition, when and how the service is used might depend on where the users and their applications are located, and on the nature of the applications themselves. It is therefore natural to consider implementing the service at the application level rather than at the link level.

If you consider implementing this service at the link level, you might need to resolve issues such as the following:

- On a message channel, how do you apply the service only to those messages that require it?
- How do you enable users and their applications to interface, or interact, with the service, if this is a requirement?
- In a multi-hop situation, where a message is sent over more than one message channel on the way to its destination, where do you invoke the components of the service?

Here are some examples of how the identification and authentication service can be implemented at the application level. The term *API exit* means either an API exit or an API-crossing exit.

- When an application puts a message on a queue, an API exit can acquire an authentication token from a trusted authentication server such as Kerberos. The API exit can add this token to the application data in the message. When the message is retrieved by the receiving application, a second API exit can ask the authentication server to authenticate the sender by checking the token.
- When an application puts a message on a queue, an API exit can append the following items to the application data in the message:
 - The digital certificate of the sender
 - The digital signature of the sender

If different algorithms for generating a message digest are available for use, the API exit can include the name of the algorithm it has used.

When the message is retrieved by the receiving application, a second API exit can perform the following checks:

- The API exit can validate the digital certificate by working through the certificate chain to the root CA certificate. To do this, the API exit must have access to a key repository that contains the remaining certificates in the certificate chain. This check provides assurance that the sender, identified by the Distinguished Name, is the genuine owner of the public key contained in the certificate.
- The API exit can check the digital signature using the public key contained in the certificate. This check authenticates the sender.

The Distinguished Name of the sender can be sent instead of the whole digital certificate. In this case, the key repository must contain the sender's certificate so that the second API exit can find the public key of the sender. Another possibility is to send all the certificates in the certificate chain.

- When an application puts a message on a queue, the *UserIdentifier* field in the message descriptor contains a user ID associated with the application. The user ID can be used to identify the sender. To enable authentication, an API exit can append some data, such as an encrypted password, to the application data in the message. When the message is retrieved by the receiving application, a second API exit can authenticate the user ID by using the data that has travelled with the message.

This technique might be considered sufficient for messages that originate in a controlled and trusted environment, and in circumstances where a trusted authentication server or PKI support is not available.

Privileged users

A privileged user is one that has full administrative authorities for WebSphere MQ.

In addition to the users listed in the following table, members of any group with +crt authority for queues are indirectly administrators. Similarly, any user that has +set authority on the queue manager, and +put authority on the command queue is an administrator.

You should not grant these privileges to ordinary users and applications.

Table 83. Privileged users by platform

Platform	Privileged users
Windows systems	<ul style="list-style-type: none">• SYSTEM• Members of the mqm group• Members of the Administrators group
UNIX and Linux systems	<ul style="list-style-type: none">• Members of the mqm group
IBM i systems	<ul style="list-style-type: none">• The profiles qmqm and qmqmadm• All members of the qmqmadm group• Any user defined with the *ALLOBJ setting
z/OS	The user ID that the channel initiator, queue manager, and advanced message security address spaces are running under.

Identifying and authenticating users using the MQCSP structure

You specify the MQCSP connection security parameters structure on an MQCONN call. This structure contains a user ID and password, which the Authorization Service can use to identify and authenticate the user.

If necessary, you can alter the MQCSP in a security exit.

Implementing identification and authentication in security exits

You can use a security exit to implement one-way or mutual authentication.

The primary purpose of a security exit is to enable the MCA at each end of a channel to authenticate its partner. At each end of a message channel, and at the server end of an MQI channel, an MCA typically acts on behalf of the queue manager to which it is connected. At the client end of an MQI channel, an MCA typically acts on behalf of the user of the WebSphere MQ MQI client application. In this situation, mutual authentication actually takes place between two queue managers, or between a queue manager and the user of a WebSphere MQ MQI client application.

The supplied security exit (the SSPI channel exit) illustrates how mutual authentication can be implemented by exchanging authentication tokens that are generated, and then checked, by a trusted authentication server such as Kerberos. For more details, see “The SSPI channel exit program” on page 487.

Mutual authentication can also be implemented by using Public Key Infrastructure (PKI) technology. Each security exit generates some random data, signs it using the private key of the queue manager or user it is representing, and sends the signed data to its partner in a security message. The partner security exit performs the authentication by checking the digital signature using the public key of the queue manager or user. Before exchanging digital signatures, the security exits might need to agree the algorithm for generating a message digest, if more than one algorithm is available for use.

When a security exit sends the signed data to its partner, it also needs to send some means of identifying the queue manager or user it is representing. This might be a Distinguished Name, or even a digital certificate. If a digital certificate is sent, the partner security exit can validate the certificate by working through the certificate chain to the root CA certificate. This provides assurance of the ownership of the public key that is used to check the digital signature.

The partner security exit can validate a digital certificate only if it has access to a key repository that contains the remaining certificates in the certificate chain. If a digital certificate for the queue manager or user is not sent, one must be available in the key repository to which the partner security exit has access. The partner security exit cannot check the digital signature unless it can find the signer's public key.

The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) use PKI techniques like the ones just described. For more information about how the Secure Sockets Layer performs authentication, see "Secure Sockets Layer (SSL) and Transport Layer Security (TLS) concepts" on page 376.

If a trusted authentication server or PKI support is not available, other techniques can be used. A common technique, which can be implemented in security exits, uses a symmetric key algorithm.

One of the security exits, exit A, generates a random number and sends it in a security message to its partner security exit, exit B. Exit B encrypts the number using its copy of a key which is known only to the two security exits. Exit B sends the encrypted number to exit A in a security message with a second random number that exit B has generated. Exit A verifies that the first random number has been encrypted correctly, encrypts the second random number using its copy of the key, and sends the encrypted number to exit B in a security message. Exit B then verifies that the second random number has been encrypted correctly. During this exchange, if either security exit is not satisfied with the authenticity of other, it can instruct the MCA to close the channel.

An advantage of this technique is that no key or password is sent over the communications connection during the exchange. A disadvantage is that it does not provide a solution to the problem of how to distribute the shared key in a secure way. One solution to this problem is described in "Implementing confidentiality in user exit programs" on page 780. A similar technique is used in SNA for the mutual authentication of two LUs when they bind to form a session. The technique is described in "Session level authentication" on page 454.

All the preceding techniques for mutual authentication can be adapted to provide one-way authentication.

Identity mapping in message exits

You can use message exits to process information to authenticate a user ID, though it might be better to implement authentication at the application level.

When an application puts a message on a queue, the *UserIdentifier* field in the message descriptor contains a user ID associated with the application. However, there is no data present that can be used to authenticate the user ID. This data can be added by a message exit at the sending end of a channel and checked by a message exit at the receiving end of the channel. The authenticating data can be an encrypted password or a digital signature, for example.

This service might be more effective if it is implemented at the application level. The basic requirement is for the user of the application that receives the message to be able to identify and authenticate the user of the application that sent the message. It is therefore natural to consider implementing this service at the application level. For more information, see "Identity mapping in the API exit and API-crossing exit" on page 694.

Identity mapping in the API exit and API-crossing exit

An application that receives a message must be able to identify and authenticate the user of the application that sent the message. This service is typically best implemented at the application level. API exits can implement the service in a number of ways.

At the level of an individual message, identification and authentication is a service that involves two users, the sender and the receiver of the message. The basic requirement is for the user of the application that receives the message to be able to identify and authenticate the user of the application that sent the message. Note that the requirement is for one way, not two way, authentication.

Depending on how it is implemented, the users and their applications might need to interface, or even interact, with the service. In addition, when and how the service is used might depend on where the users and their applications are located, and on the nature of the applications themselves. It is therefore natural to consider implementing the service at the application level rather than at the link level.

If you consider implementing this service at the link level, you might need to resolve issues such as the following:

- On a message channel, how do you apply the service only to those messages that require it?
- How do you enable users and their applications to interface, or interact, with the service, if this is a requirement?
- In a multi-hop situation, where a message is sent over more than one message channel on the way to its destination, where do you invoke the components of the service?

Here are some examples of how the identification and authentication service can be implemented at the application level. The term *API exit* means either an API exit or an API-crossing exit.

- When an application puts a message on a queue, an API exit can acquire an authentication token from a trusted authentication server such as Kerberos. The API exit can add this token to the application data in the message. When the message is retrieved by the receiving application, a second API exit can ask the authentication server to authenticate the sender by checking the token.
- When an application puts a message on a queue, an API exit can append the following items to the application data in the message:
 - The digital certificate of the sender
 - The digital signature of the sender

If different algorithms for generating a message digest are available for use, the API exit can include the name of the algorithm it has used.

When the message is retrieved by the receiving application, a second API exit can perform the following checks:

- The API exit can validate the digital certificate by working through the certificate chain to the root CA certificate. To do this, the API exit must have access to a key repository that contains the remaining certificates in the certificate chain. This check provide assurance that the sender, identified by the Distinguished Name, is the genuine owner of the public key contained in the certificate.
- The API exit can check the digital signature using the public key contained in the certificate. This check authenticates the sender.

The Distinguished Name of the sender can be sent instead of the whole digital certificate. In this case, the key repository must contain the sender's certificate so that the second API exit can find the public key of the sender. Another possibility is to send all the certificates in the certificate chain.

- When an application puts a message on a queue, the *UserIdentifier* field in the message descriptor contains a user ID associated with the application. The user ID can be used to identify the sender. To enable authentication, an API exit can append some data, such as an encrypted password, to the application data in the message. When the message is retrieved by the receiving application, a second API exit can authenticate the user ID by using the data that has travelled with the message.

This technique might be considered sufficient for messages that originate in a controlled and trusted environment, and in circumstances where a trusted authentication server or PKI support is not available.

Working with revoked certificates

Digital certificates can be revoked by Certificate Authorities. You can check the revocation status of certificates using OCSP, or CRLs on LDAP servers, depending on platform.

During the SSL handshake, the communicating partners authenticate each other with digital certificates. Authentication can include a check that the certificate received can still be trusted. Certificate Authorities (CAs) revoke certificates for various reasons, including:

- The owner has moved to a different organization
- The private key is no longer secret

CAs publish revoked personal certificates in a Certificate Revocation List (CRL). CA certificates that have been revoked are published in an Authority Revocation List (ARL).

On UNIX, Linux and Windows systems, WebSphere MQ SSL support checks for revoked certificates using OCSP (Online Certificate Status Protocol) or using CRLs and ARLs on LDAP (Lightweight Directory Access Protocol) servers. OCSP is the preferred method. However, WebSphere MQ classes for Java and WebSphere MQ classes for JMS cannot use OCSP.

On z/Os and IBM i WebSphere MQ SSL support checks for revoked certificates using CRLs and ARLs on LDAP servers only.

For more information about Certificate Authorities, see “Digital certificates” on page 371.

Revoked certificates and OCSP

WebSphere MQ determines which Online Certificate Status Protocol (OCSP) responder to use, and handles the response received. You might have to take steps to make the OCSP responder accessible.

Note: This information applies only to WebSphere MQ on Windows, UNIX and Linux systems.

To check the revocation status of a digital certificate using OCSP, WebSphere MQ can use two methods to determine which OCSP responder to contact:

- By using the AuthorityInfoAccess (AIA) certificate extension in the certificate to be checked.
- By using a URL specified in an authentication information object or specified by a client application.

A URL specified in an authentication information object or by a client application takes priority over a URL in an AIA certificate extension.

If the URL of the OCSP responder lies behind a firewall, reconfigure the firewall so the OCSP responder can be accessed or set up an OCSP proxy server. Specify the name of the proxy server by using the `SSLHTTPProxyName` variable in the SSL stanza. On client systems, you can also specify the name of the proxy server by using the environment variable `MQSSLPROXY`. For more details, see the related information.

If you are not concerned whether TLS or SSL certificates are revoked, perhaps because you are running in a test environment, you can set `OCSPCheckExtensions` to `NO` in the SSL stanza. If you set this variable, any AIA certificate extension is ignored. This solution is unlikely to be acceptable in a production environment, where you probably do not want to allow access from users presenting revoked certificates.

The call to access the OCSP responder can result in one of the following three outcomes:

Good The certificate is valid.

Revoked

The certificate is revoked.

Unknown

This outcome can arise for one of three reasons:

- WebSphere MQ cannot access the OCSP responder.
- The OCSP responder has sent a response, but WebSphere MQ cannot verify the digital signature of the response.
- The OCSP responder has sent a response that indicates that it has no revocation data for the certificate.

If WebSphere MQ receives an OCSP outcome of Unknown, its behavior depends on the setting of the OCSPAuthentication attribute. For queue managers, this attribute is held in the SSL stanza of the qm.ini file for UNIX and Linux systems, or the Windows registry. It can be set using the WebSphere MQ Explorer. For clients, it is held in the SSL stanza of the client configuration file.

If an outcome of Unknown is received and OCSPAuthentication is set to REQUIRED (the default value), WebSphere MQ rejects the connection and issues an error message of type AMQ9716. If queue manager SSL event messages are enabled, an SSL event message of type MQRC_CHANNEL_SSL_ERROR with ReasonQualifier set to MQRC_SSL_HANDSHAKE_ERROR is generated.

If an outcome of Unknown is received and OCSPAuthentication is set to OPTIONAL, WebSphere MQ allows the SSL channel to start and no warnings or SSL event messages are generated.

If an outcome of Unknown is received and OCSPAuthentication is set to WARN, the SSL channel starts but WebSphere MQ issues a warning message of type AMQ9717 in the error log. If queue manager SSL event messages are enabled, an SSL event message of type MQRC_CHANNEL_SSL_WARNING with ReasonQualifier set to MQRC_SSL_UNKNOWN_REVOCATION is generated.

Digital signing of OCSP responses

An OCSP responder can sign its responses in one of three ways. Your responder will inform you which method is used.

- The OCSP response can be digitally signed using the same CA certificate that issued the certificate that you are checking. In this case, you do not need to set up any additional certificate; the steps you have already taken to establish SSL connectivity are sufficient to verify the OCSP response.
- The OCSP response can be digitally signed using another certificate signed by the same certificate authority (CA) that issued the certificate you are checking. The signing certificate is sent together with the OCSP response in this case. The certificate flowed from the OCSP responder must have an Extended Key Usage Extension set to id-kp-OCSPSigning so that it can be trusted for this purpose. Because the OCSP response is sent with the certificate which signed it (and that certificate is signed by a CA that is already trusted for SSL connectivity), no additional certificate setup is required.
- The OCSP response can be digitally signed using another certificate that is not directly related to the certificate you are checking. In this case, the OCSP response is signed by a certificate issued by the OCSP responder itself. You must add a copy of the OCSP responder certificate to the key database of the client or queue manager which performs the OCSP checking; see “Adding a CA certificate (or the public part of a self-signed certificate) into a key repository, on UNIX, Linux or Windows systems” on page 648. When a CA certificate is added, by default it is added as a trusted root, which is the required setting in this context. If this certificate is not added, WebSphere MQ cannot verify the digital signature on the OCSP response and the OCSP check results in an Unknown outcome, which might cause WebSphere MQ to close the channel, depending on the value of OCSPAuthentication.

Related concepts:

“Manipulating authentication information objects” on page 703

“Location of an OCSP responder, and of LDAP servers that hold CRLs” on page 701

Related reference:



SSL stanza of the client configuration file (*WebSphere MQ V7.1 Installing Guide*)



MQSSLPROXY (*WebSphere MQ V7.1 Installing Guide*)

Working with Certificate Revocation Lists and Authority Revocation Lists

WebSphere MQ's support for CRLs and ARLs varies by platform.

CRL and ARL support on each platform is as follows:

- On z/OS, System SSL supports CRLs and ARLs stored in LDAP servers by the Tivoli Public Key Infrastructure product.
- On other platforms, the CRL and ARL support complies with PKIX X.509 V2 CRL profile recommendations.

WebSphere MQ maintains a cache of CRLs and ARLs that have been accessed in the preceding 12 hours.

When a queue manager or WebSphere MQ MQI client receives a certificate, it checks the CRL to confirm that the certificate is still valid. WebSphere MQ first checks in the cache, if there is a cache. If the CRL is not in the cache, WebSphere MQ interrogates the LDAP CRL server locations in the order they occur in the namelist of authentication information objects specified by the `SSLCRLNamelist` attribute, until WebSphere MQ finds an available CRL. If the namelist is not specified, or is specified with a blank value, CRLs are not checked.

For more information about LDAP, see  Using lightweight directory access protocol services with WebSphere MQ for Windows (*WebSphere MQ V7.1 Programming Guide*).

Setting up LDAP servers:

Configure the LDAP Directory Information Tree structure to reflect the hierarchy of Distinguished Names of CAs. Do this using LDAP Data Interchange Format files.

Configure the LDAP Directory Information Tree (DIT) structure to use the hierarchy corresponding to the Distinguished Names of the CAs that issue the certificates and CRLs. You can set up the DIT structure with a file that uses the LDAP Data Interchange Format (LDIF). You can also use LDIF files to update a directory.

LDIF files are ASCII text files that contain the information required to define objects within an LDAP directory. LDIF files contain one or more entries, each of which comprises a Distinguished Name, at least one object class definition and, optionally, multiple attribute definitions.

The `certificateRevocationList;binary` attribute contains a list, in binary form, of revoked user certificates. The `authorityRevocationList;binary` attribute contains a binary list of CA certificates that have been revoked. For use with WebSphere MQ SSL, the binary data for these attributes must conform to DER (Definite Encoding Rules) format. For more information about LDIF files, refer to the documentation provided with your LDAP server.

Figure 87 on page 698 shows a sample LDIF file that you might create as input to your LDAP server to load the CRLs and ARLs issued by CA1, which is an imaginary Certificate Authority with the Distinguished Name “CN=CA1, OU=Test, O=IBM, C=GB”, set up by the Test organization within IBM.

```

dn: o=IBM, c=GB
o: IBM
objectclass: top
objectclass: organization

dn: ou=Test, o=IBM, c=GB
ou: Test
objectclass: organizationalUnit

dn: cn=CA1, ou=Test, o=IBM, c=GB
cn: CA1
objectclass: cRLDistributionPoint
objectclass: certificateAuthority
authorityRevocationList;binary:: (DER format data)
certificateRevocationList;binary:: (DER format data)
caCertificate;binary:: (DER format data)

```

Figure 87. Sample LDIF file for a Certificate Authority. This might vary from implementation to implementation.

Figure 88 shows the DIT structure that your LDAP server creates when you load the sample LDIF file shown in Figure 87 together with a similar file for CA2, an imaginary Certificate Authority set up by the PKI organization, also within IBM.

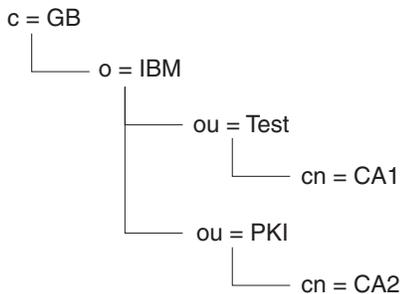


Figure 88. Example of an LDAP Directory Information Tree structure

WebSphere MQ checks both CRLs and ARLs.

Note: Ensure that the access control list for your LDAP server allows authorized users to read, search, and compare the entries that hold the CRLs and ARLs. WebSphere MQ accesses the LDAP server using the LDAPUSER and LDAPPWD properties of the AUTHINFO object.

Configuring and updating LDAP servers:

Use this procedure to configure or update your LDAP server.

1. Obtain the CRLs and ARLs in DER format from your Certification Authority, or Authorities.
2. Using a text editor or the tool provided with your LDAP server, create one or more LDIF files that contain the Distinguished Name of the CA and the required object class definitions. Copy the DER format data into the LDIF file as the values of either the `certificateRevocationList;binary` attribute for CRLs, the `authorityRevocationList;binary` attribute for ARLs, or both.
3. Start your LDAP server.
4. Add the entries from the LDIF file or files you created at step 2.

After you have configured your LDAP CRL server, check that it is set up correctly. First, try using a certificate that is not revoked on the channel, and check that the channel starts correctly. Then use a certificate that is revoked, and check that the channel fails to start.

Obtain updated CRLs from the Certification Authorities frequently. Consider doing this on your LDAP servers every 12 hours.

Accessing CRLs and ARLs with a queue manager:

A queue manager is associated with one or more authentication information objects, which hold the address of an LDAP CRL server. Websphere MQ on IBM i behaves differently from other platforms.

Note that in this section, information about Certificate Revocation Lists (CRLs) also applies to Authority Revocation Lists (ARLs).

You tell the queue manager how to access CRLs by supplying the queue manager with authentication information objects, each of which holds the address of an LDAP CRL server. The authentication information objects are held in a namelist, which is specified in the *SSLCRLNamelist* queue manager attribute.

In the following example, MQSC is used to specify the parameters:

1. Define authentication information objects using the DEFINE AUTHINFO MQSC command, with the AUTHTYPE parameter set to CRLLDAP. On IBM i, you can also use the CRTMQMAUTI CL command.

The value CRLLDAP for the AUTHTYPE parameter indicates that CRLs are accessed on LDAP servers. Each authentication information object with type CRLLDAP that you create holds the address of an LDAP server. When you have more than one authentication information object, the LDAP servers to which they point *must* contain identical information. This provides continuity of service if one or more LDAP servers fail.

Additionally, on z/OS only, all LDAP servers must be accessed using the same user ID and password. The user ID and password used are those specified in the first AUTHINFO object in the namelist.

On all platforms, the user ID and password are sent to the LDAP server unencrypted.

2. Using the DEFINE NAMELIST MQSC command, define a namelist for the names of your authentication information objects. On z/OS, ensure that the NLTYPE namelist attribute is set to AUTHINFO.
3. Using the ALTER QMGR MQSC command, supply the namelist to the queue manager. For example:
ALTER QMGR SSLCRLNL(sslcrlnlname)

where sslcrlnlname is your namelist of authentication information objects.

This command sets a queue manager attribute called *SSLCRLNamelist*. The queue manager's initial value for this attribute is blank.

On IBM i, you can specify authentication information objects, but the queue manager uses neither authentication information objects nor a namelist of authentication information objects. Only WebSphere MQ clients that use a client connection table generated by an IBM i queue manager use the authentication information specified for that IBM i queue manager. The *SSLCRLNamelist* queue manager attribute on IBM i determines what authentication information such clients use. See "Accessing CRLs and ARLs on IBM i" on page 700 for information about telling an IBM i queue manager how to access CRLs.

You can add up to 10 connections to alternative LDAP servers to the namelist, to ensure continuity of service if one or more LDAP servers fail. Note that the LDAP servers *must* contain identical information.

Accessing CRLs and ARLs on IBM i:

Use this procedure to access CRLs or ARLs on IBM i.

Note that in this section, information about Certificate Revocation Lists (CRLs) also applies to Authority Revocation Lists (ARLs).

Follow these steps to set up a CRL location for a specific certificate on IBM i:

1. Access the DCM interface, as described in “Accessing DCM” on page 620.
2. In the **Manage CRL locations** task category in the navigation panel, click **Add CRL location**. The Manage CRL Locations page is displayed in the task frame.
3. In the **CRL Location Name** field, type a CRL location name, for example LDAP Server #1
4. In the **LDAP Server** field, type the LDAP server name.
5. In the **Use Secure Sockets Layer (SSL)** field, select **Yes** if you want to connect to the LDAP server using SSL. Otherwise, select **No**.
6. In the **Port Number** field, type a port number for the LDAP server, for example 389.
7. If your LDAP server does not allow anonymous users to query the directory, type a login distinguished name for the server in the **login distinguished name** field.
8. Click **OK**. DCM informs you that it has created the CRL location.
9. In the navigation panel, click **Select a Certificate Store**. The Select a Certificate Store page is displayed in the task frame.
10. Select the **Other System Certificate Store** check box and click **Continue**. The Certificate Store and Password page is displayed.
11. In the **Certificate store path and filename** field, type the IFS path and file name you set when “Creating a certificate store on IBM i” on page 621.
12. Type a password in the **Certificate Store Password** field. Click **Continue**. The Current Certificate Store page is displayed in the task frame.
13. In the **Manage Certificates** task category in the navigation panel, click **Update CRL location assignment**. The CRL Location Assignment page is displayed in the task frame.
14. Select the radio button for the CA certificate to which you want to assign the CRL location. Click **Update CRL Location Assignment**. The Update CRL Location Assignment page is displayed in the task frame.
15. Select the radio button for the CRL location which you want to assign to the certificate. Click **Update Assignment**. DCM informs you that it has updated the assignment.

Note that DCM allows you to assign a different LDAP server by Certificate Authority.

Accessing CRLs and ARLs using WebSphere MQ Explorer:

You can use WebSphere MQ Explorer to tell a queue manager how to access CRLs.

Note that in this section, information about Certificate Revocation Lists (CRLs) also applies to Authority Revocation Lists (ARLs).

Use the following procedure to set up an LDAP connection to a CRL:

1. Ensure that you have started your queue manager.
2. Right-click the **Authentication Information** folder and click **New -> Authentication Information**. In the property sheet that opens:
 - a. On the first page **Create Authentication Information**, enter a name for the CRL(LDAP) object.
 - b. On the **General** page of **Change Properties**, select the connection type. Optionally you can enter a description.
 - c. Select the **CRL(LDAP)** page of **Change Properties**.

- d. Enter the LDAP server name as either the network name or the IP address.
 - e. If the server requires login details, provide a user ID and if necessary a password.
 - f. Click **OK**.
3. Right-click the **Namelist** folder and click **New -> Namelist**. In the property sheet that opens:
 - a. Type a name for the namelist.
 - b. Add the name of the CRL(LDAP) object (from step 2a on page 700) to the list.
 - c. Click **OK**.
 4. Right-click the queue manager, select **Properties**, and select the **SSL** page:
 - a. Select the **Check certificates received by this queue manager against Certification Revocation Lists** check box.
 - b. Type the name of the namelist (from step 3a) in the **CRL Namelist** field.

Accessing CRLs and ARLs with a WebSphere MQ MQI client:

You have three options for specifying the LDAP servers that hold CRLs for checking by a WebSphere MQ MQI client.

Note that in this section, information about Certificate Revocation Lists (CRLs) also applies to Authority Revocation Lists (ARLs).

The three ways of specifying the LDAP servers are as follows:

- Using a channel definition table
- Using the SSL configuration options structure, MQSCO, on an MQCONN call
- Using the Active Directory (on Windows systems with Active Directory support)

For more details, refer to the related information.

You can include up to 10 connections to alternative LDAP servers to ensure continuity of service if one or more LDAP servers fail. Note that the LDAP servers *must* contain identical information.

You cannot access LDAP CRLs from a WebSphere MQ MQI client channel running on Linux (zSeries platform).

Related concepts:

“Location of an OCSP responder, and of LDAP servers that hold CRLs”

Related reference:

 [setmqcrl \(WebSphere MQ V7.1 Reference\)](#)

Location of an OCSP responder, and of LDAP servers that hold CRLs:

On a WebSphere MQ MQI client system, you can specify the location of an OCSP responder, and of Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs).

You can specify these locations in three ways, listed here in order of decreasing precedence. (For IBM i, see Accessing CRLs and ARLs on IBM i.)

When a WebSphere MQ MQI client application issues an MQCONN call

You can specify an OCSP responder or an LDAP server holding CRLs on an **MQCONN** call.

On an **MQCONN** call, the connect options structure, MQCNO, can reference an SSL configuration options structure, MQSCO. In turn, the MQSCO structure can reference one or more authentication information record structures, MQAIR. Each MQAIR structure contains all the information a WebSphere MQ MQI

client requires to access an OCSP responder or an LDAP server holding CRLs. For example, one of the fields in an MQAIR structure is the URL at which a responder can be contacted. For more information about the MQAIR structure, see  MQAIR – Authentication information record (*WebSphere MQ V7.1 Reference*).

Using a client channel definition table (ccdt) to access an OCSP responder or LDAP servers

So that a WebSphere MQ MQI client can access an OCSP responder or LDAP servers that hold CRLs, include the attributes of one or more authentication information objects in a client channel definition table.

On a server queue manager, you can define one or more authentication information objects. The attributes of an authentication object contain all the information that is required to access an OCSP responder (on platforms where OCSP is supported) or an LDAP server that holds CRLs. One of the attributes specifies the OCSP responder URL, another specifies the host address, or IP address of a system on which an LDAP server runs.

An authentication information object with AUTHTYPE(OCSP) does not apply for use on IBM i or z/OS queue managers, but it can be specified on those platforms to be copied to the client channel definition table (CCDT) for client use.

To enable a WebSphere MQ MQI client to access an OCSP responder or LDAP servers that hold CRLs, the attributes of one or more authentication information objects can be included in a client channel definition table. You can include such attributes in one of the following ways:

On the server platforms AIX, HP-UX, Linux, IBM i, Solaris, and Windows

You can define a namelist that contains the names of one or more authentication information objects. You can then set the queue manager attribute, **SSLCRLNameList**, to the name of this namelist.

If you are using CRLs, more than one LDAP server can be configured to provide higher availability. The intention is that each LDAP server holds the same CRLs. If one LDAP server is unavailable when it is required, a WebSphere MQ MQI client can attempt to access another.

The attributes of the authentication information objects identified by the namelist are referred to collectively here as the *certificate revocation location*. When you set the queue manager attribute, **SSLCRLNameList**, to the name of the namelist, the certificate revocation location is copied into the client channel definition table associated with the queue manager. If the CCDT can be accessed from a client system as a shared file, or if the CCDT is then copied to a client system, the WebSphere MQ MQI client on that system can use the certificate revocation location in the CCDT to access an OCSP responder or LDAP servers that hold CRLs.

If the certificate revocation location of the queue manager is changed later, the change is reflected in the CCDT associated with the queue manager. If the queue manager attribute, **SSLCRLNameList**, is set to blank, the certificate revocation location is removed from the CCDT. These changes are not reflected in any copy of the table on a client system.

If you require the certificate revocation location at the client and server ends of an MQI channel to be different, and the server queue manager is the one that is used to create the certificate revocation location, you can do it as follows:

1. On the server queue manager, create the certificate revocation location for use on the client system.
2. Copy the CCDT containing the certificate revocation location to the client system.
3. On the server queue manager, change the certificate revocation location to what is required at the server end of the MQI channel.

On the server platform z/OS

On z/OS, a CCDT is generated by the MAKECLNT parameter of the COMMAND function of the WebSphere MQ utility program, CSQUTIL. The DISPLAY CHANNEL commands in the input data set determine which client-connection channel definitions are included in the table. Likewise, the DISPLAY AUTHINFO commands in the input data set determine which authentication information objects are used to form the certificate revocation location in the table.

The contents of a CCDT generated on z/OS do not depend on the value of any queue manager attributes, such as **SSLCRLNameList**, and cannot be updated dynamically. The only way you can change the certificate revocation location in a CCDT is to generate a new table by running **CSQUTIL** again.

Using Active Directory on Windows

On Windows systems, you can use the **setmqcr1** control command to publish the current CRL information in Active Directory.

Command **setmqcr1** does not publish OCSP information.

For information about this command and its syntax, see  [setmqcr1 \(WebSphere MQ V7.1 Reference\)](#).

Accessing CRLs and ARLs with WebSphere MQ classes for Java and WebSphere MQ classes for JMS:

WebSphere MQ classes for Java and WebSphere MQ classes for JMS access CRLs differently from other platforms.

For information about working with CRLs and ARLs with WebSphere MQ classes for Java, see  [Using certificate revocation lists \(WebSphere MQ V7.1 Programming Guide\)](#)

For information about working with CRLs and ARLs with WebSphere MQ classes for JMS, see  [SSLCERTSTORES object property \(WebSphere MQ V7.1 Programming Guide\)](#)

Manipulating authentication information objects

You can manipulate authentication information objects using MQSC or PCF commands, or the Websphere MQ Explorer.

The following MQSC commands act on authentication information objects:

- DEFINE AUTHINFO
- ALTER AUTHINFO
- DELETE AUTHINFO
- DISPLAY AUTHINFO

For a complete description of these commands, see “Script (MQSC) Commands” on page 77.

The following Programmable Command Format (PCF) commands act on authentication information objects:

- Create Authentication Information
- Copy Authentication Information
- Change Authentication Information
- Delete Authentication Information
- Inquire Authentication Information
- Inquire Authentication Information Names

For a complete description of these commands, see  Definitions of the Programmable Command Formats (*WebSphere MQ V7.1 Reference*).

On platforms where it is available, you can also use the WebSphere MQ Explorer.

Authorizing access to objects

This collection of topics contains information about using the object authority manager and channel exit programs to control access to objects.

On UNIX, Linux, and Windows systems, you control access to objects by using the object authority manager (OAM). This collection of topics contains information about using the command interface to the OAM. It also contains a checklist you can use to determine what tasks to perform to apply security to your system, and considerations for granting users the authority to administer IBM WebSphere MQ and to work with IBM WebSphere MQ objects. If the supplied security mechanisms do not meet your needs, you can develop your own channel exit programs.

Controlling access to objects by using the OAM on UNIX, Linux and Windows systems

The object authority manager (OAM) provides a command interface for granting and revoking authority to WebSphere MQ objects.

You must be suitably authorized to use these commands, as described in “Authority to administer WebSphere MQ on UNIX, Linux and Windows systems” on page 751. User IDs that are authorized to administer WebSphere MQ have *super user* authority to the queue manager, which means that you do not have to grant them further permission to issue any MQI requests or commands.

Giving access to a WebSphere MQ object on UNIX or Linux systems and Windows

Use the **setmqaut** control command, or the **MQCMD_SET_AUTH_REC** PCF command to give users, and groups of users, access to WebSphere MQ objects.

For a full definition of the **setmqaut** control command and its syntax, see  **setmqaut** (*WebSphere MQ V7.1 Reference*), and for a full definition of the **MQCMD_SET_AUTH_REC** PCF command and its syntax, see  Set Authority Record (*WebSphere MQ V7.1 Reference*).

The queue manager must be running to use this command. When you have changed access for a principal, the changes are reflected immediately by the OAM.

To give users access to an object, you need to specify:

- The name of the queue manager that owns the objects you are working with; if you do not specify the name of a queue manager, the default queue manager is assumed.
- The name and type of the object (to identify the object uniquely). You specify the name as a *profile*; this is either the explicit name of the object, or a generic name, including wildcard characters. For a detailed description of generic profiles, and the use of wildcard characters within them, see “Using OAM generic profiles on UNIX or Linux systems and Windows” on page 705.
- One or more principals and group names to which the authority applies.

If a user ID contains spaces, enclose it in quotation marks when you use this command. On Windows systems, you can qualify a user ID with a domain name. If the actual user ID contains an at sign (@) symbol, replace it with @@ to show that it is part of the user ID, not the delimiter between the user ID and the domain name.

- A list of authorizations. Each item in the list specifies a type of access that is to be granted to that object (or revoked from it). Each authorization in the list is specified as a keyword, prefixed with a

plus sign (+) or a minus sign (-). Use a plus sign to add the specified authorization, and a minus sign to remove the authorization. There must be no spaces between the + or - sign and the keyword.

You can specify any number of authorizations in a single command. For example, the list of authorizations to permit a user or group to put messages on a queue and to browse them, but to revoke access to get messages is:

```
+browse -get +put
```

Examples of using the `setmqaut` command

The following examples show how to use the `setmqaut` command to grant and revoke permission to use an object:

```
setmqaut -m saturn.queue.manager -t queue -n RED.LOCAL.QUEUE
        -g groupa +browse -get +put
```

In this example:

- saturn.queue.manager is the queue manager name
- queue is the object type
- RED.LOCAL.QUEUE is the object name
- groupa is the identifier of the group with authorizations that are to change
- +browse -get +put is the authorization list for the specified queue
 - +browse adds authorization to browse messages on the queue (to issue **MQGET** with the browse option)
 - -get removes authorization to get (**MQGET**) messages from the queue
 - +put adds authorization to put (**MQPUT**) messages on the queue

The following command revokes put authority on the queue MyQueue from principal fvuser and from groups groupa and groupb. On UNIX and Linux systems, this command also revokes put authority for all principals in the same primary group as fvuser.

```
setmqaut -m saturn.queue.manager -t queue -n MyQueue -p fvuser
        -g groupa -g groupb -put
```

Using the command with a different authorization service

If you are using your own authorization service instead of the OAM, you can specify the name of this service on the `setmqaut` command to direct the command to this service. You must specify this parameter if you have multiple installable components running at the same time; if you do not, the update is made to the first installable component for the authorization service. By default, this is the supplied OAM.

Using OAM generic profiles on UNIX or Linux systems and Windows

OAM generic profiles enable you to set the authority a user has to many objects at once, rather than having to issue separate `setmqaut` commands against each individual object when it is created.

Using generic profiles in the `setmqaut` command enables you to set a generic authority for all objects that fit that profile.

This collection of topics describes the use of generic profiles in more detail.

Using wildcard characters in OAM profiles

What makes a profile generic is the use of special characters (wildcard characters) in the profile name. For example, the question mark (?) wildcard character matches any single character in a name. So, if you specify ABC.?EF, the authorization you give to that profile applies to any objects with the names ABC.DEF, ABC.CEF, ABC.BEF, and so on.

The wildcard characters available are:

- ? Use the question mark (?) instead of any single character. For example, AB.?D applies to the objects AB.CD, AB.ED, and AB.FD.
- * Use the asterisk (*) as:
 - A *qualifier* in a profile name to match any one qualifier in an object name. A qualifier is the part of an object name delimited by a period. For example, in ABC.DEF.GHI, the qualifiers are ABC, DEF, and GHI.
For example, ABC.*.JKL applies to the objects ABC.DEF.JKL, and ABC.GHI.JKL. (Note that it does **not** apply to ABC.JKL; * used in this context always indicates one qualifier.)
 - A character within a qualifier in a profile name to match zero or more characters within the qualifier in an object name.
For example, ABC.DE*.JKL applies to the objects ABC.DE.JKL, ABC.DEF.JKL, and ABC.DEGH.JKL.
- ** Use the double asterisk (**) **once** in a profile name as:
 - The entire profile name to match all object names. For example if you use -t prcs to identify processes, then use ** as the profile name, you change the authorizations for all processes.
 - As either the beginning, middle, or ending qualifier in a profile name to match zero or more qualifiers in an object name. For example, **.ABC identifies all objects with the final qualifier ABC.

Note: When using wildcard characters on UNIX and Linux systems, you **must** enclose the profile name in single quotation marks.

Profile priorities

An important point to understand when using generic profiles is the priority that profiles are given when deciding what authorities to apply to an object being created. For example, suppose that you have issued the commands:

```
setmqaut -n AB.* -t q +put -p fred  
setmqaut -n AB.C* -t q +get -p fred
```

The first gives put authority to all queues for the principal fred with names that match the profile AB.*; the second gives get authority to the same types of queue that match the profile AB.C*.

Suppose that you now create a queue called AB.CD. According to the rules for wildcard matching, either setmqaut could apply to that queue. So, does it have put or get authority?

To find the answer, you apply the rule that, whenever multiple profiles can apply to an object, **only the most specific applies**. The way that you apply this rule is by comparing the profile names from left to right. Wherever they differ, a non-generic character is more specific than a generic character. So, in the example above, the queue AB.CD has **get** authority (AB.C* is more specific than AB.*).

When you are comparing generic characters, the order of *specificity* is:

1. ?
2. *
3. **

Dumping profile settings

For a full definition of the **dmpmqaut** control command and its syntax, see  **dmpmqaut** (*WebSphere MQ V7.1 Reference*), and for a full definition of the **MQCMD_INQUIRE_AUTH_RECS** PCF command and its syntax, see  **Inquire Authority Records** (*WebSphere MQ V7.1 Reference*).

The following examples show the use of the **dmpmqaut** control command to dump authority records for generic profiles:

1. This example dumps all authority records with a profile that matches queue a.b.c for principal user1.

```
dmpmqaut -m qm1 -n a.b.c -t q -p user1
```

The resulting dump looks something like this:

```
profile:    a.b.*
object type: queue
entity:     user1
type:       principal
authority:  get, browse, put, inq
```

Note: Although UNIX and Linux users can use the **-p** option for the **dmpmqaut** command, they must use **-g** groupname instead when defining authorizations.

2. This example dumps all authority records with a profile that matches queue a.b.c.

```
dmpmqaut -m qmgr1 -n a.b.c -t q
```

The resulting dump looks something like this:

```
profile:    a.b.c
object type: queue
entity:     Administrator
type:       principal
authority:  all
-----
profile:    a.b.*
object type: queue
entity:     user1
type:       principal
authority:  get, browse, put, inq
-----
profile:    a.**
object type: queue
entity:     group1
type:       group
authority:  get
```

3. This example dumps all authority records for profile a.b.*, of type queue.

```
dmpmqaut -m qmgr1 -n a.b.* -t q
```

The resulting dump looks something like this:

```
profile:    a.b.*
object type: queue
entity:     user1
type:       principal
authority:  get, browse, put, inq
```

4. This example dumps all authority records for queue manager qmX.

```
dmpmqaut -m qmX
```

The resulting dump looks something like this:

```
profile:    q1
object type: queue
entity:     Administrator
type:       principal
authority:  all
-----
profile:    q*
```

```

object type: queue
entity:      user1
type:       principal
authority:   get, browse
-----
profile:     name.*
object type: namelist
entity:      user2
type:       principal
authority:   get
-----
profile:     pr1
object type: process
entity:      group1
type:       group
authority:   get

```

5. This example dumps all profile names and object types for queue manager qmX.
- ```
dmpmqaut -m qmX -l
```

The resulting dump looks something like this:

```

profile: q1, type: queue
profile: q*, type: queue
profile: name.*, type: namelist
profile: pr1, type: process

```

**Note:** For WebSphere MQ for Windows only, all principals displayed include domain information, for example:

```

profile: a.b.*
object type: queue
entity: user1@domain1
type: principal
authority: get, browse, put, inq

```

### Using wildcard characters in OAM profiles:

Use wildcard characters in an object authority manager (OAM) profile name to make that profile applicable to more than one object.

What makes a profile generic is the use of special characters (wildcard characters) in the profile name. For example, the question mark (?) wildcard character matches any single character in a name. So, if you specify ABC.?EF, the authorization you give to that profile applies to any objects with the names ABC.DEF, ABC.CEF, ABC.BEF, and so on.

The wildcard characters available are:

- ? Use the question mark (?) instead of any single character. For example, AB.?D applies to the objects AB.CD, AB.ED, and AB.FD.
- \* Use the asterisk (\*) as:
  - A *qualifier* in a profile name to match any one qualifier in an object name. A qualifier is the part of an object name delimited by a period. For example, in ABC.DEF.GHI, the qualifiers are ABC, DEF, and GHI.  
For example, ABC.\*.JKL applies to the objects ABC.DEF.JKL, and ABC.GHI.JKL. (Note that it does **not** apply to ABC.JKL; \* used in this context always indicates one qualifier.)
  - A character within a qualifier in a profile name to match zero or more characters within the qualifier in an object name.  
For example, ABC.DE\*.JKL applies to the objects ABC.DE.JKL, ABC.DEF.JKL, and ABC.DEGH.JKL.

- \*\* Use the double asterisk (\*\*) **once** in a profile name as:
- The entire profile name to match all object names. For example if you use -t prcs to identify processes, then use \*\* as the profile name, you change the authorizations for all processes.
  - As either the beginning, middle, or ending qualifier in a profile name to match zero or more qualifiers in an object name. For example, \*\*.ABC identifies all objects with the final qualifier ABC.

**Note:** When using wildcard characters on UNIX and Linux systems, you **must** enclose the profile name in single quotation marks.

### Profile priorities:

More than one generic profile can apply to a single object. Where this is the case, the most specific rule applies.

An important point to understand when using generic profiles is the priority that profiles are given when deciding what authorities to apply to an object being created. For example, suppose that you have issued the commands:

```
setmqaut -n AB.* -t q +put -p fred
setmqaut -n AB.C* -t q +get -p fred
```

The first gives put authority to all queues for the principal fred with names that match the profile AB.\*; the second gives get authority to the same types of queue that match the profile AB.C\*.

Suppose that you now create a queue called AB.CD. According to the rules for wildcard matching, either setmqaut could apply to that queue. So, does it have put or get authority?

To find the answer, you apply the rule that, whenever multiple profiles can apply to an object, **only the most specific applies**. The way that you apply this rule is by comparing the profile names from left to right. Wherever they differ, a non-generic character is more specific than a generic character. So, in the example above, the queue AB.CD has **get** authority (AB.C\* is more specific than AB.\*).

When you are comparing generic characters, the order of *specificity* is:

1. ?
2. \*
3. \*\*

### Dumping profile settings:

Use the **dmpmqaut** control command or the **MQCMD\_INQUIRE\_AUTH\_RECS** PCF command to dump the current authorizations associated with a specified profile.

For a full definition of the **dmpmqaut** control command and its syntax, see  **dmpmqaut** (*WebSphere MQ V7.1 Reference*), and for a full definition of the **MQCMD\_INQUIRE\_AUTH\_RECS** PCF command and its syntax, see  **Inquire Authority Records** (*WebSphere MQ V7.1 Reference*).

The following examples show the use of the **dmpmqaut** control command to dump authority records for generic profiles:

1. This example dumps all authority records with a profile that matches queue a.b.c for principal user1.  

```
dmpmqaut -m qm1 -n a.b.c -t q -p user1
```

The resulting dump looks something like this example:

```
profile: a.b.*
object type: queue
entity: user1
type: principal
authority: get, browse, put, inq
```

**Note:** UNIX and Linux users cannot use the `-p` option; they must use `-g` groupname instead.

2. This example dumps all authority records with a profile that matches queue a.b.c.

```
dmpmqaut -m qmgr1 -n a.b.c -t q
```

The resulting dump looks something like this example:

```
profile: a.b.c
object type: queue
entity: Administrator
type: principal
authority: all

profile: a.b.*
object type: queue
entity: user1
type: principal
authority: get, browse, put, inq

profile: a.**
object type: queue
entity: group1
type: group
authority: get
```

3. This example dumps all authority records for profile a.b.\*, of type queue.

```
dmpmqaut -m qmgr1 -n a.b.* -t q
```

The resulting dump looks something like this example:

```
profile: a.b.*
object type: queue
entity: user1
type: principal
authority: get, browse, put, inq
```

4. This example dumps all authority records for queue manager qmX.

```
dmpmqaut -m qmX
```

The resulting dump looks something like this example:

```
profile: q1
object type: queue
entity: Administrator
type: principal
authority: all

profile: q*
object type: queue
entity: user1
type: principal
authority: get, browse

profile: name.*
object type: namelist
entity: user2
type: principal
```

```

authority: get

profile: pr1
object type: process
entity: group1
type: group
authority: get

```

5. This example dumps all profile names and object types for queue manager qmX.

```
dmpmqaut -m qmX -l
```

The resulting dump looks something like this example:

```

profile: q1, type: queue
profile: q*, type: queue
profile: name.*, type: namelist
profile: pr1, type: process

```

**Note:** For WebSphere MQ for Windows only, all principals displayed include domain information, for example:

```

profile: a.b.*
object type: queue
entity: user1@domain1
type: principal
authority: get, browse, put, inq

```

## Displaying access settings

Use the **dspmqa** control command, or the **MQCMD\_INQUIRE\_ENTITY\_AUTH** PCF command to view the authorizations that a specific principal or group has for a particular object.

The queue manager must be running to use this command. When you change access for a principal, the changes are reflected immediately by the OAM. Authorization can be displayed for only one group or principal at a time. For a full definition of the **dmpmqaut** control command and its syntax, see

 **dmpmqaut** (*WebSphere MQ V7.1 Reference*), and for a full definition of the **MQCMD\_INQUIRE\_ENTITY\_AUTH** PCF command and its syntax, see  **Inquire Entity Authority** (*WebSphere MQ V7.1 Reference*).

The following example shows the use of the **dspmqa** control command to display the authorizations that the group GpAdmin has to a process definition named Annuities that is on queue manager QueueMan1.

```
dspmqa -m QueueMan1 -t process -n Annuities -g GpAdmin
```

## Changing and revoking access to a WebSphere MQ object

To change the level of access that a user or group has to an object, use the **setmqaut** command. To revoke the access of a particular user that is a member of a group that has authorization, remove the user from the group.

The process of removing the user from a group is described in:

- “Creating and managing groups on Windows” on page 465
- “Creating and managing groups on HP-UX” on page 467
- “Creating and managing groups on AIX” on page 468
- “Creating and managing groups on Solaris” on page 470
- “Creating and managing groups on Linux” on page 470

The user ID that creates a WebSphere MQ object is granted full control authorities to that object. If you remove this user ID from the local mqm group (or the Administrators group on Windows systems) these

authorities are not revoked. Use the **setmqaut** control command or the **MQCMD\_DELETE\_AUTH\_REC** PCF command to revoke access to an object for the user ID that created it, after removing it from the mqm or Administrators group. For a full definition of the setmqaut control command and its syntax, see

 **setmqaut** (*WebSphere MQ V7.1 Reference*), and for a full definition of the **MQCMD\_INQUIRE\_ENTITY\_AUTH** PCF command and its syntax, see  **Inquire Entity Authority** (*WebSphere MQ V7.1 Reference*).

On Windows, delete the OAM entries corresponding to a particular Windows user account before deleting the user profile. It is impossible to remove the OAM entries after removing the user account.

## Preventing security access checks on Windows, UNIX and Linux systems

To turn off all security checking you can disable the OAM. This might be suitable for a test environment. Having disabled or removed the OAM, you cannot add an OAM to an existing queue manager.

If you decide that you do not want to perform security checks (for example, in a test environment), you can disable the OAM in one of two ways:

- Before you create a queue manager, set the operating system environment variable MQSNOAUT. (If you do this, you cannot add an OAM later.)

See  **Environment Variables** (*WebSphere MQ V7.1 Reference*) for more information about the implications of setting the MQSNOAUT variable.

- Edit the queue manager configuration file (or the registry on Windows) to remove the service. (If you do this, you cannot add an OAM later.)

If you use setmqaut, or dspmqaut while the OAM is disabled, note the following points:

- The OAM does not validate the specified principal, or group, meaning that the command can accept invalid values.
- The OAM does not perform security checks and indicates that all principals and groups are authorized to perform all applicable object operations.

When an OAM is removed, it cannot be put back on an existing queue manager. This is because the OAM needs to be in place at object creation time. To use the WebSphere MQ OAM again after it has been removed, the queue manager needs to be rebuilt.

## Granting required access to resources

Use this topic to determine what tasks to perform to apply security to your WebSphere MQ system.

### About this task

During this task, you decide what actions are necessary to apply the appropriate level of security to the elements of your WebSphere MQ installation. Each individual task you are referred to gives step-by-step instructions for all platforms.

### Procedure

1. Do you need to limit access to your queue manager to certain users?
  - a. No: Take no further action.
  - b. Yes: Go to the next question.
2. Do these users need partial administrative access on a subset of queue manager resources?
  - a. No: Go to the next question.
  - b. Yes: See “Granting partial administrative access on a subset of queue manager resources” on page 713.
3. Do these users need full administrative access on a subset of queue manager resources?
  - a. No: Go to the next question.

- b. Yes: See “Granting full administrative access on a subset of queue manager resources” on page 721.
- 4. Do these users need read only access to all queue manager resources?
  - a. No: Go to the next question.
  - b. Yes: See “Granting read-only access to all resources on a queue manager” on page 728.
- 5. Do these users need full administrative access on all queue manager resources?
  - a. No: Go to the next question.
  - b. Yes: See “Granting full administrative access to all resources on a queue manager” on page 729.
- 6. Do you need user applications to connect to your queue manager?
  - a. No: Disable connectivity, as described in “Removing connectivity to the queue manager” on page 730
  - b. Yes: See “Allowing user applications to connect to your queue manager” on page 731.

### Granting partial administrative access on a subset of queue manager resources

You need to give certain users partial administrative access to some, but not all, queue manager resources. Use this table to determine the actions you need to take.

| The users need to administer objects of this type | Perform this action                                                                                                                                   |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queues                                            | Grant partial administrative access to the required queues, as described in “Granting limited administrative access to some queues”                   |
| Topics                                            | Grant partial administrative access to the required topics, as described in “Granting limited administrative access to some topics” on page 715       |
| Channels                                          | Grant partial administrative access to the required channels, as described in “Granting limited administrative access to some channels” on page 716   |
| The queue manager                                 | Grant partial administrative access to the queue manager, as described in “Granting limited administrative access to a queue manager” on page 717     |
| Processes                                         | Grant partial administrative access to the required processes, as described in “Granting limited administrative access to some processes” on page 718 |
| Namelists                                         | Grant partial administrative access to the required namelists, as described in “Granting limited administrative access to some namelists” on page 719 |
| Services                                          | Grant partial administrative access to the required services, as described in “Granting limited administrative access to some services” on page 720   |

### Granting limited administrative access to some queues:

Grant partial administrative access to some queues on a queue manager, to each group of users with a business need for it.

#### About this task

To grant limited administrative access to some queues for some actions, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:
 

```
setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName ReqdAction
```
- For IBM i, issue the following command:
 

```
GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(ReqdAction)
MQMNAME('QMgrName')
```

- For z/OS, issue the following commands:

```
RDEFINE MQADMIN QMgrName.QUEUE.ObjectProfile UACC(NONE)
PERMIT QMgrName.QUEUE.ObjectProfile CLASS(MQADMIN) ID(GroupName) ACCESS(ALTER)
```

These commands grant access to the specified queue. To determine which MQSC commands the user can perform on the queue, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.QType UACC(NONE)
PERMIT QMgrName.ReqdAction.QType CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY QUEUE command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.QType UACC(NONE)
PERMIT QMgrName.DISPLAY.QType CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

### **GroupName**

The name of the group to be granted access.

### **ReqdAction**

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +dlt, +dsp. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMDLT, \*ADM DSP. The authorization \*ALLADM is equivalent to all these individual authorizations.
- On z/OS, one of the values ALTER, CLEAR, DELETE, or MOVE.

**Note:** Granting +crt for queues indirectly makes the user or group an administrator. Do not use +crt authority to grant limited administrative access to some queues.

### **QType**

For the DISPLAY command, one of the values QUEUE, QLOCAL, QALIAS, QMODEL, QREMOTE, or QCLUSTER.

For other values of *ReqdAction*, one of the values QLOCAL, QALIAS, QMODEL, or QREMOTE.

### **Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### **Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### **Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Granting limited administrative access to some topics:

Grant partial administrative access to some topics on a queue manager, to each group of users with a business need for it.

### About this task

To grant limited administrative access to some topics for some actions, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t topic -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*TOPIC) USER(GroupName) AUT(ReqdAction)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.TOPIC.ObjectProfile UACC(NONE)  
PERMIT QMgrName.TOPIC.ObjectProfile CLASS(MQADMIN) ID(GroupName) ACCESS(ALTER)`

These commands grant access to the specified topic. To determine which MQSC commands the user can perform on the topic, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.TOPIC UACC(NONE)
PERMIT QMgrName.ReqdAction.TOPIC CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY TOPIC command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.TOPIC UACC(NONE)
PERMIT QMgrName.DISPLAY.TOPIC CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

#### **ReqdAction**

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +dsp, +ctrl. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMDLT, \*ADMDSP, \*CTRL. The authorization \*ALLADM is equivalent to all these individual authorizations.
- On z/OS, one of the values ALTER, CLEAR, DEFINE, DELETE, or MOVE.

## Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

## Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

## Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Granting limited administrative access to some channels:

Grant partial administrative access to some channels on a queue manager, to each group of users with a business need for it.

### About this task

To grant limited administrative access to some channels for some actions, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t channel -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*CHL) USER(GroupName) AUT(ReqdAction)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.CHANNEL.ObjectProfile UACC(NONE)  
PERMIT QMgrName.CHANNEL.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

These commands grant access to the specified channel. To determine which MQSC commands the user can perform on the channel, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.CHANNEL UACC(NONE)
PERMIT QMgrName.ReqdAction.CHANNEL CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY CHANNEL command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.CHANNEL UACC(NONE)
PERMIT QMgrName.DISPLAY.CHANNEL CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

#### **ReqdAction**

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +dsp, +ctrl, +ctrlx. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMDLT, \*ADMDSP, \*CTRL, \*CTRLx. The authorization \*ALLADM is equivalent to all these individual authorizations.
- On z/OS, one of the values ALTER, CLEAR, DEFINE, DELETE, or MOVE.

#### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### Related reference:

 setmqaut (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

#### Related information:

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

### Granting limited administrative access to a queue manager:

Grant partial administrative access to a queue manager, to each group of users with a business need for it.

#### About this task

To grant limited administrative access to perform some actions on the queue manager, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t qmgr -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*MQM) USER(GroupName) AUT(ReqdAction)  
 MQMNAME('QMgrName')`

#### Results

To determine which MQSC commands the user can perform on the queue manager, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.QMGR UACC(NONE)
PERMIT QMgrName.ReqdAction.QMGR CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY QMGR command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.QMGR UACC(NONE)
PERMIT QMgrName.DISPLAY.QMGR CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### QMgrName

The name of the queue manager.

## ObjectProfile

The name of the object or generic profile for which to change authorizations.

## GroupName

The name of the group to be granted access.

## ReqdAction

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +dsp. The authorization +alladm is equivalent to +chg +clr +dlt +dsp. Although +set is an MQI authorization and not normally considered administrative, granting +set on the queue manager can indirectly lead to full administrative authority. Do not grant +set to ordinary users and applications.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMDLT, \*ADMDSP. The authorization \*ALLADM is equivalent to all these individual authorizations.

## Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

## Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

## Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Granting limited administrative access to some processes:

Grant partial administrative access to some processes on a queue manager, to each group of users with a business need for it.

## About this task

To grant limited administrative access to some processes for some actions, use the appropriate commands for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t process -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*PRC) USER(GroupName) AUT(ReqdAction)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.PROCESS.ObjectProfile UACC(NONE)  
PERMIT QMgrName.PROCESS.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

These commands grant access to the specified channel. To determine which MQSC commands the user can perform on the channel, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.PROCESS UACC(NONE)
PERMIT QMgrName.ReqdAction.PROCESS CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY PROCESS command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.PROCESS UACC(NONE)
PERMIT QMgrName.DISPLAY.PROCESS CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**ReqdAction**

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +dsp. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMCLT, \*ADM DSP. The authorization \*ALLADM is equivalent to all these individual authorizations.
- On z/OS, one of the values ALTER, CLEAR, DEFINE, DELETE, or MOVE.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**



setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

**Related information:**



Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

**Granting limited administrative access to some namelists:**

Grant partial administrative access to some namelists on a queue manager, to each group of users with a business need for it.

**About this task**

To grant limited administrative access to some namelists for some actions, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t namelist -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*NMLIST) USER(GroupName) AUT(ReqdAction)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:

```
RDEFINE MQADMIN QMgrName.NAMELIST.ObjectProfile UACC(NONE)
PERMIT QMgrName.NAMELIST.ObjectProfile CLASS(MQADMIN) ID(GroupName) ACCESS(ALTER)
```

These commands grant access to the specified namelist. To determine which MQSC commands the user can perform on the namelist, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.NAMELIST UACC(NONE)
PERMIT QMgrName.ReqdAction.NAMELIST CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY NAMELIST command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.NAMELIST UACC(NONE)
PERMIT QMgrName.DISPLAY.NAMELIST CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

#### **ReqdAction**

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +ctrl, +ctrlx, +dsp. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMDLT, \*ADMDSP, \*CTRL, \*CTRLX. The authorization \*ALLADM is equivalent to all these individual authorizations.
- On z/OS, one of the values ALTER, CLEAR, DEFINE, DELETE, or MOVE.

#### **Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### **Related reference:**



setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

#### **Related information:**



Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

#### **Granting limited administrative access to some services:**

Grant partial administrative access to some services on a queue manager, to each group of users with a business need for it.

#### **About this task**

To grant limited administrative access to some services for some actions, use the appropriate commands for your operating system.

**Note:** Service objects do not exist on z/OS.

### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t service -g GroupName ReqdAction`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*SVC) USER(GroupName) AUT(ReqdAction)  
MQMNAME('QMgrName')`

### Results

These commands grant access to the specified service. To determine which MQSC commands the user can perform on the service, issue the following commands for each MQSC command:

```
RDEFINE MQCMDS QMgrName.ReqdAction.SERVICE UACC(NONE)
PERMIT QMgrName.ReqdAction.SERVICE CLASS(MQCMDS) ID(GroupName) ACCESS(ALTER)
```

To permit the user to use the DISPLAY SERVICE command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.DISPLAY.SERVICE UACC(NONE)
PERMIT QMgrName.DISPLAY.SERVICE CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### QMgrName

The name of the queue manager.

#### ObjectProfile

The name of the object or generic profile for which to change authorizations.

#### GroupName

The name of the group to be granted access.

#### ReqdAction

The action you are allowing the group to take:

- On UNIX, Linux and Windows systems, any combination of the following authorizations: +chg, +clr, +crt, +dlt, +ctrl, +ctrlx, +dsp. The authorization +alladm is equivalent to +chg +clr +dlt +dsp.
- On IBM i, any combination of the following authorizations: \*ADMCHG, \*ADMCLR, \*ADMCRT, \*ADMDLT, \*ADMDSP, \*CTRL, \*CTRLX. The authorization \*ALLADM is equivalent to all these individual authorizations.

#### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### Related reference:



setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

#### Related information:



Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

## Granting full administrative access on a subset of queue manager resources

You need to give certain users full administrative access to some, but not all, queue manager resources. Use these tables to determine the actions you need to take.

| The users need to administer objects of this type | Perform this action                                                                                                                             |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Queues                                            | Grant full administrative access to the required queues, as described in “Granting full administrative access to some queues”                   |
| Topics                                            | Grant full administrative access to the required topics, as described in “Granting full administrative access to some topics” on page 723       |
| Channels                                          | Grant full administrative access to the required channels, as described in “Granting full administrative access to some channels” on page 724   |
| The queue manager                                 | Grant full administrative access to the queue manager, as described in “Granting full administrative access to a queue manager” on page 725     |
| Processes                                         | Grant full administrative access to the required processes, as described in “Granting full administrative access to some processes” on page 726 |
| Namelist                                          | Grant full administrative access to the required namelists, as described in “Granting full administrative access to some namelists” on page 726 |
| Services                                          | Grant full administrative access to the required services, as described in “Granting full administrative access to some services” on page 727   |

### Granting full administrative access to some queues:

Grant full administrative access to some queues on a queue manager, to each group of users with a business need for it.

#### About this task

To grant full administrative access to some queues, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.QUEUE.ObjectProfile UACC(NONE)  
PERMIT QMgrName.QUEUE.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

The variable names have the following meanings:

#### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### ObjectProfile

The name of the object or generic profile for which to change authorizations.

#### GroupName

The name of the group to be granted access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:



setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

### Related information:



Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

### Granting full administrative access to some topics:

Grant full administrative access to some topics on a queue manager, to each group of users with a business need for it.

#### About this task

To grant full administrative access to some topics for some actions, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t topic -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*TOPIC) USER(GroupName) AUT(ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.TOPIC.ObjectProfile UACC(NONE)  
PERMIT QMgrName.TOPIC.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 `setmqaut` (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

**Related information:**

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

**Granting full administrative access to some channels:**

Grant full administrative access to some channels on a queue manager, to each group of users with a business need for it.

**About this task**

To grant full administrative access to some channels, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t channel -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*CHL) USER(GroupName) AUT(ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.CHANNEL.ObjectProfile UACC(NONE)  
PERMIT QMgrName.CHANNEL.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

**Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

**Granting full administrative access to a queue manager:**

Grant full administrative access to a queue manager, to each group of users with a business need for it.

**About this task**

To grant full administrative access to the queue manager, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -t qmgr -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*MQM) USER(GroupName) AUT(*ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.QMGR UACC(NONE)  
PERMIT QMgrName.QMGR CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

**Related information:**

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

### Granting full administrative access to some processes:

Grant full administrative access to some processes on a queue manager, to each group of users with a business need for it.

#### About this task

To grant full administrative access to some processes, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t process -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*PRC) USER(GroupName) AUT(*ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.CHANNEL.ObjectProfile UACC(NONE)  
PERMIT QMgrName.PROCESS.ObjectProfile CLASS(MQADMIN) ID(GroupName) ACCESS(ALTER)`

The variable names have the following meanings:

#### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### ObjectProfile

The name of the object or generic profile for which to change authorizations.

#### GroupName

The name of the group to be granted access.

#### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### Related reference:

 `setmqaut` (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

#### Related information:

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

### Granting full administrative access to some namelists:

Grant full administrative access to some namelists on a queue manager, to each group of users with a business need for it.

#### About this task

To grant full administrative access to some namelists, use the appropriate commands for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t namelist -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*NMLIST) USER(GroupName)  
AUT(*ALLADM) MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.NAMELIST.ObjectProfile UACC(NONE)  
PERMIT QMgrName.NAMELIST.ObjectProfile CLASS(MQADMIN) ID(GroupName )  
ACCESS(ALTER)`

The variable names have the following meanings:

### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### ObjectProfile

The name of the object or generic profile for which to change authorizations.

### GroupName

The name of the group to be granted access.

## Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

## Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

## Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Granting full administrative access to some services:

Grant full administrative access to some services on a queue manager, to each group of users with a business need for it.

## About this task

To grant full administrative access to some services, use the appropriate commands for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t service -g GroupName +alladm`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*SVC) USER(GroupName) AUT(*ALLADM)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQADMIN QMgrName.SERVICE.ObjectProfile UACC(NONE)  
PERMIT QMgrName.SERVICE.ObjectProfile CLASS(MQADMIN) ID(GroupName ) ACCESS(ALTER)`

The variable names have the following meanings:

### QMGrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### ObjectProfile

The name of the object or generic profile for which to change authorizations.

### GroupName

The name of the group to be granted access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Granting read-only access to all resources on a queue manager

Grant read-only access to all the resources on a queue manager, to each user or group of users with a business need for it.

### About this task

Use the Add Role Based Authorities wizard or the appropriate commands for your operating system.

### Procedure

- Using the wizard:
  - In the WebSphere MQ Explorer Navigator pane, right-click the queue manager and click **Object Authorities > Add Role Based Authorities**. The Add Role Based Authorities wizard opens.
- For UNIX and Windows systems, issue the following commands:

```
setmqaut -m QMgrName -n ** -t queue -g GroupName +browse +dsp
setmqaut -m QMgrName -n SYSTEM.ADMIN.COMMAND.QUEUE -t queue -g GroupName +dsp +inq +put
setmqaut -m QMgrName -n SYSTEM.MQEXPLORER.REPLY.MODEL -t queue -g GroupName +dsp +inq +get
setmqaut -m QMgrName -n ** -t topic -g GroupName +dsp
setmqaut -m QMgrName -n ** -t channel -g GroupName +dsp
setmqaut -m QMgrName -n ** -t clntconn -g GroupName +dsp
setmqaut -m QMgrName -n ** -t authinfo -g GroupName +dsp
setmqaut -m QMgrName -n ** -t listener -g GroupName +dsp
setmqaut -m QMgrName -n ** -t namelist -g GroupName +dsp
setmqaut -m QMgrName -n ** -t process -g GroupName +dsp
setmqaut -m QMgrName -n ** -t service -g GroupName +dsp
setmqaut -m QMgrName -t qmgr -g GroupName +dsp +inq +connect
```

The specific authorities to SYSTEM.ADMIN.COMMAND.QUEUE and SYSTEM.MQEXPLORER.REPLY.MODEL are necessary only if you want to use the MQ Explorer.

- For IBM i, issue the following commands:

```
GRTMQMAUT OBJ(*ALL) OBJTYPE(*Q) USER('GroupName') AUT(*ADM DSP *BROWSE) MQMNAME('QMGrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*TOPIC) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMGrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*CHL) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMGrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*CLTCN) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMGrName')
```

```
GRTMQMAUT OBJ(*ALL) OBJTYPE(*AUTHINFO) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMgrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*LSR) USER('GroupName') AUT(*ADM DSP)MQMNAME('QMgrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*NMLIST) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMgrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*PRC) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMgrName')
GRTMQMAUT OBJ(*ALL) OBJTYPE(*SVC) USER('GroupName') AUT(*ADM DSP) MQMNAME('QMgrName')
GRTMQMAUT OBJ('object-name') OBJTYPE(*MQM) USER('GroupName') AUT(*ADM DSP *CONNECT *INQ)
MQMNAME('QMgrName')
```

- For z/OS, issue the following commands:

```
RDEFINE MQQUEUE QMgrName.** UACC(NONE)
PERMIT QMgrName.** CLASS(MQQUEUE) ID(GroupName) ACCESS(READ)
RDEFINE MQTOPIC QMgrName.** UACC(NONE)
PERMIT QMgrName.** CLASS(MQTOPIC) ID(GroupName) ACCESS(READ)
RDEFINE MQPROC QMgrName.** UACC(NONE)
PERMIT QMgrName.** CLASS(MQPROC) ID(GroupName) ACCESS(READ)
RDEFINE MQNLIST QMgrName.** UACC(NONE)
PERMIT QMgrName.** CLASS(MQNLIST) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.BATCH UACC(NONE)
PERMIT QMgrName.BATCH CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.CICS UACC(NONE)
PERMIT QMgrName.CICS CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.IMS UACC(NONE)
PERMIT QMgrName.IMS CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.CHIN UACC(NONE)
PERMIT QMgrName.CHIN CLASS(MQCONN) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **GroupName**

The name of the group to be granted access.

#### **Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### **Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

#### **Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## **Granting full administrative access to all resources on a queue manager**

Grant full administrative access to all the resources on a queue manager, to each user or group of users with a business need for it.

### **About this task**

Use the Add Role Based Authorities wizard or the appropriate commands for your operating system.

### **Procedure**

- Using the wizard:
  1. In the WebSphere MQ Explorer Navigator pane, right-click the queue manager and click **Object Authorities > Add Role Based Authorities**. The Add Role Based Authorities wizard opens.

- For UNIX and Linux systems, issue the following commands:

```
setmqaut -m QMgrName -n '*' -t queue -g GroupName +alladm +browse
setmqaut -m QMgrName -n @class -t queue -g GroupName +crt
setmqaut -m QMgrName -n SYSTEM.ADMIN.COMMAND.QUEUE -t queue -g GroupName +dsp +inq +put
setmqaut -m QMgrName -n SYSTEM.MQEXPLORER.REPLY.QUEUE -t queue -g GroupName +dsp +inq +get
setmqaut -m QMgrName -n '*' -t topic -g GroupName +alladm
setmqaut -m QMgrName -n @class -t topic -g GroupName +crt
setmqaut -m QMgrName -n '*' -t channel -g GroupName +alladm
setmqaut -m QMgrName -n @class -t channel -g GroupName +crt
setmqaut -m QMgrName -n '*' -t clntconn -g GroupName +alladm
setmqaut -m QMgrName -n @class -t clntconn -g GroupName +crt
setmqaut -m QMgrName -n '*' -t authinfo -g GroupName +alladm
setmqaut -m QMgrName -n @class -t authinfo -g GroupName +crt
setmqaut -m QMgrName -n '*' -t listener -g GroupName +alladm
setmqaut -m QMgrName -n @class -t listener -g GroupName +crt
setmqaut -m QMgrName -n '*' -t namelist -g GroupName +alladm
setmqaut -m QMgrName -n @class -t namelist -g GroupName +crt
setmqaut -m QMgrName -n '*' -t process -g GroupName +alladm
setmqaut -m QMgrName -n @class -t process -g GroupName +crt
setmqaut -m QMgrName -n '*' -t service -g GroupName +alladm
setmqaut -m QMgrName -n @class -t service -g GroupName +crt
setmqaut -m QMgrName -t qmgr -g GroupName +alladm +conn
```

- For Windows systems, issue the same commands as for UNIX and Linux systems, but using the profile name @CLASS instead of @class.
- For IBM i, issue the following command:  
GRTMQMAUT OBJ(\*ALL) OBJTYPE(\*ALL) USER('GroupName') AUT(\*ALLADM)  
MQMNAME('QMgrName')
- For z/OS, issue the following commands:  
RDEFINE MQADMIN QMgrName.\*.\* UACC(NONE)  
PERMIT QMgrName.\*.\* CLASS(MQADMIN) ID(GroupName) ACCESS(ALTER)

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **GroupName**

The name of the group to be granted access.

#### **Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

#### **Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

#### **Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## **Removing connectivity to the queue manager**

If you do not want user applications to connect to your queue manager, remove their authority to connect to it.

### **About this task**

Revoke the authority of all users to connect to the queue manager by using the appropriate command for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -t qmgr -g GroupName -connect`
- For IBM i, issue the following command:  
`RVKMQMAUT OBJ ('QMgrName') OBJTYPE(*MQM) USER(*ALL) AUT(*CONNECT)`
- For z/OS, issue the following commands:  
`RDEFINE MQCONN QMgrName.BATCH UACC(NONE)`  
`RDEFINE MQCONN QMgrName.CHIN UACC(NONE)`  
`RDEFINE MQCONN QMgrName.CICS UACC(NONE)`  
`RDEFINE MQCONN QMgrName.IMS UACC(NONE)`

Do not issue any PERMIT commands. The variable names have the following meanings:

### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### GroupName

The name of the group to be denied access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:



[setmqaut \(WebSphere MQ V7.1 Reference\)](#)



[Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### Related information:



[Revoke MQ Object Authority \(RVKMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

## Allowing user applications to connect to your queue manager

You want to allow user application to connect to your queue manager. Use the tables in this topic to determine what actions to take.

First, determine whether client applications will connect to your queue manager.

If none of the applications that will connect to your queue manager are client applications, disable remote access as described in “Disabling remote access to the queue manager” on page 739.

If one or more of the applications that will connect to your queue manager are client applications, secure remote connectivity as described in “Securing remote connectivity to the queue manager” on page 732.

In both cases, set up connection security as described in “Setting up connection security” on page 740

If you want to control access to resources for each user connecting to the queue manager, see the following table. If the statement in the first column is true, take the action listed in the second column.

| Statement                                                      | Take this action                                                    |
|----------------------------------------------------------------|---------------------------------------------------------------------|
| You have applications that make use of queues                  | See “Controlling user access to queues” on page 741                 |
| You have applications that make use of topics                  | See “Controlling user access to topics” on page 747.                |
| You have applications that inquire on the queue manager object | See “Granting authority to inquire on a queue manager” on page 749. |
| You have applications that use process objects                 | See “Granting authority to access processes” on page 750            |
| You have applications that make use of namelists               | See “Granting authority to access namelists” on page 750            |

### Securing remote connectivity to the queue manager:

You can secure remote connectivity to the queue manager using SSL or TLS, a security exit, channel authentication records, or a combination of these methods.

#### About this task

You connect a client to the queue manager by using a client-connection channel on the client workstation and a server-connection channel on the server. Secure such connections in one of the following ways.

#### Procedure

1. Using SSL or TLS with channel authentication records:
  - a. Prevent any Distinguished Name (DN) from opening a channel, by using an SSLPEERMAP channel authentication record to map all DNs to USERSRC(NOACCESS).
  - b. Allow specific DNs or sets of DNs to open a channel by using an SSLPEERMAP channel authentication record to map them to USERSRC(CHANNEL).
2. Using SSL or TLS with a security exit:
  - a. Set MCAUSER on the server-connection channel to a user identifier with no privileges.
  - b. Write a security exit to assign an MCAUSER value depending on the value of SSL DN it receives in the SSLPeerNamePtr and SSLPeerNameLength fields passed to the exit in the MQCD structure.
3. Using SSL or TLS with fixed channel definition values:
  - a. Set SSLPEER on the server-connection channel to a specific value or narrow range of values.
  - b. Set MCAUSER on the server-connection channel to the user ID the channel should run with.
4. Using channel authentication records on channels that do not use SSL or TLS:
  - a. Prevent any IP address from opening channels, by using an address-mapping channel authentication record with ADDRESS(\*) and USERSRC(NOACCESS).
  - b. Allow specific IP addresses to open channels, by using address-mapping channel authentication records for those addresses with USERSRC(CHANNEL).
5. Using a security exit:
  - a. Write a security exit to authorize connections based on any property you choose, for example, the originating IP address.
6. It is also possible to use channel authentication records with a security exit, or to use all three methods, if your particular circumstances require it.

## Related concepts:

 WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*)

“Implementing access control in security exits” on page 758

“Link level security using a security exit” on page 438

 Writing a security exit (*WebSphere MQ V7.1 Programming Guide*)

## Related reference:

 ALTER CHANNEL (*WebSphere MQ V7.1 Reference*)

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

 DISPLAY CHLAUTH (*WebSphere MQ V7.1 Reference*)

## Blocking specific IP addresses:

You can prevent a specific channel accepting an inbound connection from an IP address, or prevent the whole queue manager from allowing access from an IP address, by using a channel authentication record.

## Before you begin

Enable channel authentication records by running the following command:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## About this task

To disallow specific channels from accepting an inbound connection and ensure that connections are only accepted when using the correct channel name, one type of rule can be used to block IP addresses. To disallow an IP address access to the whole queue manager, you would normally use a firewall to permanently block it.

However, another type of rule can be used to allow you to block a few addresses temporarily, for example while you are waiting for the firewall to be updated.

Note that this type of rule will hold the inbound connection open for 30 seconds before it is failed. To avoid this delay, and then only if the inbound connections to be blocked are channels, use the type of rule described in the following example.

This rule operates at the listener, before the channel name requested is known. You must therefore specify the channel name as an asterisk (\*), matching all channels.

You can specify the IP address in the channel authentication record singly, as a generic IP address pattern, or as a list of IP addresses.

## Procedure

To block IP addresses from using a specific channel, or queue manager, set a channel authentication record by using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**.

```
SET CHLAUTH(generic-channel-name) TYPE(AADDRESSMAP) ADDRESS(generic-ip-address)
USERSRC(NOACCESS)
```

There are three parts to the command:

## SET CHLAUTH (*generic-channel-name*)

You use this part of the command to control whether you want to block a connection for the entire queue manager, single channel or range of channels. What you put in here determines which areas are covered.

For example:

- SET CHLAUTH('\*') - blocks every channel on a queue manager, that is, the entire queue manager
- SET CHLAUTH('SYSTEM.\*') - blocks every channel that begins with SYSTEM.
- SET CHLAUTH('SYSTEM.DEF.SVRCONN') - blocks the channel SYSTEM.DEF.SVRCONN

## Type of CHLAUTH rule

Use this part of the command to specify the type of command and determines whether you want to supply a single address or list of addresses.

For example:

- TYPE (ADDRESSMAP) - Use ADDRESSMAP if you want to supply a single address or wild card address. For example, ADDRESS('192.168.\*') blocks any connections coming from an IP address starting in 192.168.

For more information about filtering IP addresses with patterns, see  [Generic IP addresses \(WebSphere MQ V7.1 Reference\)](#)

- TYPE (BLOCKADDR) - Use BLOCKADDR if you want to supply a list of address to block.

## Additional parameters

These parameters are dependent upon the type of rule you used in the second part of the command:

- For TYPE (ADDRESSMAP) you use ADDRESS
- For TYPE (BLOCKADDR) you use ADDRLIST

## Related tasks:

*“Temporarily blocking specific IP addresses if the queue manager is not running”*

## Related reference:

 [SET CHLAUTH \(WebSphere MQ V7.1 Reference\)](#)

*Temporarily blocking specific IP addresses if the queue manager is not running:*

You might want to block particular IP addresses, or ranges of addresses, when the queue manager is not running and you cannot therefore issue MQSC commands. You can temporarily block IP addresses on an exceptional basis by modifying the `blockaddr.ini` file.

## About this task

The `blockaddr.ini` file contains a copy of the BLOCKADDR definitions that are used by the queue manager. This file is read by the listener if the listener is started before the queue manager. In these circumstances, the listener uses any values that you have manually added to the `blockaddr.ini` file.

However, be aware that when the queue manager is started, it writes the set of BLOCKADDR definitions to the `blockaddr.ini` file, over-writing any manual editing you might have done. Similarly, every time you add or delete a BLOCKADDR definition by using the **SET CHLAUTH** command, the `blockaddr.ini` file is updated. You can therefore make permanent changes to the BLOCKADDR definitions only by using the **SET CHLAUTH** command when the queue manager is running.

## Procedure

1. Open the `blockaddr.ini` file in a text editor. The file is located in the data directory of the queue manager.

2. Add IP addresses as simple keyword-value pairs, where the keyword is Addr. For information about filtering IP addresses with patterns, see  Generic IP addresses (*WebSphere MQ V7.1 Reference*). For example:

```
Addr = 192.0.2.0
Addr = 192.0.*
Addr = 192.0.2.1-8
```

#### Related tasks:

“Blocking specific IP addresses” on page 733

#### Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

#### Blocking specific user IDs:

You can prevent specific users from using a channel by specifying user IDs that, if asserted, cause the channel to end. Do this by setting a channel authentication record.

#### Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

#### Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE(BLOCKUSER) USERLIST(userID1, userID2)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

The user list provided on a TYPE(BLOCKUSER) only applies to SVRCONN channels and not queue manager to queue manager channels.

*userID1* and *userID2* are each the ID of a user that is to be prevented from using the channel. You can also specify the special value \*MQADMIN to refer to privileged administrative users. For more information about privileged users, see “Privileged users” on page 692. For more information about

\*MQADMIN, see  SET CHLAUTH (*WebSphere MQ V7.1 Reference*).

#### Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

#### Mapping a remote queue manager to an MCAUSER user ID:

You can use a channel authentication record to set the MCAUSER attribute of a channel, according to the queue manager from which the channel is connecting.

#### Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

#### About this task

Optionally, you can restrict the IP addresses to which the rule applies.

Note that this technique does not apply to server-connection channels. If you specify the name of a server-connection channel in the commands shown below, it has no effect.

## Procedure

- Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE (QMGRMAP)
QMNAME(generic-partner-qmgr-name) USERSRC(MAP) MCAUSER(user)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*generic-partner-qmgr-name* is either the name of the queue manager, or a pattern including the asterisk (\*) symbol as a wildcard that matches the queue manager name.

*user* is the user ID to be used for all connections from the specified queue manager.

- To restrict this command to certain IP addresses, include the **ADDRESS** parameter, as follows:

```
SET CHLAUTH('generic-channel-name') TYPE (QMGRMAP)
QMNAME(generic-partner-qmgr-name) USERSRC(MAP)
MCAUSER(user) ADDRESS(generic-ip-address)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*generic-ip-address* is either a single address, or a pattern including the asterisk (\*) symbol as a wildcard or the hyphen (-) to indicate a range, that matches the address. For more information

about generic IP addresses, see  [Generic IP addresses \(WebSphere MQ V7.1 Reference\)](#).

## Related reference:

 [SET CHLAUTH \(WebSphere MQ V7.1 Reference\)](#)

*Mapping a client asserted user ID to an MCAUSER user ID:*

You can use a channel authentication record to change the MCAUSER attribute of a server-connection channel, according to the original user ID received from a client.

## Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## About this task

Note that this technique applies only to server-connection channels. It has no effect on other channel types.

## Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE (USERMAP) CLNTUSER(client-user-name)
USERSRC(MAP) MCAUSER(user)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*client-user-name* is the user ID asserted by the client.

*user* is the user ID to be used instead of the client user name.

## Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

*Mapping an SSL or TLS Distinguished Name to an MCAUSER user ID:*

You can use a channel authentication record to set the MCAUSER attribute of a channel, according to the Distinguished Name (DN) received.

## Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE (SSLPEERMAP) SSLPEER(generic-ssl-peer-name)
USERSRC(MAP) MCAUSER(user)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*generic-ssl-peer-name* is a string following the standard IBM WebSphere MQ rules for SSLPEER values.

See  WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*).

*user* is the user ID to be used for all connections using the specified DN.

## Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

*Blocking access from a remote queue manager:*

You can use a channel authentication record to prevent a remote queue manager from starting channels.

## Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## About this task

Note that this technique does not apply to server-connection channels. If you specify the name of a server-connection channel in the command shown below, it has no effect.

## Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE(QMGRMAP)
QMNAME('generic-partner-qmgr-name') USERSRC(NOACCESS)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*generic-partner-qmgr-name* is either the name of the queue manager, or a pattern including the asterisk (\*) symbol as a wildcard that matches the queue manager name.

## Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

*Blocking access for a client asserted user ID:*

You can use a channel authentication record to prevent a client asserted user ID from starting channels.

## Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## About this task

Note that this technique applies only to server-connection channels. It has no effect on other channel types.

## Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE(USERMAP) CLNTUSER('client-user-name')
USERSRC(NOACCESS)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*client-user-name* is the user ID asserted by the client.

## Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

*Blocking access for an SSL Distinguished Name:*

You can use a channel authentication record to prevent an SSL Distinguished Name from starting channels.

## Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

## Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE(SSLPEERMAP) SSLPEER('generic-ssl-peer-name')
USERSRC(NOACCESS)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*generic-ssl-peer-name* is a string following the standard IBM WebSphere MQ rules for SSLPEER values.

See  WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*).

### Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

*Mapping an IP address to an MCAUSER user ID:*

You can use a channel authentication record to set the MCAUSER attribute of a channel, according to the IP address from which the connection is received.

### Before you begin

Ensure that channel authentication records are enabled as follows:

```
ALTER QMGR CHLAUTH(ENABLED)
```

### Procedure

Set a channel authentication record using the MQSC command **SET CHLAUTH**, or the PCF command **Set Channel Authentication Record**. For example, you can issue the MQSC command:

```
SET CHLAUTH('generic-channel-name') TYPE(ADDRESSMAP) ADDRESS('generic-ip-address')
USERSRC(MAP) MCAUSER(user)
```

*generic-channel-name* is either the name of a channel to which you want to control access, or a pattern including the asterisk (\*) symbol as a wildcard that matches the channel name.

*user* is the user ID to be used for all connections using the specified DN.

*generic-ip-address* is either the address from which the connection is being made, or a pattern including the asterisk (\*) as a wildcard or the hyphen (-) to indicate a range, that matches the address.

### Related reference:

 SET CHLAUTH (*WebSphere MQ V7.1 Reference*)

### Related information:

 Generic IP addresses (*WebSphere MQ V7.1 Reference*)

### Disabling remote access to the queue manager:

If you do not want client applications to connect to your queue manager, disable remote access to it.

### About this task

Prevent client applications connecting to the queue manager in one of the following ways:

### Procedure

- Delete all server-connection channels using the MQSC command **DELETE CHANNEL**.
- Set the message channel agent user identifier (MCAUSER) of the channel to a user ID with no access rights, using the MQSC command **ALTER CHANNEL**.

**Related reference:**

 ALTER CHANNEL (*WebSphere MQ V7.1 Reference*)

 DELETE CHANNEL (*WebSphere MQ V7.1 Reference*)

**Setting up connection security:**

Grant the authority to connect to the queue manager to each user or group of users with a business need to do so.

**About this task**

To set up connection security, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:

```
setmqaut -m QMgrName -t qmgr -g GroupName +connect
```

- For IBM i, issue the following command:

```
GRTMQMAUT OBJ('QMgrName') OBJTYPE(*MQM) USER('GroupName') AUT(*CONNECT)
```

- For z/OS, issue the following commands:

```
RDEFINE MQCONN QMgrName.BATCH UACC(NONE)
PERMIT QMgrName.BATCH CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.CICS UACC(NONE)
PERMIT QMgrName.CICS CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.IMS UACC(NONE)
PERMIT QMgrName.IMS CLASS(MQCONN) ID(GroupName) ACCESS(READ)
RDEFINE MQCONN QMgrName.CHIN UACC(NONE)
PERMIT QMgrName.CHIN CLASS(MQCONN) ID(GroupName) ACCESS(READ)
```

These commands give authority to connect for batch, CICS, IMS and the channel initiator (CHIN). If you do not use a particular type of connection, omit the relevant commands. The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 setmqaut (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

**Related information:**

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

## Controlling user access to queues:

You want to control application access to queues. Use this topic to determine what actions to take.

For each true statement in the first column, take the action indicated in the second column.

| Statement                                          | Action                                                                         |
|----------------------------------------------------|--------------------------------------------------------------------------------|
| The application gets messages from a queue         | See "Granting authority to get messages from queues"                           |
| The application sets context                       | See "Granting authority to set context" on page 742                            |
| The application passes context                     | See "Granting authority to pass context" on page 743                           |
| The application puts messages on a clustered queue | See "Authorizing putting messages on remote cluster queues" on page 802        |
| The application puts messages on a local queue     | See "Granting authority to put messages to a local queue" on page 744          |
| The application puts messages on a model queue     | See "Granting authority to put messages to a model queue" on page 745          |
| The application puts messages on a remote queue    | See "Granting authority to put messages to a remote cluster queue" on page 746 |

### *Granting authority to get messages from queues:*

Grant the authority to get messages from a queue or set of queues, to each group of users with a business need for it.

### About this task

To grant the authority to get messages from some queues, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +get`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*GET)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)`

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

## Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

## Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

## Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

*Granting authority to set context:*

Grant the authority to set context on a message that is being put, to each group of users with a business need for it.

## About this task

To grant the authority to set context on some queues, use the appropriate commands for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue one of the following commands:
  - To set identity context only:

```
setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +setid
```
  - To set all context:

```
setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +setall
```
- For IBM i, issue one of the following commands:
  - To set identity context only:

```
GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*SETID) MQMNAME('QMgrName')
```
  - To set all context:

```
GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*SETALL) MQMNAME('QMgrName')
```
- For z/OS, issue one of the following sets of commands:
  - To set identity context only:

```
RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)
```
  - To set all context:

```
RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(CONTROL)
```

The variable names have the following meanings:

### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### ObjectProfile

The name of the object or generic profile for which to change authorizations.

## GroupName

The name of the group to be granted access.

### Related concepts:

“Profiles for context security” on page 549

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:



setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

### Related information:



Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

*Granting authority to pass context:*

Grant the authority to pass context from a retrieved message to one that is being put, to each group of users with a business need for it.

### About this task

To grant the authority to pass context on some queues, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue one of the following commands:
  - To pass identity context only:  
`setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +passid`
  - To pass all context:  
`setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +passall`
- For IBM i, issue one of the following commands:
  - To pass identity context only:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*PASSID)  
MQMNAME('QMgrName')`
  - To pass all context:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*PASSALL)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands to pass identity context or all context:  
`RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)`

The variable names have the following meanings:

#### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### ObjectProfile

The name of the object or generic profile for which to change authorizations.

#### GroupName

The name of the group to be granted access.

**Related concepts:**

“Profiles for context security” on page 549

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

**Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

*Granting authority to put messages to a local queue:*

Grant the authority to put messages to a local queue or set of queues, to each group of users with a business need for it.

**About this task**

To grant the authority to put messages to some local queues, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +put`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*PUT)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)`

The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

**Related information:**

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

*Granting authority to put messages to a model queue:*

Grant the authority to put messages to a model queue or set of model queues, to each group of users with a business need for it.

### About this task

Model queues are used to create dynamic queues. You must therefore grant authority to both the model and dynamic queues. To grant these authorities, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue the following commands:

```
setmqaut -m QMgrName -n ModelQueueName -t queue -g GroupName +put
setmqaut -m QMgrName -n ObjectProfile -t queue -g GroupName +put
```

- For IBM i, issue the following commands:

```
GRTMQMAUT OBJ('ModelQueueName') OBJTYPE(*Q) USER(GroupName) AUT(*PUT)
MQMNAME('QMgrName')
GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*Q) USER(GroupName) AUT(*PUT)
MQMNAME('QMgrName')
```

- For z/OS, issue the following commands:

```
RDEFINE MQQUEUE QMgrName.ModelQueueName UACC(NONE)
PERMIT QMgrName.ModelQueueName CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)
RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ModelQueueName**

The name of the model queue on which dynamic queues are based.

#### **ObjectProfile**

The name of the dynamic queue or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

*Granting authority to put messages to a remote cluster queue:*

Grant the authority to put messages to a remote cluster queue or set of queues, to each group of users with a business need for it.

### About this task

To put a message on a remote cluster queue, you can either put it on a local definition of a remote queue, or a fully qualified remote queue. If you are using a local definition of a remote queue, you need authority to put to the local object: see “Granting authority to put messages to a local queue” on page 744. If you are using a fully qualified remote queue, you need authority to put to the remote queue. Grant this authority using the appropriate commands for your operating system.

The default behavior is to perform access control against the `SYSTEM.CLUSTER.TRANSMIT.QUEUE`. Note that this behavior applies, even if you are using multiple transmission queues.

The specific behavior described in this topic applies only when you have configured the

**ClusterQueueAccessControl** attribute in the `qm.ini` file to be *RQMName*, as described in the  Security stanza (*WebSphere MQ V7.1 Installing Guide*) topic, and restarted the queue manager.

On UNIX, UNIX, Windows systems, and IBM i, you can also use the `SET AUTHREC` command.

### Procedure

- For UNIX, UNIX and Windows systems, issue the following command:

```
setmqaut -m QMgrName -t rqmname -n ObjectProfile -g GroupName +put
```

Note that you can use the *rqmname* object for remote cluster queues only.

- For IBM i, issue the following command:

```
GRTMQMAUT OBJTYPE(*RMTMQMNAME) OBJ('ObjectProfile') USER(GroupName) AUT(*PUT)
MQMNAME('QMgrName')
```

Note that you can use the `RMTMQMNAME` object for remote cluster queues only.

- For z/OS, issue the following commands:

```
RDEFINE MQQUEUE QMgrName.ObjectProfile UACC(NONE)
PERMIT QMgrName.ObjectProfile CLASS(MQQUEUE) ID(GroupName) ACCESS(UPDATE)
```

Note that you can use the name of the remote queue manager (or queue-sharing group) for remote cluster queues only. The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the remote queue manager or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

### Related concepts:

“Profiles for queue security” on page 533

“Profiles to control queue-sharing group or queue manager level security” on page 525

 Opening remote queues (*WebSphere MQ V7.1 Programming Guide*)

### Related reference:

 setmqaut (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

### Related information:

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

### Controlling user access to topics:

You need to control the access of applications to topics. Use this topic to determine what actions to take.

For each true statement in the first column, take the action indicated in the second column.

| Statement                                     | Action                                                      |
|-----------------------------------------------|-------------------------------------------------------------|
| The application publishes messages to a topic | See “Granting authority to publish messages to a topic”     |
| The application subscribes to a topic         | See “Granting authority to subscribe to topics” on page 748 |

*Granting authority to publish messages to a topic:*

Grant the authority to publish messages to a topic or set of topics, to each group of users with a business need for it.

### About this task

To grant the authority to publish messages to some topics, use the appropriate commands for your operating system.

### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t topic -g GroupName +pub`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*TOPIC) USER(GroupName) AUT(*PUB)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQTOPIC QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQTOPIC) ID(GroupName) ACCESS(UPDATE)`

The variable names have the following meanings:

#### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### ObjectProfile

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

**Related concepts:**

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

**Related reference:**

setmqaut (*WebSphere MQ V7.1 Reference*)



Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

**Related information:**

Grant MQ Object Authority (GRMQMAUT) (*WebSphere MQ V7.1 Reference*)

*Granting authority to subscribe to topics:*

Grant the authority to subscribe to a topic or set of topics, to each group of users with a business need for it.

**About this task**

To grant the authority to subscribe to some topics, use the appropriate commands for your operating system.

**Procedure**

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t topic -g GroupName +sub`
- For IBM i, issue the following command:  
`GRMQMAUT OBJ('ObjectProfile') OBJTYPE(*TOPIC) USER(GroupName) AUT(*SUB)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQTOPIC QMgrName.SUBSCRIBE.ObjectProfile UACC(NONE)  
PERMIT QMgrName.SUBSCRIBE.ObjectProfile CLASS(MQTOPIC) ID(GroupName) ACCESS(UPDATE)`

The variable names have the following meanings:

**QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

**ObjectProfile**

The name of the object or generic profile for which to change authorizations.

**GroupName**

The name of the group to be granted access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

### Granting authority to inquire on a queue manager:

Grant the authority to inquire on a queue manager, to each group of users with a business need for it.

#### About this task

To grant the authority to inquire on a queue manager, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t qmgr -g GroupName +inq`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*MQM) USER(GroupName) AUT(*INQ)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQCMDS QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQCMDS) ID(GroupName ) ACCESS(READ)`

These commands grant access to the specified queue manager. To permit the user to use the MQINQ command, issue the following commands:

```
RDEFINE MQCMDS QMgrName.MQINQ.QMGR UACC(NONE)
PERMIT QMgrName.MQINQ.QMGR CLASS(MQCMDS) ID(GroupName) ACCESS(READ)
```

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

### Related concepts:

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:

 [setmqaut \(WebSphere MQ V7.1 Reference\)](#)

 [Security controls and options \(WebSphere MQ V7.1 Product Overview Guide\)](#)

### Related information:

 [Grant MQ Object Authority \(GRTMQMAUT\) \(WebSphere MQ V7.1 Reference\)](#)

### Granting authority to access processes:

Grant the authority to access a process or set of processes, to each group of users with a business need for it.

#### About this task

To grant the authority to access some processes, use the appropriate commands for your operating system.

#### Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t process -g GroupName +all`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*PRC) USER(GroupName) AUT(*ALL)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQPROC QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQPROC) ID(GroupName) ACCESS(READ)`

The variable names have the following meanings:

#### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

#### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

#### **GroupName**

The name of the group to be granted access.

### Granting authority to access namelists:

Grant the authority to access a namelist or set of namelists, to each group of users with a business need for it.

#### About this task

To grant the authority to access some namelists, use the appropriate commands for your operating system.

## Procedure

- For UNIX, Linux and Windows systems, issue the following command:  
`setmqaut -m QMgrName -n ObjectProfile -t namelist -g GroupName +all`
- For IBM i, issue the following command:  
`GRTMQMAUT OBJ('ObjectProfile') OBJTYPE(*NMLIST) USER(GroupName) AUT(*ALL)  
MQMNAME('QMgrName')`
- For z/OS, issue the following commands:  
`RDEFINE MQNLIST QMgrName.ObjectProfile UACC(NONE)  
PERMIT QMgrName.ObjectProfile CLASS(MQNLIST) ID(GroupName) ACCESS(READ)`

The variable names have the following meanings:

### **QMgrName**

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### **ObjectProfile**

The name of the object or generic profile for which to change authorizations.

### **GroupName**

The name of the group to be granted access.

## Authority to administer WebSphere MQ on UNIX, Linux and Windows systems

WebSphere MQ administrators can use all WebSphere MQ commands and grant authorities for other users. When administrators issue commands to remote queue managers, they must have the required authority on the remote queue manager. Further considerations apply to Windows systems.

WebSphere MQ administrators have authority to use all WebSphere MQ commands (including the commands to grant WebSphere MQ authorities for other users)

To be a WebSphere MQ administrator, you must be a member of a special group called the *mqm* group (or a member of the Administrators group on Windows systems). The *mqm* group is created automatically when WebSphere MQ is installed; add further users to the group to allow them to perform administration. All members of this group have access to all resources. This access can be revoked only by removing a user from the *mqm* group and issuing the REFRESH SECURITY command. Administrators can use control commands to administer WebSphere MQ. One of these control commands is **setmqaut**, which is used to grant authorities to other users to enable them to access or control WebSphere MQ resources. The PCF commands for managing authority records are available to non-administrators who have been granted dsp and chg authorities on the queue manager. For more information about managing authorities using PCF commands, see Programmable Command Formats.

Administrators can use the control command **runmqsc** to issue WebSphere MQ Script (MQSC) commands. When **runmqsc** is used in indirect mode to send MQSC commands to a remote queue manager, each MQSC command is encapsulated within an Escape PCF command. Administrators must have the required authorities for the MQSC commands to be processed by the remote queue manager. The WebSphere MQ Explorer issues PCF commands to perform administration tasks. Administrators require no additional authorities to use the WebSphere MQ Explorer to administer a queue manager on the local system. When the WebSphere MQ Explorer is used to administer a queue manager on another system, administrators must have the required authorities for the PCF commands to be processed by the remote queue manager.

For more information about authority checks when PCF and MQSC commands are processed, see the following topics:

- For PCF commands that operate on queue managers, queues, processes, namelists, and authentication information objects, see Authority to work with WebSphere MQ objects. Refer to this section for the equivalent MQSC commands encapsulated within Escape PCF commands.
- For PCF commands that operate on channels, channel initiators, listeners, and clusters, see  Channel security. Refer to this section for the equivalent MQSC commands encapsulated within Escape PCF commands.
- For PCF commands that operate on authority records, see Authority checking for PCF commands
- For MQSC commands that are processed by the command server on WebSphere MQ for z/OS, see Command security and command resource security on z/OS.

Additionally, on Windows systems, the SYSTEM account has full access to WebSphere MQ resources.

On UNIX and Linux platforms, a special user ID of mqm is also created, for use by the product only. It must never be available to non-privileged users. All WebSphere MQ objects are owned by user ID mqm.

On Windows systems, members of the Administrators group can also administer any queue manager, as can the SYSTEM account. You can also create a domain mqm group on the domain controller that contains all privileged user IDs active within the domain, and add it to the local mqm group. Some commands, for example **crtmqm**, manipulate authorities on WebSphere MQ objects and so need authority to work with these objects (as described in the following sections). Members of the mqm group have authority to work with all objects, but there might be circumstances on Windows systems when authority is denied if you have a local user and a domain-authenticated user with the same name. This is described in “Principals and groups” on page 756.

Windows versions with a User Account Control (UAC) feature restricts the actions users can perform on certain operating system facilities, even if they are members of the Administrators group. If your userid is in the Administrators group but not the mqm group you must use an elevated command prompt to issue WebSphere MQ admin commands such as **crtmqm**, otherwise the error "AMQ7077: You are not authorized to perform the requested operation" is generated. To open an elevated command prompt, right-click the start menu item, or icon, for the command prompt, and select "Run as administrator".

You do not need to be a member of the mqm group to do the following:

- Issue commands from an application program that issues PCF commands, or MQSC commands within an Escape PCF command, unless the commands manipulate channel initiators. (These commands are described in “Protecting channel initiator definitions” on page 448).
- Issue MQI calls from an application program (unless you want to use the fast path bindings on the **MQCONN** call).
- Use the **crtmqcvx** command to create a fragment of code that performs data conversion on data type structures.
- Use the **dspmq** command to display queue managers.
- Use the **dspmqtrc** command to display WebSphere MQ formatted trace output.

A 12 character limitation applies to both group and user IDs.

UNIX and Linux platforms generally restrict the length of a user ID to 12 characters. AIX Version 5.3 has raised this limit but WebSphere MQ continues to observe a 12 character restriction on all UNIX and Linux platforms. If you use a user ID of greater than 12 characters, WebSphere MQ replaces it with the value UNKNOWN. Do not define a user ID with a value of UNKNOWN.

## Related concepts:

“Authority to work with WebSphere MQ objects on UNIX, Linux and Windows systems”

## Managing the mqm group

Users in the mqm group are granted full administrative privileges over WebSphere MQ. For this reason, you should not enrol applications and ordinary users in the mqm group. The mqm group should contain the accounts of the WebSphere MQ administrators only.

These tasks are described in:

- Creating and managing groups on Windows
- Creating and managing groups on HP-UX
- Creating and managing groups on AIX
- Creating and managing groups on Solaris
- Creating and managing groups on Linux

If your domain controller runs on Windows 2000 or Windows 2003, your domain administrator might have to set up a special account for WebSphere MQ to use. This is described in the  [Configuring WebSphere MQ accounts \(WebSphere MQ V7.1 Installing Guide\)](#).

## Authority to work with WebSphere MQ objects on UNIX, Linux and Windows systems

All objects are protected by WebSphere MQ, and principals must be given appropriate authority to access them. Different principals need different access rights to different objects.

Queue managers, queues, process definitions, namelists, channels, client connection channels, listeners, services, and authentication information objects are all accessed from applications that use MQI calls or PCF commands. These resources are all protected by WebSphere MQ, and applications need to be given permission to access them. The entity making the request might be a user, an application program that issues an MQI call, or an administration program that issues a PCF command. The identifier of the requester is referred to as the *principal*.

Different groups of principals can be granted different types of access authority to the same object. For example, for a specific queue, one group might be allowed to perform both put and get operations; another group might be allowed only to browse the queue (**MQGET** with browse option). Similarly, some groups might have put and get authority to a queue, but not be allowed to alter attributes of the queue or delete it.

Some operations are particularly sensitive and should be limited to privileged users. For example:

- Accessing some special queues, such as transmission queues or the command queue `SYSTEM.ADMIN.COMMAND.QUEUE`
- Running programs that use full MQI context options
- Creating and deleting application queues

Full access permission to an object is automatically given to the user ID that created the object and to all members of the mqm group (and to the members of the local Administrators group on Windows systems).

## When security checks are made on UNIX, Linux and Windows systems

Security checks are typically made on connecting to a queue manager, opening or closing objects, and putting or getting messages.

The security checks made for a typical application are as follows:

### Connecting to the queue manager (MQCONN or MQCONNX calls)

This is the first time that the application is associated with a particular queue manager. The queue manager interrogates the operating environment to discover the user ID associated with the application. WebSphere MQ then verifies that the user ID is authorized to connect to the queue manager and retains the user ID for future checks.

Users do not have to sign on to WebSphere MQ; WebSphere MQ assumes that users have signed on to the underlying operating system and have been authenticated by that.

### Opening the object (MQOPEN or MQPUT1 calls)

WebSphere MQ objects are accessed by opening the object and issuing commands against it. All resource checks are performed when the object is opened, rather than when it is actually accessed. This means that the **MQOPEN** request must specify the type of access required (for example, whether the user wants only to browse the object or perform an update like putting messages onto a queue).

WebSphere MQ checks the resource that is named in the **MQOPEN** request. For an alias or remote queue object, the authorization used is that of the object itself, not the queue to which the alias or remote queue resolves. This means that the user does not need permission to access it. Limit the authority to create queues to privileged users. If you do not, users might bypass the normal access control simply by creating an alias. If a remote queue is referred to explicitly with both the queue and queue manager names, the transmission queue associated with the remote queue manager is checked.

The authority to a dynamic queue is based on that of the model queue from which it is derived, but is not necessarily the same. This is described in Note 1 on page 474.

The user ID used by the queue manager for access checks is the user ID obtained from the operating environment of the application connected to the queue manager. A suitably authorized application can issue an **MQOPEN** call specifying an alternative user ID; access control checks are then made on the alternative user ID. This does not change the user ID associated with the application, only that used for access control checks.

### Putting and getting messages (MQPUT or MQGET calls)

No access control checks are performed.

### Closing the object (MQCLOSE)

No access control checks are performed, unless the **MQCLOSE** results in a dynamic queue being deleted. In this case, there is a check that the user ID is authorized to delete the queue.

### Subscribing to a topic (MQSUB)

When an application subscribes to a topic, it specifies the type of operation that it needs to perform. It is either creating a new subscription, altering an existing subscription, or resuming an existing subscription without changing it. For each type of operation, the queue manager checks that the user ID that is associated with the application has the authority to perform the operation.

When an application subscribes to a topic, the authority checks are performed against the topic objects that are found in the topic tree at, or above, the point in the topic tree at which the application subscribed. The authority checks might involve checks on more than one topic object.

The user ID that the queue manager uses for the authority checks is the user ID obtained from the operating system when the application connects to the queue manager.

The queue manager performs authority checks on subscriber queues but not on managed queues.

## How access control is implemented by WebSphere MQ on UNIX, Linux and Windows systems

WebSphere MQ uses the security services provided by the underlying operating system, using the object authority manager. WebSphere MQ supplies commands to create and maintain access control lists.

An access control interface called the Authorization Service Interface is part of WebSphere MQ. WebSphere MQ supplies an implementation of an access control manager (conforming to the Authorization Service Interface) known as the *object authority manager (OAM)*. This is automatically installed and enabled for each queue manager you create, unless you specify otherwise (as described in “Preventing security access checks on Windows, UNIX and Linux systems” on page 712). The OAM can be replaced by any user or vendor written component that conforms to the Authorization Service Interface.

The OAM exploits the security features of the underlying operating system, using operating system user and group IDs. Users can access WebSphere MQ objects only if they have the correct authority. “Controlling access to objects by using the OAM on UNIX, Linux and Windows systems” on page 704 describes how to grant and revoke this authority.

The OAM maintains an access control list (ACL) for each resource that it controls. Authorization data is stored on a local queue called SYSTEM.AUTH.DATA.QUEUE. Access to this queue is restricted to users in the mqm group, and additionally on Windows, to users in the Administrators group, and users logged in with the SYSTEM ID. User access to the queue cannot be changed.

WebSphere MQ supplies commands to create and maintain access control lists. For more information on these commands, see “Controlling access to objects by using the OAM on UNIX, Linux and Windows systems” on page 704.

WebSphere MQ passes the OAM a request containing a principal, a resource name, and an access type. The OAM grants or rejects access based on the ACL that it maintains. WebSphere MQ follows the decision of the OAM; if the OAM cannot make a decision, WebSphere MQ does not allow access.

### Identifying the user ID on Windows, UNIX and Linux systems

The object authority manager identifies the principal that is requesting access to a resource. The user ID used as the principal varies according to context.

The object authority manager (OAM) must be able to identify who is requesting access to a particular resource. WebSphere MQ uses the term *principal* to refer to this identifier. The principal is established when the application first connects to the queue manager; it is determined by the queue manager from the user ID associated with the connecting application. (If the application issues XA calls without connecting to the queue manager, then the user ID associated with the application that issues the xa\_open call is used for authority checks by the queue manager.)

On UNIX and Linux systems, the authorization routines checks either the real (logged-in) user ID, or the effective user ID associated with the application. The user ID checked can be dependent on the bind type, for details see  Installable services (*WebSphere MQ V7.1 Installing Guide*).

WebSphere MQ propagates the user ID received from the system in the message header (MQMD structure) of each message as identification of the user. This identifier is part of the message context information and is described in “Context authority on UNIX, Linux and Windows systems” on page 758. Applications cannot alter this information unless they have been authorized to change context information.

## Principals and groups:

Principals can belong to groups. You can grant access to a particular resource to groups rather than to individuals, to reduce the amount of administration required. On UNIX and Linux systems all Access Control Lists (ACLs) are based on groups, but on Windows systems, ACLs are based on user IDs and groups.

For example, you might define a group consisting of users who want to run a particular application. Other users can be given access to all the resources they require by adding their user ID to the appropriate group. This process is described in:

- Creating and managing groups on Windows
- Creating and managing groups on HP-UX
- Creating and managing groups on AIX
- Creating and managing groups on Solaris
- Creating and managing groups on Linux

A principal can belong to more than one group (its group set). It has the aggregate of all the authorities granted to each group in its group set. These authorities are cached, so any changes you make to the group membership of the principal are not recognized until the queue manager is restarted, unless you issue the MQSC command REFRESH SECURITY (or the PCF equivalent).

### UNIX and Linux systems

All ACLs are based on groups. When a user is granted access to a particular resource, the primary group of the user ID is included in the ACL. The individual user ID is not included and authority is granted to all members of that group. Because of this, be aware that you can inadvertently change the authority of a principal by changing the authority of another principal in the same group.

All users are nominally assigned to the default user group *nobody* and by default, no authorizations are given to this group. You can change the authorization in the *nobody* group to grant access to WebSphere MQ resources to users without specific authorizations.

Do not define a user ID with the value "UNKNOWN". The value "UNKNOWN" is used when a user ID is too long, so arbitrary user IDs would use the access authorities of UNKNOWN.

User IDs can contain up to 12 characters and group names up to 12 characters.

### Windows systems

ACLs are based on both user IDs and groups. Checks are the same as for UNIX systems except that individual user IDs can be displayed in the ACL as well. You can have different users on different domains with the same user ID. WebSphere MQ permits user IDs to be qualified by a domain name so that these users can be given different levels of access.

The group name can optionally include a domain name, specified in the following formats:

GroupName@domain  
domain\GroupName

Global groups are checked by the OAM in two cases only:

1. The queue manager security stanza includes the setting: GroupModel=GlobalGroups; see

 Security (*WebSphere MQ V7.1 Installing Guide*).

2. The queue manager is using an alternate security access group; see  crtmqm (*WebSphere MQ V7.1 Reference*).

User IDs can contain up to 20 characters, domain names up to 15 characters, and group names up to 64 characters.

The OAM first checks the local security database, then the database of the primary domain, and finally the database of any trusted domains. The first user ID encountered is used by the OAM for checking. Each of these user IDs might have different group memberships on a particular computer.

Some control commands (for example, **crtmqm**) change authorities on WebSphere MQ objects using the object authority manager (OAM). The OAM searches the security databases in the order given in the preceding paragraph to determine the authority rights for a particular user ID. As a result, the authority determined by the OAM might override the fact that a user ID is a member of the local mqm group. For example, if you issue the **crtmqm** command from a user ID authenticated by a domain controller that has membership of the local mqm group through a global group, the command fails if the system has a local user of the same name who is not in the local mqm group.

### Windows security identifiers (SIDs):

WebSphere MQ on Windows uses the SID where it is available. If a Windows SID is not supplied with an authorization request, WebSphere MQ identifies the user based on the user name alone, but this might result in the wrong authority being granted.

On Windows systems, the security identifier (SID) is used to supplement the user ID. The SID contains information that identifies the full user account details on the Windows security account manager (SAM) database where the user is defined. When a message is created on WebSphere MQ for Windows, WebSphere MQ stores the SID in the message descriptor. When WebSphere MQ on Windows performs authorization checks, it uses the SID to query the full information from the SAM database. (The SAM database in which the user is defined must be accessible for this query to succeed.)

By default, if a Windows SID is not supplied with an authorization request, WebSphere MQ identifies the user based on the user name alone. It does this by searching the security databases in the following order:

1. The local security database
2. The security database of the primary domain
3. The security database of trusted domains

If the user name is not unique, incorrect WebSphere MQ authority might be granted. To prevent this problem, include an SID in each authorization request; the SID is used by WebSphere MQ to establish user credentials.

To specify that all authorization requests must include an SID, use **regedit**. Set the SecurityPolicy to NTSIDsRequired.

### Alternate-user authority on UNIX, Linux and Windows systems

You can specify that a user ID can use the authority of another user when accessing a WebSphere MQ object. This is called *alternate-user authority*, and you can use it on any WebSphere MQ object.

Alternate-user authority is essential where a server receives requests from a program and wants to ensure that the program has the required authority for the request. The server might have the required authority, but it needs to know whether the program has the authority for the actions it has requested.

For example, assume that a server program running under user ID PAYSERV retrieves a request message from a queue that was put on the queue by user ID USER1. When the server program gets the request message, it processes the request and puts the reply back into the reply-to queue specified with the request message. Instead of using its own user ID (PAYSERV) to authorize opening the reply-to queue, the server can specify a different user ID, in this case, USER1. In this example, you can use alternate-user authority to control whether PAYSERV is allowed to specify USER1 as an alternate-user ID when it opens the reply-to queue.

The alternate-user ID is specified on the **AlternateUserId** field of the object descriptor.

## Context authority on UNIX, Linux and Windows systems

Context is information that applies to a particular message and is contained in the message descriptor, MQMD, which is part of the message. Applications can specify the context data when either an MQOPEN or MQPUT call is made.

The context information comes in two sections:

### Identity section

Who the message came from. It consists of the `UserIdentifier`, `AccountingToken`, and `AppIdentityData` fields.

### Origin section

Where the message came from, and when it was put onto the queue. It consists of the `PutAppType`, `PutAppName`, `PutDate`, `PutTime`, and `AppOriginData` fields.

Applications can specify the context data when either an MQOPEN or MQPUT call is made. This data might be generated by the application, passed on from another message, or generated by the queue manager by default. For example, context data can be used by server programs to check the identity of the requester, testing whether the message came from an application running under an authorized user ID.

A server program can use the `UserIdentifier` to determine the user ID of an alternative user. You use context authorization to control whether the user can specify any of the context options on any MQOPEN or MQPUT1 call.

See  [Controlling context information \(WebSphere MQ V7.1 Programming Guide\)](#) for information about the context options, and  [Overview for MQMD \(WebSphere MQ V7.1 Reference\)](#) for descriptions of the message descriptor fields relating to context.

## Implementing access control in security exits

You can implement access control in a security exit by use of the `MCAUserIdentifier` or the object authority manager.

### MCAUserIdentifier

Every instance of a channel that is current has an associated channel definition structure, MQCD. The initial values of the fields in MQCD are determined by the channel definition that is created by a WebSphere MQ administrator. In particular, the initial value of one of the fields, `MCAUserIdentifier`, is determined by the value of the `MCAUSER` parameter on the `DEFINE CHANNEL` command, or by the equivalent to `MCAUSER` if the channel definition is created in another way. `MCAUserIdentifier` contains the first 12 bytes of the MCA user identifier. If the MCA user identifier is not blank, it specifies the user identifier to be used by the message channel agent for authorization to access MQ resources. Ensure that the `MCAUSER` is less than 12 characters on Windows platform.

The MQCD structure is passed to a channel exit program when it is called by an MCA. When a security exit is called by an MCA, the security exit can change the value of `MCAUserIdentifier`, replacing any value that was specified in the channel definition.

On IBM i, UNIX, Linux and Windows systems, unless the value of `MCAUserIdentifier` is blank, the queue manager uses the value of `MCAUserIdentifier` as the user ID for authority checks when an MCA attempts to access the queue manager's resources after it has connected to the queue manager. If the value of `MCAUserIdentifier` is blank, the queue manager uses the default user ID of the MCA instead. This applies

to RCVR, RQSTR, CLUSRCVR and SVRCONN channels. For sending MCAs, the default user ID is always used for authority checks, even if the value of *MCAUserIdentifier* is not blank.

On z/OS, the queue manager might use the value of *MCAUserIdentifier* for authority checks, provided it is not blank. For receiving MCAs and server connection MCAs, whether the queue manager uses the value of *MCAUserIdentifier* for authority checks depends on:

- The value of the PUTAUT parameter in the channel definition
- The RACF profile used for the checks
- The access level of the channel initiator address space user ID to the RESLEVEL profile

For sending MCAs, it depends on:

- Whether the sending MCA is a caller or a responder
- The access level of the channel initiator address space user ID to the RESLEVEL profile

The user ID that a security exit stores in *MCAUserIdentifier* can be acquired in various ways. Here are some examples:

- Provided there is no security exit at the client end of an MQI channel, a user ID associated with the WebSphere MQ client application flows from the client connection MCA to the server connection MCA when the client application issues an MQCONN call. The server connection MCA stores this user ID in the *RemoteUserIdentifier* field in the channel definition structure, MQCD. If the value of *MCAUserIdentifier* is blank at this time, the MCA stores the same user ID in *MCAUserIdentifier*. If the MCA does not store the user ID in *MCAUserIdentifier*, a security exit can do it later by setting *MCAUserIdentifier* to the value of *RemoteUserIdentifier*.

If the user ID that flows from the client system is entering a new security domain and is not valid on the server system, the security exit can substitute the user ID for one that is valid and store the substituted user ID in *MCAUserIdentifier*.

- The user ID can be sent by the partner security exit in a security message.

On a message channel, a security exit called by the sending MCA can send the user ID under which the sending MCA is running. A security exit called by the receiving MCA can then store the user ID in *MCAUserIdentifier*. Similarly, on an MQI channel, a security exit at the client end of the channel can send the user ID associated with the WebSphere MQ MQI client application. A security exit at the server end of the channel can then store the user ID in *MCAUserIdentifier*. As in the previous example, if the user ID is not valid on the target system, the security exit can substitute the user ID for one that is valid and store the substituted user ID in *MCAUserIdentifier*.

If a digital certificate is received as part of the identification and authentication service, a security exit can map the Distinguished Name in the certificate to a user ID that is valid on the target system. It can then store the user ID in *MCAUserIdentifier*.

- If SSL is used on the channel, the partner's Distinguished Name (DN) is passed to the exit in the *SSLPeerNamePtr* field of the MQCD, and the DN of the issuer of that certificate is passed to the exit in the *SSLRemCertIssNamePtr* field of the MQCXP.

For more information about the *MCAUserIdentifier* field, the channel definition structure, MQCD, and the channel exit parameter structure, MQCXP, see  Channel-exit calls and data structures (*WebSphere MQ V7.1 Reference*). For more information about the user ID that flows from a client system on an MQI channel, see Access control.

**Note:** Security exit applications constructed prior to the release of WebSphere MQ v7.1 may require updating. For more information see  Channel security exit programs (*WebSphere MQ V7.1 Programming Guide*).

## WebSphere MQ object authority manager user authentication

On WebSphere MQ MQI client connections, security exits can be used to modify or create the MQCSP structure used in object authority manager (OAM) user authentication. This is described in

 Channel-exit programs for messaging channels (*WebSphere MQ V7.1 Programming Guide*)

## Implementing access control in message exits

You might need to use a message exit to substitute one user ID with another.

Consider a client application that sends a message to a server application. The server application can extract the user ID from the *UserIdentifier* field in the message descriptor and, provided it has alternate user authority, ask the queue manager to use this user ID for authority checks when it accesses WebSphere MQ resources on behalf of the client.

If the PUTAUT parameter is set to CTX (or ALTMCA on z/OS) in the channel definition, the user ID in the *UserIdentifier* field of each incoming message is used for authority checks when the MCA opens the destination queue.

In certain circumstances, when a report message is generated, it is put using the authority of the user ID in the *UserIdentifier* field of the message causing the report. In particular, confirm-on-delivery (COD) reports and expiration reports are always put with this authority.

Because of these situations, it might be necessary to substitute one user ID for another in the *UserIdentifier* field as a message enters a new security domain. This can be done by a message exit at the receiving end of the channel. Alternatively, you can ensure that the user ID in the *UserIdentifier* field of an incoming message is defined in the new security domain.

If an incoming message contains a digital certificate for the user of the application that sent the message, a message exit can validate the certificate and map the Distinguished Name in the certificate to a user ID that is valid on the receiving system. It can then set the *UserIdentifier* field in the message descriptor to this user ID.

If it is necessary for a message exit to change the value of the *UserIdentifier* field in an incoming message, it might be appropriate for the message exit to authenticate the sender of the message at the same time. For more details, see "Identity mapping in message exits" on page 693.

## Implementing access control in the API exit and API-crossing exit

An API or API-crossing exit can provide access controls to supplement those provided by WebSphere MQ. In particular, the exit can provide access control at the message level. The exit can ensure that an application puts on a queue, or gets from a queue, only those messages that satisfy certain criteria.

Consider the following examples:

- A message contains information about an order. When an application attempts to put a message on a queue, an API or API-crossing exit can check that the total value of the order is less than some prescribed limit.
- Messages arrive on a destination queue from remote queue managers. When an application attempts to get a message from the queue, an API or API-crossing exit can check that the sender of the message is authorized to send a message to the queue.

---

## Confidentiality of messages

To maintain confidentiality, encrypt your messages. There are various methods of encrypting messages in WebSphere MQ depending on your needs.

Your choice of CipherSpec determines what level of confidentiality you have.

If you need application-level, end-to-end data protection for your point to point messaging infrastructure, you can use WebSphere MQ Advanced Message Security to encrypt the messages, or write your own API exit or API-crossing exit.

If you need to encrypt messages only while they are being transported through a channel, because you have adequate security on your queue managers, you can use SSL or TLS, or you can write your own security exit, message exit, or send and receive exit programs.

For more information about WebSphere MQ Advanced Message Security, see “WebSphere MQ Advanced Message Security” on page 439. The use of SSL and TLS with WebSphere MQ is described at “WebSphere MQ support for SSL and TLS” on page 386. The use of exit programs in message encryption is described at “Implementing confidentiality in user exit programs” on page 780.

## Connecting two queue managers using SSL or TLS

Secure communications that use the SSL or TLS cryptographic security protocols involve setting up the communication channels and managing the digital certificates that you will use for authentication.

To set up your SSL or TLS installation you must define your channels to use SSL or TLS. You must also obtain and manage your digital certificates. On a test system, you can use self-signed certificates or certificates issued by a local certificate authority (CA). On a production system, do not use self-signed certificates. For more information, see “Protecting channels with SSL” on page 449.

For full information about creating and managing certificates, see “Working with SSL or TLS on IBM i” on page 619, “Working with SSL or TLS on UNIX, Linux and Windows systems” on page 631, or “Working with SSL or TLS on z/OS” on page 674.

This collection of topics introduces the tasks involved in setting up SSL communications, and provides step-by-step guidance on completing those tasks.

You might also want to test SSL or TLS client authentication, which are an optional part of the protocols. During the SSL or TLS handshake, the SSL or TLS client always obtains and validates a digital certificate from the server. With the WebSphere MQ implementation, the SSL or TLS server always requests a certificate from the client.

### Notes:

1. In this context, an SSL client refers to the connection initiating the handshake.
2. See the  Glossary (*WebSphere MQ V7.1 Product Overview Guide*) for further details.

On IBM i, UNIX, Linux and Windows systems, the SSL or TLS client sends a certificate only if it has one labeled in the correct WebSphere MQ format, which is `ibmwebsphermq` followed by the name of your queue manager changed to lowercase. For example, for QM1, `ibmwebsphermqqm1`.

On z/OS, the SSL or TLS client sends a certificate only if it has one of the following certificates:

- For a shared channel only, a certificate with a label in the format `ibmWebSphereMQ` followed by the name of your queue-sharing group, for example `ibmWebSphereMQQSG1`
- A certificate with a label in the format `ibmWebSphereMQ` followed by the name of your queue manager, for example `ibmWebSphereMQQM1`

- A default certificate (which might be the `ibmWebSphereMQ` certificate).

If the channel is shared, the channel first tries to find a certificate for the queue-sharing group. If it does not find a certificate for a queue-sharing group, it tries to find a certificate for the queue manager.

WebSphere MQ uses the `ibmwebspheremq` prefix on a label to avoid confusion with certificates for other products. Ensure that you specify the entire certificate label in lowercase.

The SSL or TLS server always validates the client certificate if one is sent. If the client does not send a certificate, authentication fails only if the end of the channel that is acting as the SSL or TLS server is defined with either the `SSLCAUTH` parameter set to `REQUIRED` or an `SSLPEER` parameter value set. For more information about connecting a queue manager anonymously, that is, when the SSL or TLS client does not send a certificate, see “Connecting two queue managers using one-way authentication” on page 766.

## Using self-signed certificates for mutual authentication of two queue managers

Follow these sample instructions to implement mutual authentication between two queue managers, using self-signed SSL or TLS certificates.

### About this task

Scenario:

- You have two queue managers, QM1 and QM2, which need to communicate securely. You require mutual authentication to be carried out between QM1 and QM2.
- You have decided to test your secure communication using self-signed certificates.

The resulting configuration looks like this:

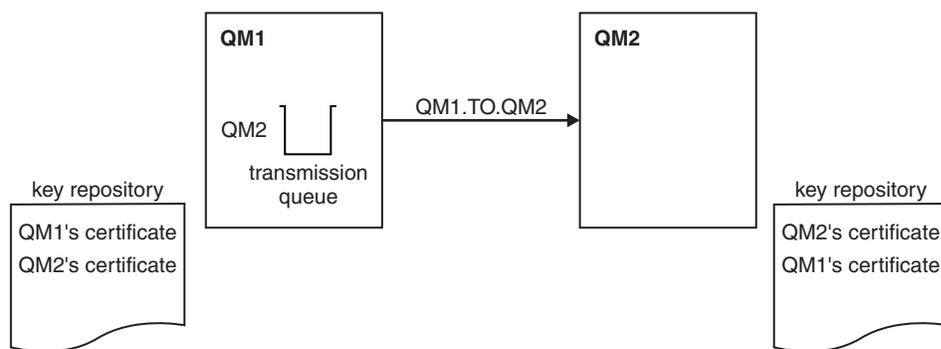


Figure 89. Configuration resulting from this task

In Figure 89, the key repository for QM1 contains the certificate for QM1 and the public certificate from QM2. The key repository for QM2 contains the certificate for QM2 and the public certificate from QM1.

### Procedure

1. Prepare the key repository on each queue manager, according to operating system:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
2. Create a self-signed certificate for each queue manager:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.

3. Extract a copy of each certificate:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
4. Transfer the public part of the QM1 certificate to the QM2 system and vice versa, using a utility such as FTP, as described in “Exchanging self-signed certificates” on page 683.
5. Add the partner certificate to the key repository for each queue manager:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
6. On QM1, define a sender channel and associated transmission queue, by issuing commands like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QM1.MACH.COM)
XMITQ(QM2)
SSLCIPH(RC4_MD5_US) DESCR('Sender channel using SSL from QM1 to QM2')

DEFINE QLOCAL(QM2) USAGE(XMITQ)
```

This example uses CipherSpec RC4\_MD5. The CipherSpecs at each end of the channel must be the same.

7. On QM2, define a receiver channel, by issuing a command like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC4_MD5_US)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from QM1 to QM2')
```

The channel must have the same name as the sender channel you defined in step 6, and use the same CipherSpec.

8. Start the channel, as described in “Starting the sender channel” on page 684.

## Results

Key repositories and channels are created as illustrated in Figure 89 on page 762

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following examples.

From queue manager QM1, enter the following command:

```
DISPLAY CHS(QM1.TO.QM2) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(QM1.TO.QM2) CHLTYPE(SDR)
CONNAME(9.20.25.40) CURRENT
RQMNAME(QM2)
SSLCERTI("CN=QM2,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5E:02,CN=QM2,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(MQGET)
XMITQ(QM2)
```

From queue manager QM2, enter the following command:

```
DISPLAY CHS(QM1.TO.QM2) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```

DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(QM2.TO.QM1) CHLTYPE(RCVR)
CONNAME(9.20.35.92) CURRENT
RQMNAME(QM1)
SSLCERTI("CN=QM1,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QM1,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
XMITQ()

```

In each case, the value of SSLPEER must match that of the DN in the partner certificate that was created in Step 2. The issuer's name matches the peer name because the certificate is self-signed.

SSLPEER is optional. If it is specified, its value must be set so that the DN in the partner certificate (created in step 2) is allowed. For more information about the use of SSLPEER, see  WebSphere MQ rules for SSLPEER values (*WebSphere MQ V7.1 Reference*).

### Using CA-signed certificates for mutual authentication of two queue managers

Follow these sample instructions to implement mutual authentication between two queue managers, using CA-signed SSL or TLS certificates.

#### About this task

Scenario:

- You have two queue managers called QMA and QMB, which need to communicate securely. You require mutual authentication to be carried out between QMA and QMB.
- In the future you are planning to use this network in a production environment, and therefore you have decided to use CA-signed certificates from the beginning.

The resulting configuration looks like this:

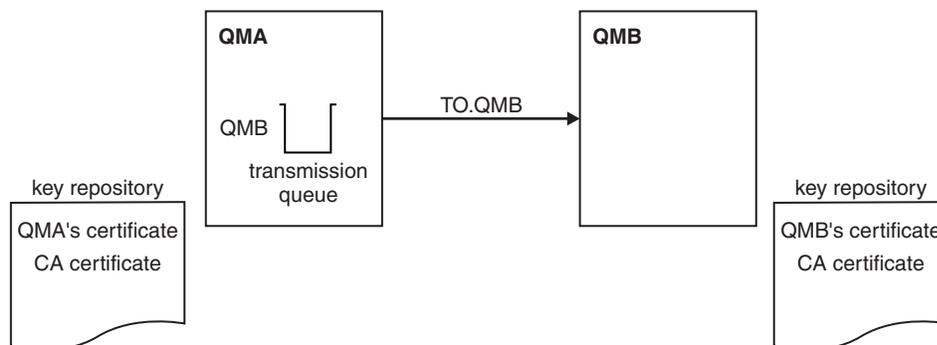


Figure 90. Configuration resulting from this task

In Figure 90, the key repository for QMA contains QMA's certificate and the CA certificate. The key repository for QMB contains QMB's certificate and the CA certificate. In this example both QMA's certificate and QMB's certificate were issued by the same CA. If QMA's certificate and QMB's certificate were issued by different CAs then the key repositories for QMA and QMB must contain both CA certificates.

## Procedure

1. Prepare the key repository on each queue manager, according to operating system:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
2. Request a CA-signed certificate for each queue manager. You might use different CAs for the two queue managers.
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
3. Add the Certificate Authority certificate to the key repository for each queue manager: If the Queue managers are using different Certificate Authorities then the CA certificate for each Certificate Authority must be added to both key repositories.
  - Do not perform this step on IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
4. Add the CA-signed certificate to the key repository for each queue manager:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
5. On QMA, define a sender channel and associated transmission queue by issuing commands like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QMB.MACH.COM)
XMITQ(QMB)
SSLCIPH(RC2_MD5_EXPORT) DESC('Sender channel using SSL from QMA to QMB')

DEFINE QLOCAL(QMB) USAGE(XMITQ)
```

This example uses CipherSpec RC4\_MD5. The CipherSpecs at each end of the channel must be the same.

6. On QMB, define a receiver channel by issuing a command like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC2_MD5_EXPORT)
SSLCAUTH(REQUIRED) DESC('Receiver channel using SSL to QMB')
```

The channel must have the same name as the sender channel you defined in step 6, and use the same CipherSpec.

7. Start the channel:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.

## Results

Key repositories and channels are created as illustrated in Figure 90 on page 764.

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is like that shown in the following examples.

From queue manager QMA, enter the following command:

```
DISPLAY CHS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(TO.QMB) CHLTYPE(SDR)
CONNNAME(9.20.25.40) CURRENT
RQMNAME(QMB)
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMB,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(MQGET)
XMITQ(QMB)
```

From the queue manager QMB, enter the following command:

```
DISPLAY CHS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(TO.QMB) CHLTYPE(RCVR)
CONNNAME(9.20.35.92) CURRENT
RQMNAME(QMA)
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMA,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
XMITQ()
```

In each case, the value of SSLPEER must match that of the Distinguished Name (DN) in the partner certificate that was created in Step 2. The issuer name matches the subject DN of the CA certificate that signed the personal certificate added in Step 4.

## Connecting two queue managers using one-way authentication

Follow these sample instructions to modify a system with mutual authentication to allow a queue manager to connect using one-way authentication to another; that is, when the SSL or TLS client does not send a certificate.

### About this task

Scenario:

- Your two queue managers (QM1 and QM2) have been set up as in “Using CA-signed certificates for mutual authentication of two queue managers” on page 764.
- You want to change QM1 so that it connects using one-way authentication to QM2.

The resulting configuration looks like this:

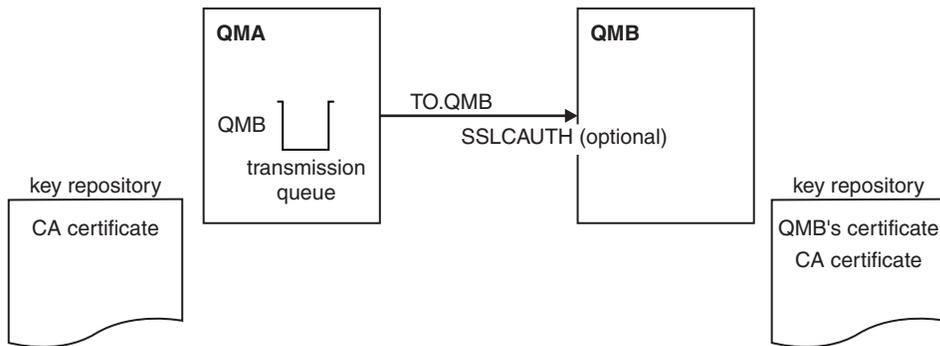


Figure 91. Queue managers allowing one-way authentication

## Procedure

1. Remove QM1's personal certificate from its key repository, according to operating system:
  - On IBM i systems. The certificate is labeled as follows:
    - `ibmwebspheremq` followed by the name of your queue manager folded to lowercase. For example, for QM1, `ibmwebspheremqqm1`.
  - On UNIX, Linux, and Windows systems. The certificate is labeled as follows:
    - `ibmwebspheremq` followed by the name of your queue manager folded to lowercase. For example, for QM1, `ibmwebspheremqqm1`.
  - On z/OS systems. Perform this step twice, to remove the personal certificate for QMA and the default certificate. The certificates are labeled as follows:
    - `ibmWebSphereMQ` followed by the name of your queue manager or queue-sharing group. For example, for QM1, `ibmWebSphereMQM1`.
    - The default certificate (which might be the `ibmWebSphereMQ` certificate).
2. Optional: On QM1, if any SSL or TLS channels have run previously, refresh the SSL or TLS environment, as described in “Refreshing the SSL or TLS environment” on page 684.
3. Allow anonymous connections on the receiver, as described in “Allowing anonymous connections on a receiver channel” on page 684.

## Results

Key repositories and channels are changed as illustrated in Figure 91

## What to do next

If the sender channel was running and you issued the `REFRESH SECURITY TYPE(SSL)` command (in step 2), the channel restarts automatically. If the sender channel was not running, start it.

At the server end of the channel, the presence of the peer name parameter value on the channel status display indicates that a client certificate has flowed.

Verify that the task has been completed successfully by issuing some `DISPLAY` commands. If the task was successful, the resulting output is similar to that shown in the following examples:

From the QM1 queue manager, enter the following command:

```
DISPLAY CHS(TO.QM2) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(TO.QMB) SSLPEER
AMQ8417: Display Channel Status details.
CHANNEL(TO.QM2) CHLTYPE(SDR)
CONNAME(9.20.25.40) CURRENT
RQMNAME(QM2)
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMB,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(MQGET)
XMITQ(QM2)
```

From the QM2 queue manager, enter the following command:

```
DISPLAY CHS(TO.QM2) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```
DISPLAY CHSTATUS(TO.QM2) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(TO.QM2) CHLTYPE(RCVR)
CONNAME(9.20.35.92) CURRENT
RQMNAME(QMA) SSLCERTI()
SSLPEER() STATUS(RUNNING)
SUBSTATE(RECEIVE) XMITQ()
```

On QM2, the SSLPEER field is empty, showing that QM1 did not send a certificate. On QM1, the value of SSLPEER matches that of the DN in QM2's personal certificate.

## Connecting a client to a queue manager securely

Secure communications that use the SSL or TLS cryptographic security protocols involve setting up the communication channels and managing the digital certificates that you will use for authentication.

To set up your SSL or TLS installation you must define your channels to use SSL or TLS. You must also obtain and manage your digital certificates. On a test system, you can use self-signed certificates or certificates issued by a local certificate authority (CA). On a production system, do not use self-signed certificates. For more information, see "Protecting channels with SSL" on page 449.

For full information about creating and managing certificates, see "Working with SSL or TLS on IBM i" on page 619, "Working with SSL or TLS on UNIX, Linux and Windows systems" on page 631, or "Working with SSL or TLS on z/OS" on page 674.

This collection of topics introduces the tasks involved in setting up SSL communications, and provides step-by-step guidance on completing those tasks.

You might also want to test SSL or TLS client authentication, which are an optional part of the protocols. During the SSL or TLS handshake, the SSL or TLS client always obtains and validates a digital certificate from the server. With the WebSphere MQ implementation, the SSL or TLS server always requests a certificate from the client.

On IBM i, UNIX, Linux, and Windows systems, the SSL or TLS client sends a certificate only if it has one labeled in the correct WebSphere MQ format, which is `ibmwebsphermq` followed by your logon user ID changed to lowercase, for example `ibmwebsphermqmyuserid`.

WebSphere MQ uses the `ibmwebsphermq` prefix on a label to avoid confusion with certificates for other products. Ensure that you specify the entire certificate label in lowercase.

The SSL or TLS server always validates the client certificate if one is sent. If the client does not send a certificate, authentication fails only if the end of the channel that is acting as the SSL or TLS server is defined with either the SSLCAUTH parameter set to REQUIRED or an SSLPEER parameter value set. For more information about connecting a queue manager anonymously, see “Connecting a client to a queue manager anonymously” on page 773.

## Using self-signed certificates for mutual authentication of a client and queue manager

Follow these sample instructions to implement mutual authentication between a client and a queue manager, by using self-signed SSL or TLS certificates.

### About this task

Scenario:

- You have a client, C1, and a queue manager, QM1, which need to communicate securely. You require mutual authentication to be carried out between C1 and QM1.
- You have decided to test your secure communication by using self-signed certificates.

DCM on IBM i does not support self-signed certificates, so this task is not applicable on IBM i systems.

The resulting configuration looks like this:

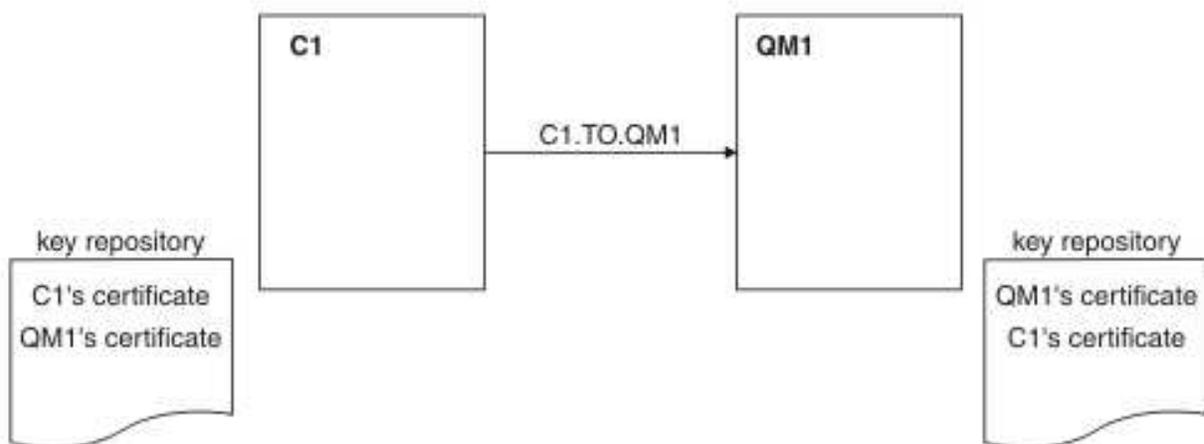


Figure 92. Configuration resulting from this task

In Figure 92, the key repository for QM1 contains the certificate for QM1 and the public certificate from C1. The key repository for C1 contains the certificate for C1 and the public certificate from QM1.

### Procedure

1. Prepare the key repository on the client and queue manager, according to operating system:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
2. Create self-signed certificates for the client and queue manager:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
3. Extract a copy of each certificate:

- On UNIX, Linux, and Windows systems.
  - On z/OS systems.
4. Transfer the public part of the C1 certificate to the QM1 system and vice versa, using a utility such as FTP, as described in “Exchanging self-signed certificates” on page 683.
  5. Add the partner certificate to the key repository for the client and queue manager:
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems.
  6. Define a client-connection channel in either of the following ways:
    - Using the MQCONNX call with the MQSCO structure on C1, as described in  [Creating a client-connection channel on the WebSphere MQ MQI client \(WebSphere MQ V7.1 Installing Guide\)](#).
    - Using a client channel definition table, as described in  [Creating server-connection and client-connection definitions on the server \(WebSphere MQ V7.1 Installing Guide\)](#).
  7. On QM1, define a server-connection channel, by issuing a command like the following example:
 

```
DEFINE CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(RC4_MD5_US)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from C1 to QM1')
```

The channel must have the same name as the client-connection channel you defined in step 6, and use the same CipherSpec.

## Results

Key repositories and channels are created as illustrated in Figure 92 on page 769

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following example.

From queue manager QM1, enter the following command:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNAME(9.20.35.92) CURRENT
SSLCERTI("CN=QM1,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5E:02,CN=QM2,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
```

It is optional to set the SSLPEER filter attribute of the channel definitions. If the channel definition SSLPEER is set, its value must match the subject DN in the partner certificate that was created in Step 2. After a successful connection, the SSLPEER field in the DISPLAY CHSTATUS output shows the subject DN of the remote client certificate.

## Using CA-signed certificates for mutual authentication of a client and queue manager

Follow these sample instructions to implement mutual authentication between a client and a queue manager, by using CA-signed SSL or TLS certificates.

### About this task

Scenario:

- You have a client, C1, and a queue manager, QM1, which need to communicate securely. You require mutual authentication to be carried out between C1 and QM1.
- In the future you are planning to use this network in a production environment, and therefore you have decided to use CA-signed certificates from the beginning.

The resulting configuration looks like this:

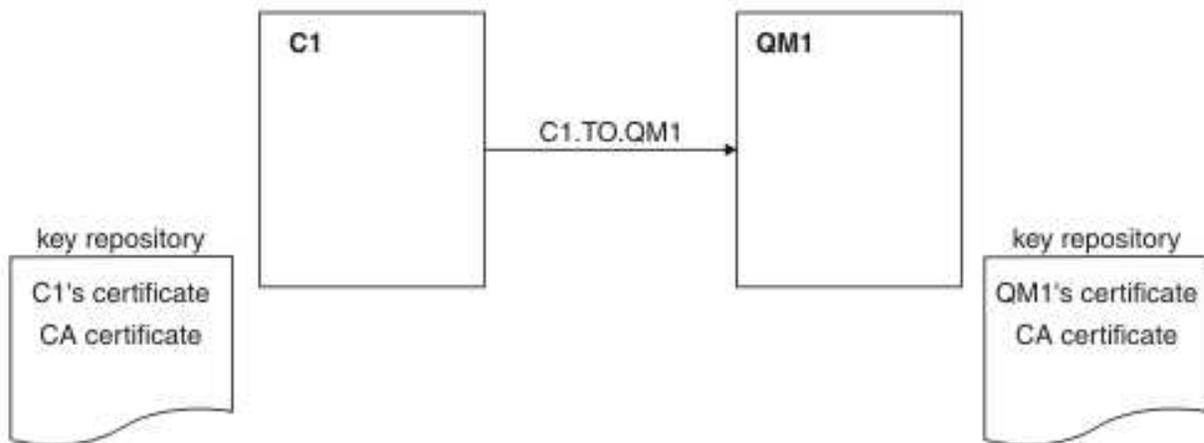


Figure 93. Configuration resulting from this task

In Figure 93, the key repository for C1 contains certificate for C1 and the CA certificate. The key repository for QM1 contains the certificate for QM1 and the CA certificate. In this example both C1's certificate and QM1's certificate were issued by the same CA. If C1's certificate and QM1's certificate were issued by different CAs then the key repositories for C1 and QM1 must contain both CA certificates.

### Procedure

1. Prepare the key repository on the client and queue manager, according to operating system:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
2. Request a CA-signed certificate for the client and queue manager. You might use different CAs for the client and queue manager.
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
3. Add the certificate authority certificate to the key repository for the client and queue manager. If the client and queue manager are using different Certificate Authorities then the CA certificate for each Certificate Authority must be added to both key repositories.

- Do not perform this step on IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
4. Add the CA-signed certificate to the key repository for the client and queue manager:
    - On IBM i systems.
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems (queue manager only).
  5. Define a client-connection channel in either of the following ways:
    - Using the MQCONN call with the MQSCO structure on C1, as described in  [Creating a client-connection channel on the WebSphere MQ MQI client \(WebSphere MQ V7.1 Installing Guide\)](#).
    - Using a client channel definition table, as described in  [Creating server-connection and client-connection definitions on the server \(WebSphere MQ V7.1 Installing Guide\)](#).
  6. On QM1, define a server-connection channel by issuing a command like the following example:
 

```
DEFINE CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(RC2_MD5_EXPORT)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from C1 to QM1')
```

The channel must have the same name as the client-connection channel you defined in step 6, and use the same CipherSpec.

## Results

Key repositories and channels are created as illustrated in Figure 93 on page 771.

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is like that shown in the following example.

From the queue manager QM1, enter the following command:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNNAME(9.20.35.92) CURRENT
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMA,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
```

The SSLPEER field in the DISPLAY CHSTATUS output shows the subject DN of the remote client certificate that was created in Step 2. The issuer name matches the subject DN of the CA certificate that signed the personal certificate added in Step 4.

## Connecting a client to a queue manager anonymously

Follow these sample instructions to modify a system with mutual authentication to allow a queue manager to connect anonymously to another.

### About this task

Scenario:

- Your queue manager and client (QM1 and C1) have been set up as in “Using CA-signed certificates for mutual authentication of a client and queue manager” on page 771.
- You want to change C1 so that it connects anonymously to QM1.

The resulting configuration looks like this:

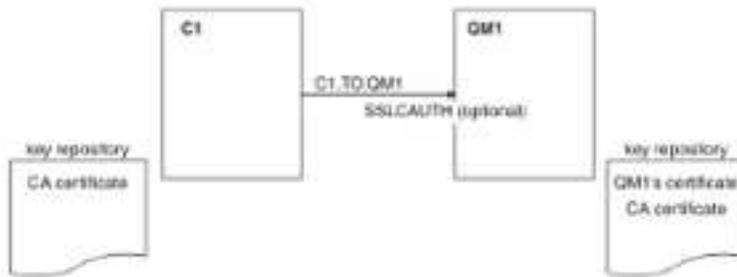


Figure 94. Client and queue manager allowing anonymous connection

### Procedure

1. Remove the personal certificate from key repository for C1, according to operating system:
  - On IBM i systems. The certificate is labeled as follows:
    - `ibmwebsphermq` followed by your logon user ID folded to lowercase, for example `ibmwebsphermqmyuserid`.
  - On UNIX, Linux, and Windows systems. The certificate is labeled as follows:
    - `ibmwebsphermq` followed by your logon user ID folded to lowercase, for example `ibmwebsphermqmyuserid`.
2. Restart the client application, or cause the client application to close and reopen all SSL or TLS connections.
3. Allow anonymous connections on the queue manager, by issuing the following command:  

```
ALTER CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) SSLCAUTH(OPTIONAL)
```

### Results

Key repositories and channels are changed as illustrated in Figure 94

### What to do next

At the server end of the channel, the presence of the peer name parameter value on the channel status display indicates that a client certificate has flowed.

Verify that the task has been completed successfully by issuing some DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following example:

From queue manager QM1, enter the following command:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```

DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNAME(9.20.35.92) CURRENT
SSLCERTI() SSLPEER()
STATUS(RUNNING) SUBSTATE(RECEIVE)

```

The SSLCERTI and SSLPEER fields are empty, showing that C1 did not send a certificate.

## Specifying CipherSpecs

Specify a CipherSpec by using the **SSLCIPH** parameter in either the **DEFINE CHANNEL** MQSC command or the **ALTER CHANNEL** MQSC command.

Some of the CipherSpecs that you can use with IBM WebSphere MQ are FIPS compliant. Others, such as NULL\_MD5, are not. Similarly, some of the FIPS compliant CipherSpecs are also Suite B compliant although others, are not. All Suite B compliant CipherSpecs are also FIPS compliant. All Suite B compliant CipherSpecs fall into two groups: 128 bit (for example, ECDHE\_ECDSA\_AES\_128\_GCM\_SHA256) and 192 bit (for example, ECDHE\_ECDSA\_AES\_256\_GCM\_SHA384),

The following diagram illustrates the relationship between these subsets:



Cipher specifications that you can use with IBM WebSphere MQ SSL and TLS support are listed in the following table. When you request a personal certificate, you specify a key size for the public and private key pair. The key size that is used during the SSL handshake is the size stored in the certificate unless it is determined by the CipherSpec, as noted in the table.

| CipherSpec name                 | Protocol used | Data integrity | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|---------------------------------|---------------|----------------|----------------------|-----------------|-------------------|---------|
| AES_SHA_US                      | SSL 3.0       | SHA-1          | AES                  | 128             | No                | No      |
| DES_SHA_EXPORT <sup>2</sup>     | SSL 3.0       | SHA-1          | DES                  | 56              | No                | No      |
| DES_SHA_EXPORT1024 <sup>3</sup> | SSL 3.0       | SHA-1          | DES                  | 56              | No                | No      |
| FIPS_WITH_DES_CBC_SHA           | SSL 3.0       | SHA-1          | DES                  | 56              | No <sup>6</sup>   | No      |
| NULL_MD5 <sup>a</sup>           | SSL 3.0       | MD5            | None                 | 0               | No                | No      |
| NULL_SHA <sup>a</sup>           | SSL 3.0       | SHA-1          | None                 | 0               | No                | No      |
| RC2_MD5_EXPORT <sup>2 a</sup>   | SSL 3.0       | MD5            | RC2                  | 40              | No                | No      |
| RC4_MD5_EXPORT <sup>2 a</sup>   | SSL 3.0       | MD5            | RC4                  | 40              | No                | No      |
| RC4_MD5_US <sup>a</sup>         | SSL 3.0       | MD5            | RC4                  | 128             | No                | No      |
| RC4_SHA_US <sup>a</sup>         | SSL 3.0       | SHA-1          | RC4                  | 128             | No                | No      |

| Cipherspec name                               | Protocol used | Data integrity         | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|-----------------------------------------------|---------------|------------------------|----------------------|-----------------|-------------------|---------|
| RC4_56_SHA_EXPORT1024 <sup>3 b</sup>          | SSL 3.0       | SHA-1                  | RC4                  | 56              | No                | No      |
| TLS_RSA_EXPORT_WITH_RC2_40_MD5 <sup>c</sup>   | TLS 1.0       | MD5                    | RC2                  | 40              | No                | No      |
| TLS_RSA_EXPORT_WITH_RC4_40_MD5 <sup>2 c</sup> | TLS 1.0       | MD5                    | RC4                  | 40              | No                | No      |
| TLS_RSA_WITH_AES_128_CBC_SHA <sup>a</sup>     | TLS 1.0       | SHA-1                  | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_CBC_SHA <sup>4 a</sup>   | TLS 1.0       | SHA-1                  | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_DES_CBC_SHA <sup>a</sup>         | TLS 1.0       | SHA-1                  | DES                  | 56              | No <sup>5</sup>   | No      |
| TLS_RSA_WITH_NULL_MD5 <sup>c</sup>            | TLS 1.0       | MD5                    | None                 | 0               | No                | No      |
| TLS_RSA_WITH_NULL_SHA <sup>c</sup>            | TLS 1.0       | SHA-1                  | None                 | 0               | No                | No      |
| TLS_RSA_WITH_RC4_128_MD5 <sup>c</sup>         | TLS 1.0       | MD5                    | RC4                  | 128             | No                | No      |
| ECDHE_ECDSA_AES_128_CBC_SHA256 <sup>d</sup>   | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| ECDHE_ECDSA_AES_256_CBC_SHA384 <sup>d</sup>   | TLS 1.2       | SHA-384                | AES                  | 256             | Yes               | No      |
| ECDHE_ECDSA_AES_128_GCM_SHA256 <sup>b</sup>   | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | 128 bit |
| ECDHE_ECDSA_AES_256_GCM_SHA384 <sup>b</sup>   | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | 192 bit |
| ECDHE_ECDSA_NULL_SHA256 <sup>b</sup>          | TLS 1.2       | SHA-1                  | None                 | 0               | No                | No      |
| ECDHE_ECDSA_RC4_128_SHA256 <sup>b</sup>       | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |
| ECDHE_RSA_AES_128_CBC_SHA256 <sup>d</sup>     | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| ECDHE_RSA_AES_256_CBC_SHA384 <sup>d</sup>     | TLS 1.2       | SHA-384                | AES                  | 256             | Yes               | No      |
| ECDHE_RSA_AES_128_GCM_SHA256 <sup>b</sup>     | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | No      |
| ECDHE_RSA_AES_256_GCM_SHA384 <sup>b</sup>     | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | No      |
| ECDHE_RSA_NULL_SHA256 <sup>b</sup>            | TLS 1.2       | SHA-1                  | None                 | 0               | No                | No      |
| ECDHE_RSA_RC4_128_SHA256 <sup>b</sup>         | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |
| TLS_RSA_WITH_AES_128_CBC_SHA256 <sup>d</sup>  | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_CBC_SHA256 <sup>d</sup>  | TLS 1.2       | SHA-256                | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_AES_128_GCM_SHA256 <sup>b</sup>  | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_GCM_SHA384 <sup>b</sup>  | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_NULL_NULL <sup>b</sup>           | TLS 1.2       | None                   | None                 | 0               | No                | No      |
| TLS_RSA_WITH_NULL_SHA256 <sup>d</sup>         | TLS 1.2       | SHA-256                | None                 | 0               | No                | No      |
| TLS_RSA_WITH_RC4_128_SHA256 <sup>b</sup>      | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |

| CipherSpec name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Protocol used | Data integrity | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------------|----------------------|-----------------|-------------------|---------|
| <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. Specifies whether the CipherSpec is FIPS-certified on a FIPS-certified platform. See “Federal Information Processing Standards (FIPS)” on page 393 for an explanation of FIPS.</li> <li>2. The maximum handshake key size is 512 bits. If either of the certificates exchanged during the SSL handshake has a key size greater than 512 bits, a temporary 512-bit key is generated for use during the handshake.</li> <li>3. The handshake key size is 1024 bits.</li> <li>4. This CipherSpec cannot be used to secure a connection from the WebSphere MQ Explorer to a queue manager unless the appropriate unrestricted policy files are applied to the JRE used by the Explorer.</li> <li>5. This CipherSpec was FIPS 140-2 certified before 19 May 2007.</li> <li>6. This CipherSpec was FIPS 140-2 certified before 19 May 2007. The name FIPS_WITH_DES_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.</li> <li>7. This CipherSpec can be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, either avoid using triple DES, or enable secret key reset when using this CipherSpec.</li> </ol> <p><b>Platform support:</b></p> <ul style="list-style-type: none"> <li>• a Available on all supported platforms.</li> <li>• b Available only on UNIX, Linux, and Windows platforms.</li> <li>• c Available only on IBM i platforms.</li> <li>• d Available on UNIX, Linux, and Windows platforms, z/OS, and IBM i V7R1M0 TR6 or later.</li> </ul> <p>For further information, on using these CipherSpecs on the IBM i platform, see  Using TLS Version 1.2</p> <p>To use these CipherSpecs on z/OS, you must be using z/OS V1R13, and install the IBM WebSphere MQ APAR  PM77341 in conjunction with the System SSL APAR  OA39422.</p> |               |                |                      |                 |                   |         |

**Related concepts:**

“Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403

**Related reference:**

 DEFINE CHANNEL (*WebSphere MQ V7.1 Reference*)

 ALTER CHANNEL (*WebSphere MQ V7.1 Reference*)

**Obtaining information about CipherSpecs using WebSphere MQ Explorer**

You can use WebSphere MQ Explorer to display descriptions of CipherSpecs.

Use the following procedure to obtain information about the CipherSpecs in “Specifying CipherSpecs” on page 774:

1. Open **WebSphere MQ Explorer** and expand the **Queue Managers** folder.
2. Ensure that you have started your queue manager.
3. Select the queue manager you want to work with and click **Channels**.
4. Right-click the channel you want to work with and select **Properties**.
5. Select the **SSL** property page.
6. Select from the list the CipherSpec you want to work with. A description is displayed in the window below the list.

## Alternatives for specifying CipherSpecs

For those platforms where the operating system provides the SSL support, your system might support new CipherSpecs. You can specify a new CipherSpec with the SSLCIPH parameter, but the value you supply depends on your platform.

**Note:** This section does not apply to UNIX, Linux or Windows systems, because the CipherSpecs are provided with the WebSphere MQ product, so new CipherSpecs do not become available after shipment.

For those platforms where the operating system provides the SSL support, your system might support new CipherSpecs that are not included in “Specifying CipherSpecs” on page 774. You can specify a new CipherSpec with the SSLCIPH parameter, but the value you supply depends on your platform. In all cases the specification *must* correspond to an SSL CipherSpec that is both valid and supported by the version of SSL your system is running.

**IBM i** A two-character string representing a hexadecimal value.

For more information about the permitted values, refer to the appropriate product documentation (search on *cipher\_spec*). For all IBM i versions, start here: [IBM i](#)

You can use either the CHGMQMCHL or the CRTMQMCHL command to specify the value, for example:

```
CRTMQMCHL CHLNAME('channel name') SSLCIPH('hexadecimal value')
```

You can also use the ALTER QMGR MQSC command to set the SSLCIPH parameter.

**z/OS** A two-character string representing a hexadecimal value. The hexadecimal codes correspond to the values defined in the SSL protocol.

For more information, refer to the description of `gsk_environment_open()` in the API reference chapter of *z/OS Cryptographic Services System SSL Programming*, SC24-5901, where there is a list of all the supported SSL V3.0 and TLS V1.0 cipher specifications in the form of 2-digit hexadecimal codes.

## Considerations for WebSphere MQ clusters

With WebSphere MQ clusters it is safest to use the CipherSpec names in “Specifying CipherSpecs” on page 774. If you use an alternative specification, be aware that the specification might not be valid on other platforms. For more information, refer to “SSL and clusters” on page 806.

## Specifying a CipherSpec for a WebSphere MQ MQI client

You have three options for specifying a CipherSpec for a WebSphere MQ MQI client.

These options are as follows:

- Using a channel definition table
- Using the SSL configuration options structure, MQSCO, on an MQCONN call
- Using the Active Directory (on Windows systems with Active Directory support)

#### Related reference:

“Specifying that an MQI channel uses SSL” on page 398

## Specifying a CipherSuite with WebSphere MQ classes for Java and WebSphere MQ classes for JMS

WebSphere MQ classes for Java and WebSphere MQ classes for JMS specify CipherSuites differently from other platforms.

For information about specifying a CipherSuite with WebSphere MQ classes for Java, see  [Secure Sockets Layer \(SSL\) support](#) (*WebSphere MQ V7.1 Programming Guide*)

For information about specifying a CipherSuite with WebSphere MQ classes for JMS, see  [Using Secure Sockets Layer \(SSL\) with WebSphere MQ classes for JMS](#) (*WebSphere MQ V7.1 Programming Guide*)

#### Related concepts:

“Cryptographic security protocols: SSL and TLS” on page 376

## Resetting SSL and TLS secret keys

IBM WebSphere MQ supports the resetting of secret keys on queue managers and clients.

Secret keys are reset when a specified number of encrypted bytes of data have flowed across the channel, or after the channel has been idle for a period.

### Queue manager

For a queue manager, use the command **ALTER QMGR** with the parameter **SSLRKEYC** to set the values used during key renegotiation. On IBM i, use **CHGMQM** with the **SSLRSTCNT** parameter.

### MQI client

By default, MQI clients do not renegotiate the secret key. You can make an MQI client renegotiate the key in any of three ways. In the following list, the methods are shown in order of priority. If you specify multiple values, the highest priority value is used.

1. By using the KeyResetCount field in the MQSCO structure on an MQCONN call
2. By using the environment variable MQSSLRESET
3. By setting the SSLKeyResetCount attribute in the MQI client configuration file

These variables can be set to an integer in the range 0 through 999 999 999, representing the number of unencrypted bytes sent and received within an SSL or TLS conversation before the SSL or TLS secret key is renegotiated. Specifying a value of 0 indicates that SSL or TLS secret keys are never renegotiated. If you specify an SSL or TLS secret key reset count in the range 1 byte through 32 KB, SSL or TLS channels will use a secret key reset count of 32 KB. This is to avoid excessive key resets which would occur for small SSL or TLS secret key reset values.

If a value greater than zero is specified and channel heartbeats are enabled for the channel, the secret key is also renegotiated before message data is sent or received following a channel heartbeat.

The count of bytes until the next secret key renegotiation is reset after each successful renegotiation.

For full details of the MQSCO structure, see  [KeyResetCount \(MQLONG\)](#) (*WebSphere MQ V7.1 Reference*). For full details of MQSSLRESET, see  [MQSSLRESET](#) (*WebSphere MQ V7.1 Installing Guide*).

For more information about the use of SSL or TLS in the client configuration file, see  [SSL stanza of the client configuration file](#) (*WebSphere MQ V7.1 Installing Guide*).

## Java

For IBM WebSphere MQ classes for Java, an application can reset the secret key in either of the following ways:

- By setting the `sslResetCount` field in the `MQEnvironment` class.
- By setting the environment property `MQC.SSL_RESET_COUNT_PROPERTY` in a `Hashtable` object. The application then assigns the hashtable to the `properties` field in the `MQEnvironment` class, or passes the hashtable to an `MQQueueManager` object on its constructor.

If the application uses more than one of these ways, the usual precedence rules apply. See [Class com.ibm.mq.MQEnvironment](#) for the precedence rules.

The value of the `sslResetCount` field or environment property `MQC.SSL_RESET_COUNT_PROPERTY` represents the total number of bytes sent and received by the WebSphere MQ classes for Java client code before the secret key is renegotiated. The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by the WebSphere MQ classes for Java client.

If the reset count is zero, which is the default value, the secret key is never renegotiated. The reset count is ignored if no `CipherSuite` is specified.

## JMS

For IBM WebSphere MQ classes for JMS, the `SSLRESETCOUNT` property represents the total number of bytes sent and received by a connection before the secret key that is used for encryption is renegotiated. The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by IBM WebSphere MQ classes for JMS. For example, to configure a `ConnectionFactory` object that can be used to create a connection over an SSL or TLS enabled MQI channel with a secret key that is renegotiated after 4 MB of data have flowed, issue the following command to `JMSAdmin`:

```
ALTER CF(my.cf) SSLRESETCOUNT(4194304)
```

If the value of `SSLRESETCOUNT` is zero, which is the default value, the secret key is never renegotiated. The `SSLRESETCOUNT` property is ignored if `SSLCIPHERSUITE` is not set.

## .NET

For .NET unmanaged clients, the integer property `SSLKeyResetCount` indicates the number of unencrypted bytes sent and received within an SSL or TLS conversation before the secret key is renegotiated.

For information about the use of object properties in IBM WebSphere MQ classes for .NET, see

 [Getting and setting attribute values \(WebSphere MQ V7.1 Programming Guide\)](#).

## XMS .NET

For XMS .NET unmanaged clients, see the documentation for `Message Service Client for .NET`

### Related reference:

 ALTER QMGR (*WebSphere MQ V7.1 Reference*)

 DISPLAY QMGR (*WebSphere MQ V7.1 Reference*)

### Related information:

 Change Message Queue Manager (CHGMQM) (*WebSphere MQ V7.1 Reference*)

 Display Message Queue Manager (DSPMQM) (*WebSphere MQ V7.1 Reference*)

## Implementing confidentiality in user exit programs

### Implementing confidentiality in security exits

Security exits can play a role in the confidentiality service by generating and distributing the symmetric key for encrypting and decrypting the data that flows on the channel. A common technique for doing this uses PKI technology.

One security exit generates a random data value, encrypts it with the public key of the queue manager or user that the partner security exit is representing, and sends the encrypted data to its partner in a security message. The partner security exit decrypts the random data value with the private key of the queue manager or user it is representing. Each security exit can now use the random data value to derive the symmetric key independently of the other by using an algorithm known to both of them. Alternatively, they can use the random data value as the key.

If the first security exit has not authenticated its partner by this time, the next security message sent by the partner can contain an expected value encrypted with the symmetric key. The first security exit can now authenticate its partner by checking that the partner security exit was able to encrypt the expected value correctly.

The security exits can also use this opportunity to agree the algorithm for encrypting and decrypting the data that flows on the channel, if more than one algorithm is available for use.

### Implementing confidentiality in message exits

A message exit at the sending end of a channel can encrypt the application data in a message and another message exit at the receiving end of the channel can decrypt the data. For performance reasons, a symmetric key algorithm is normally used for this purpose. For more information about how the symmetric key can be generated and distributed, see “Implementing confidentiality in user exit programs.”

Headers in a message, such as the transmission queue header, MQXQH, which includes the embedded message descriptor, must not be encrypted by a message exit. This is because data conversion of the message headers takes place either after a message exit is called at the sending end or before a message exit is called at the receiving end. If the headers are encrypted, data conversion fails and the channel stops.

### Implementing confidentiality in send and receive exits

Send and receive exits can be used to encrypt and decrypt the data that flows on a channel. They are more appropriate than message exits for providing this service for the following reasons:

- On a message channel, message headers can be encrypted as well as the application data in the messages.

- Send and receive exits can be used on MQI channels as well as message channels. Parameters on MQI calls might contain sensitive application data that needs to be protected while it flows on an MQI channel. You can therefore use the same send and receive exits on both kinds of channels.

## Implementing confidentiality in the API exit and API-crossing exit

The application data in a message can be encrypted by an API or API-crossing exit when the message is put by the sending application and decrypted by a second exit when the message is retrieved by the receiving application. For performance reasons, a symmetric key algorithm is typically used for this purpose. However, at the application level, where many users might be sending messages to each other, the problem is how to ensure that only the intended receiver of a message is able to decrypt the message. One solution is to use a different symmetric key for each pair of users that send messages to each other. But this solution might be difficult and time consuming to administer, particularly if the users belong to different organizations. A standard way of solving this problem is known as *digital enveloping* and uses PKI technology.

When an application puts a message on a queue, an API or API-crossing exit generates a random symmetric key and uses the key to encrypt the application data in the message. The exit encrypts the symmetric key with the public key of the intended receiver. It then replaces the application data in the message with the encrypted application data and the encrypted symmetric key. In this way, only the intended receiver can decrypt the symmetric key and therefore the application data. If an encrypted message has more than one possible intended receiver, the exit can encrypt a copy of the symmetric key for each intended receiver.

If different algorithms for encrypting and decrypting the application data are available for use, the exit can include the name of the algorithm it has used.

---

## Data integrity of messages

To maintain data integrity, you can use various types of user exit program to provide message digests or digital signatures for your messages.

### Data integrity

#### Implementing data integrity in messages

When you use SSL or TLS, your choice of CipherSpec determines the level of data integrity in the enterprise. If you use the WebSphere MQ Advanced Message Service (AMS) you can specify the integrity for a unique message.

#### Implementing data integrity in message exits

A message can be digitally signed by a message exit at the sending end of a channel. The digital signature can then be checked by a message exit at the receiving end of a channel to detect whether the message has been deliberately modified.

Some protection can be provided by using a message digest instead of a digital signature. A message digest might be effective against casual or indiscriminate tampering, but it does not prevent the more informed individual from changing or replacing the message, and generating a completely new digest for it. This is particularly true if the algorithm that is used to generate the message digest is a well known one.

#### Implementing data integrity in send and receive exits

On a message channel, message exits are more appropriate for providing this service because a message exit has access to a whole message. On an MQI channel, parameters on MQI calls might contain application data that needs to be protected and only send and receive exits can provide this protection.

#### Implementing data integrity in the API exit or API-crossing exit

A message can be digitally signed by an API or API-crossing exit when the message is put by the sending application. The digital signature can then be checked by a second exit when the message is retrieved by the receiving application to detect whether the message has been deliberately modified.

Some protection can be provided by using a message digest instead of a digital signature. A message digest might be effective against casual or indiscriminate tampering, but it does not prevent the more informed individual from changing or replacing the message, and generating a completely new digest for it. This is particularly true if the algorithm that is used to generate the message digest is a well known one,

## Connecting two queue managers using SSL or TLS

Secure communications that use the SSL or TLS cryptographic security protocols involve setting up the communication channels and managing the digital certificates that you will use for authentication.

To set up your SSL or TLS installation you must define your channels to use SSL or TLS. You must also obtain and manage your digital certificates. On a test system, you can use self-signed certificates or certificates issued by a local certificate authority (CA). On a production system, do not use self-signed certificates. For more information, see “Protecting channels with SSL” on page 449.

For full information about creating and managing certificates, see “Working with SSL or TLS on IBM i” on page 619, “Working with SSL or TLS on UNIX, Linux and Windows systems” on page 631, or “Working with SSL or TLS on z/OS” on page 674.

This collection of topics introduces the tasks involved in setting up SSL communications, and provides step-by-step guidance on completing those tasks.

You might also want to test SSL or TLS client authentication, which are an optional part of the protocols. During the SSL or TLS handshake, the SSL or TLS client always obtains and validates a digital certificate from the server. With the WebSphere MQ implementation, the SSL or TLS server always requests a certificate from the client.

### Notes:

1. In this context, an SSL client refers to the connection initiating the handshake.
2. See the  Glossary (*WebSphere MQ V7.1 Product Overview Guide*) for further details.

On IBM i, UNIX, Linux and Windows systems, the SSL or TLS client sends a certificate only if it has one labeled in the correct WebSphere MQ format, which is `ibmwebspheremq` followed by the name of your queue manager changed to lowercase. For example, for QM1, `ibmwebspheremqqm1`.

On z/OS, the SSL or TLS client sends a certificate only if it has one of the following certificates:

- For a shared channel only, a certificate with a label in the format `ibmWebSphereMQ` followed by the name of your queue-sharing group, for example `ibmWebSphereMQQSG1`
- A certificate with a label in the format `ibmWebSphereMQ` followed by the name of your queue manager, for example `ibmWebSphereMQQM1`
- A default certificate (which might be the `ibmWebSphereMQ` certificate).

If the channel is shared, the channel first tries to find a certificate for the queue-sharing group. If it does not find a certificate for a queue-sharing group, it tries to find a certificate for the queue manager.

WebSphere MQ uses the `ibmwebspheremq` prefix on a label to avoid confusion with certificates for other products. Ensure that you specify the entire certificate label in lowercase.

The SSL or TLS server always validates the client certificate if one is sent. If the client does not send a certificate, authentication fails only if the end of the channel that is acting as the SSL or TLS server is

defined with either the SSLCAUTH parameter set to REQUIRED or an SSLPEER parameter value set. For more information about connecting a queue manager anonymously, that is, when the SSL or TLS client does not send a certificate, see “Connecting two queue managers using one-way authentication” on page 766.

## **Digital certificate labels, understanding the requirements**

When setting up SSL and TLS to use digital certificates, there might be specific label requirements that you must follow, depending on the platform used and the method you use to connect.

### **About this task**

#### **What is the certificate label?**

A certificate label is a unique identifier representing a digital certificate stored in a key repository, and provides a convenient human-readable name with which to refer to a particular certificate when performing key management functions. You assign the certificate label when adding a certificate to a key repository for the first time.

The certificate label is separate from the certificate's *Subject Distinguished Name* or *Subject Common Name* fields. Note that the *Subject Distinguished Name* and *Subject Common Name* are fields within the certificate itself. These are defined when the certificate is created and cannot be changed. However, you can change the label associated with a digital certificate if necessary.

#### **How is the certificate label used?**

IBM WebSphere MQ uses certificate labels to locate a personal certificate that is sent during the SSL handshake. This eliminates ambiguity when more than one personal certificate exists in the key repository.

Certificate labels follow a naming convention; you need to ensure that you use the correct label naming convention corresponding to the platform that you are using.

In this context, an SSL or TLS client refers to the connection partner initiating the handshake, which might be a IBM WebSphere MQ client or another queue manager.

During the SSL or TLS handshake, the SSL or TLS client always obtains and validates a digital certificate from the server. With the IBM WebSphere MQ implementation, the SSL or TLS server always requests a certificate from the client and the client always provide a certificate to the server if a one is found. If the client is unable to locate a personal certificate, the client sends a no certificate response to the server.

The SSL or TLS server always validates the client certificate if one is sent. If the client does not send a certificate, authentication fails if the end of the channel that is acting as the SSL or TLS server is defined with either the SSLCAUTH parameter set to REQUIRED or an SSLPEER parameter value set.

For more information about connecting a queue manager using one-way authentication, that is, when the SSL or TLS client does not send a certificate, see “Connecting two queue managers using one-way authentication” on page 766.

## **IBM i, UNIX, Linux, and Windows systems:**

### **About this task**

On IBM i, UNIX, Linux, and Windows systems, the SSL or TLS server sends a certificate to the client, only if the server finds one labeled in the correct IBM WebSphere MQ format. On these systems, the correct format is `ibmwebspheremq`, followed by the name of your queue manager changed to lowercase.

For example, for a queue manager named QM1, the certificate label requirement is:

```
ibmwebspheremqmqm1
```

If no certificate is found in the queue manager's key repository, matching the required label in the correct case and format, an error occurs and the SSL or TLS handshake fails.

## **z/OS®:**

### **About this task**

WebSphere MQ Clients are not supported on IBM z/OS systems. However, the z/OS queue manager can act in the role of the SSL or TLS client when initiating a connection, or the SSL or TLS server when accepting a connection request. Certificate label requirements for z/OS queue managers apply in both of these roles, and differ from the requirements on distributed platforms.

On z/OS systems, the correct format for labels is `ibmWebSphereMQ` in mixed case, followed by the name of the queue manager or queue sharing group in uppercase as described below.

On z/OS, the queue manager sends a certificate, only if it has one of the following certificates:

- For a shared channel only, a certificate with a label in the format `ibmWebSphereMQ`, followed by the name of your queue-sharing group, for example, `ibmWebSphereMQQSG1`
- A certificate with a label in the format `ibmWebSphereMQ` followed by the name of your queue manager, for example `ibmWebSphereMQQM1`
- The default certificate, if one is specified and no match has been found, based on the label.

If the channel is shared, the channel first tries to find a certificate for the queue-sharing group. If the channel does not find a certificate for a queue-sharing group, the channel tries to find a certificate for the queue manager.

## **WebSphere MQ client:**

### **About this task**

When connecting from a IBM WebSphere MQ client application, the SSL or TLS client sends a certificate only if it has one a certificate with a label in the format `ibmwebspheremq`, followed by the username of the user running the client application process.

For example, for the username `wasadmin`, the certificate label requirement is as shown, folded to lowercase:

```
ibmwebspheremqwasadmin
```

The above label requirement applies to Message Service Clients for C, or C++, and .NET.

## **IBM WebSphere MQ Java or IBM WebSphere MQ JMS client:**

### **About this task**

IBM WebSphere MQ Java or IBM WebSphere MQ JMS clients use the facilities of their Java Secure Socket Extension (JSSE) provider to select a personal certificate during the SSL or TLS handshake and therefore are not subject to certificate label requirements.

The default behavior, is that the JSSE client iterates through the certificates in the key repository, selecting the first acceptable personal certificate found. However, this behavior is only a default, and is dependent on the implementation of the JSSE provider.

In addition, the JSSE interface is highly customizable through configuration and direct access at runtime by the application. Consult the documentation supplied by your JSSE provider for specific details.

For troubleshooting, or to better understand the handshake performed by the IBM WebSphere MQ Java client application in combination with your specific JSSE provider, you can enable debugging by setting `javax.net.debug=ssl`

in the JVM environment.

You can use `-Djavax.net.debug=ssl` on the command line, or set the variable within the application, or through configuration.

#### Related concepts:

“Importing a personal certificate into a key repository on UNIX, Linux or Windows systems” on page 650

## Using self-signed certificates for mutual authentication of two queue managers

Follow these sample instructions to implement mutual authentication between two queue managers, using self-signed SSL or TLS certificates.

### About this task

Scenario:

- You have two queue managers, QM1 and QM2, which need to communicate securely. You require mutual authentication to be carried out between QM1 and QM2.
- You have decided to test your secure communication using self-signed certificates.

The resulting configuration looks like this:

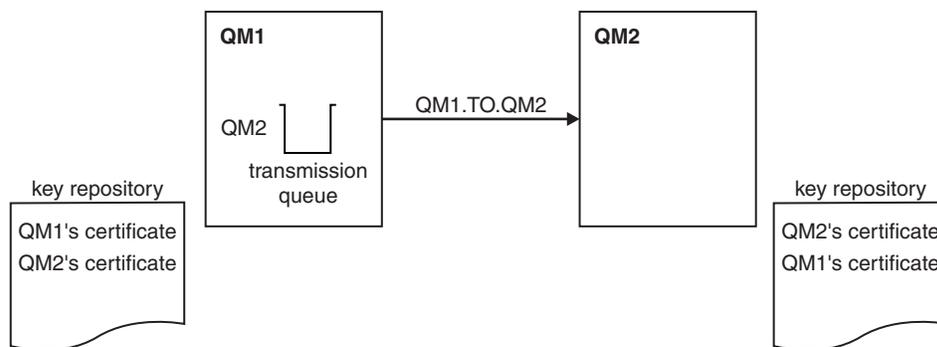


Figure 95. Configuration resulting from this task

In Figure 95, the key repository for QM1 contains the certificate for QM1 and the public certificate from QM2. The key repository for QM2 contains the certificate for QM2 and the public certificate from QM1.

### Procedure

1. Prepare the key repository on each queue manager, according to operating system:
  - On UNIX, Linux, and Windows systems.

- On z/OS systems.
2. Create a self-signed certificate for each queue manager:
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems.
  3. Extract a copy of each certificate:
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems.
  4. Transfer the public part of the QM1 certificate to the QM2 system and vice versa, using a utility such as FTP, as described in “Exchanging self-signed certificates” on page 683.
  5. Add the partner certificate to the key repository for each queue manager:
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems.
  6. On QM1, define a sender channel and associated transmission queue, by issuing commands like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QM1.MACH.COM)
XMITQ(QM2)
SSLCIPH(RC4_MD5_US) DESCR('Sender channel using SSL from QM1 to QM2')

DEFINE QLOCAL(QM2) USAGE(XMITQ)
```

This example uses CipherSpec RC4\_MD5. The CipherSpecs at each end of the channel must be the same.

7. On QM2, define a receiver channel, by issuing a command like the following example:

```
DEFINE CHANNEL(QM1.TO.QM2) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC4_MD5_US)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from QM1 to QM2')
```

The channel must have the same name as the sender channel you defined in step 6, and use the same CipherSpec.

8. Start the channel, as described in “Starting the sender channel” on page 684.

## Results

Key repositories and channels are created as illustrated in Figure 95 on page 785

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following examples.

From queue manager QM1, enter the following command:

```
DISPLAY CHS(QM1.TO.QM2) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(QM1.TO.QM2) CHLTYPE(SDR)
CONNAME(9.20.25.40) CURRENT
RQMNAME(QM2)
SSLCERTI("CN=QM2,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5E:02,CN=QM2,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(MQGET)
XMITQ(QM2)
```

From queue manager QM2, enter the following command:

```
DISPLAY CHS(QM1.TO.QM2) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(QM1.TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(QM2.TO.QM1) CHLTYPE(RCVR)
CONNNAME(9.20.35.92) CURRENT
RQMNAME(QM1)
SSLCERTI("CN=QM1,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QM1,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
XMITQ()
```

In each case, the value of SSLPEER must match that of the DN in the partner certificate that was created in Step 2. The issuer's name matches the peer name because the certificate is self-signed.

SSLPEER is optional. If it is specified, its value must be set so that the DN in the partner certificate (created in step 2) is allowed. For more information about the use of SSLPEER, see  [WebSphere MQ rules for SSLPEER values \(WebSphere MQ V7.1 Reference\)](#).

## Using CA-signed certificates for mutual authentication of two queue managers

Follow these sample instructions to implement mutual authentication between two queue managers, using CA-signed SSL or TLS certificates.

### About this task

Scenario:

- You have two queue managers called QMA and QMB, which need to communicate securely. You require mutual authentication to be carried out between QMA and QMB.
- In the future you are planning to use this network in a production environment, and therefore you have decided to use CA-signed certificates from the beginning.

The resulting configuration looks like this:

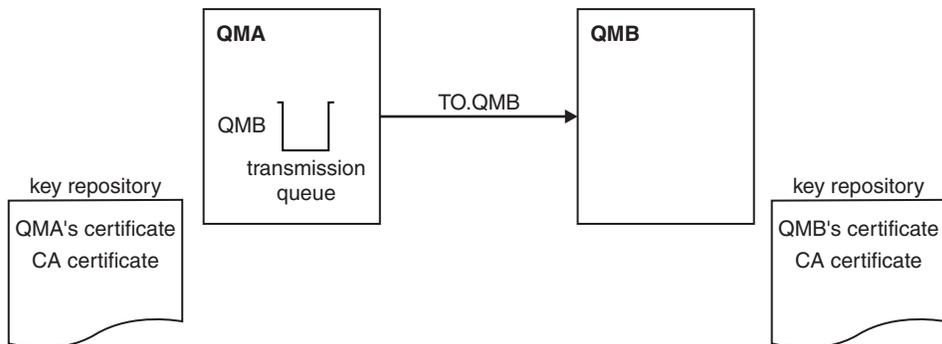


Figure 96. Configuration resulting from this task

In Figure 96, the key repository for QMA contains QMA's certificate and the CA certificate. The key repository for QMB contains QMB's certificate and the CA certificate. In this example both QMA's certificate and QMB's certificate were issued by the same CA. If QMA's certificate and QMB's certificate

were issued by different CAs then the key repositories for QMA and QMB must contain both CA certificates.

## Procedure

1. Prepare the key repository on each queue manager, according to operating system:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
2. Request a CA-signed certificate for each queue manager. You might use different CAs for the two queue managers.
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
3. Add the Certificate Authority certificate to the key repository for each queue manager: If the Queue managers are using different Certificate Authorities then the CA certificate for each Certificate Authority must be added to both key repositories.
  - Do not perform this step on IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.
4. Add the CA-signed certificate to the key repository for each queue manager:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.

5. On QMA, define a sender channel and associated transmission queue by issuing commands like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(SDR) TRPTYPE(TCP) CONNAME(QMB.MACH.COM)
XMITQ(QMB)
SSLCIPH(RC2_MD5_EXPORT) DESC('Sender channel using SSL from QMA to QMB')

DEFINE QLOCAL(QMB) USAGE(XMITQ)
```

This example uses CipherSpec RC4\_MD5. The CipherSpecs at each end of the channel must be the same.

6. On QMB, define a receiver channel by issuing a command like the following example:

```
DEFINE CHANNEL(TO.QMB) CHLTYPE(RCVR) TRPTYPE(TCP) SSLCIPH(RC2_MD5_EXPORT)
SSLCAUTH(REQUIRED) DESC('Receiver channel using SSL to QMB')
```

The channel must have the same name as the sender channel you defined in step 6, and use the same CipherSpec.

7. Start the channel:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems.

## Results

Key repositories and channels are created as illustrated in Figure 96 on page 787.

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is like that shown in the following examples.

From queue manager QMA, enter the following command:

```
DISPLAY CHS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(TO.QMB) CHLTYPE(SDR)
CONNAME(9.20.25.40) CURRENT
RQMNAME(QMB)
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMB,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(MQGET)
XMITQ(QMB)
```

From the queue manager QMB, enter the following command:

```
DISPLAY CHS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(TO.QMB) CHLTYPE(RCVR)
CONNAME(9.20.35.92) CURRENT
RQMNAME(QMA)
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMA,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
XMITQ()
```

In each case, the value of SSLPEER must match that of the Distinguished Name (DN) in the partner certificate that was created in Step 2. The issuer name matches the subject DN of the CA certificate that signed the personal certificate added in Step 4.

## Connecting two queue managers using one-way authentication

Follow these sample instructions to modify a system with mutual authentication to allow a queue manager to connect using one-way authentication to another; that is, when the SSL or TLS client does not send a certificate.

### About this task

Scenario:

- Your two queue managers (QM1 and QM2) have been set up as in “Using CA-signed certificates for mutual authentication of two queue managers” on page 764.
- You want to change QM1 so that it connects using one-way authentication to QM2.

The resulting configuration looks like this:

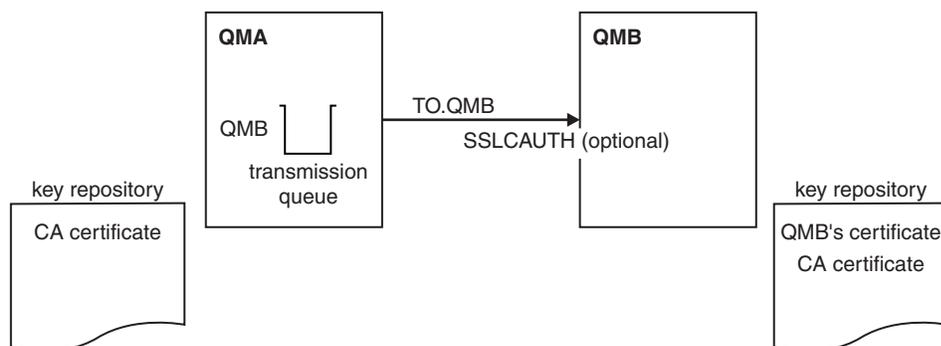


Figure 97. Queue managers allowing one-way authentication

## Procedure

1. Remove QM1's personal certificate from its key repository, according to operating system:
  - On IBM i systems. The certificate is labeled as follows:
    - `ibmwebspheremq` followed by the name of your queue manager folded to lowercase. For example, for QM1, `ibmwebspheremqqm1`.
  - On UNIX, Linux, and Windows systems. The certificate is labeled as follows:
    - `ibmwebspheremq` followed by the name of your queue manager folded to lowercase. For example, for QM1, `ibmwebspheremqqm1`.
  - On z/OS systems. Perform this step twice, to remove the personal certificate for QMA and the default certificate. The certificates are labeled as follows:
    - `ibmWebSphereMQ` followed by the name of your queue manager or queue-sharing group. For example, for QM1, `ibmWebSphereMQM1`.
    - The default certificate (which might be the `ibmWebSphereMQ` certificate).
2. Optional: On QM1, if any SSL or TLS channels have run previously, refresh the SSL or TLS environment, as described in “Refreshing the SSL or TLS environment” on page 684.
3. Allow anonymous connections on the receiver, as described in “Allowing anonymous connections on a receiver channel” on page 684.

## Results

Key repositories and channels are changed as illustrated in Figure 97

## What to do next

If the sender channel was running and you issued the `REFRESH SECURITY TYPE(SSL)` command (in step 2), the channel restarts automatically. If the sender channel was not running, start it.

At the server end of the channel, the presence of the peer name parameter value on the channel status display indicates that a client certificate has flowed.

Verify that the task has been completed successfully by issuing some `DISPLAY` commands. If the task was successful, the resulting output is similar to that shown in the following examples:

From the QM1 queue manager, enter the following command:

```
DISPLAY CHS(TO.QM2) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
 4 : DISPLAY CHSTATUS(TO.QMB) SSLPEER
AMQ8417: Display Channel Status details.
 CHANNEL(TO.QM2) CHLTYPE(SDR)
 CONNAME(9.20.25.40) CURRENT
 RQMNAME(QM2)
 SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
 SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMB,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
 STATUS(RUNNING) SUBSTATE(MQGET)
 XMITQ(QM2)
```

From the QM2 queue manager, enter the following command:

```
DISPLAY CHS(TO.QM2) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```
DISPLAY CHSTATUS(TO.QM2) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(TO.QM2) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
 CHANNEL(TO.QM2) CHLTYPE(RCVR)
 CONNAME(9.20.35.92) CURRENT
 RQMNAME(QMA) SSLCERTI()
 SSLPEER() STATUS(RUNNING)
 SUBSTATE(RECEIVE) XMITQ()
```

On QM2, the SSLPEER field is empty, showing that QM1 did not send a certificate. On QM1, the value of SSLPEER matches that of the DN in QM2's personal certificate.

## Connecting a client to a queue manager securely

Secure communications that use the SSL or TLS cryptographic security protocols involve setting up the communication channels and managing the digital certificates that you will use for authentication.

To set up your SSL or TLS installation you must define your channels to use SSL or TLS. You must also obtain and manage your digital certificates. On a test system, you can use self-signed certificates or certificates issued by a local certificate authority (CA). On a production system, do not use self-signed certificates. For more information, see "Protecting channels with SSL" on page 449.

For full information about creating and managing certificates, see "Working with SSL or TLS on IBM i" on page 619, "Working with SSL or TLS on UNIX, Linux and Windows systems" on page 631, or "Working with SSL or TLS on z/OS" on page 674.

This collection of topics introduces the tasks involved in setting up SSL communications, and provides step-by-step guidance on completing those tasks.

You might also want to test SSL or TLS client authentication, which are an optional part of the protocols. During the SSL or TLS handshake, the SSL or TLS client always obtains and validates a digital certificate from the server. With the WebSphere MQ implementation, the SSL or TLS server always requests a certificate from the client.

On IBM i, UNIX, Linux, and Windows systems, the SSL or TLS client sends a certificate only if it has one labeled in the correct WebSphere MQ format, which is `ibmwebsphermq` followed by your logon user ID changed to lowercase, for example `ibmwebsphermqmyuserid`.

WebSphere MQ uses the `ibmwebsphermq` prefix on a label to avoid confusion with certificates for other products. Ensure that you specify the entire certificate label in lowercase.

The SSL or TLS server always validates the client certificate if one is sent. If the client does not send a certificate, authentication fails only if the end of the channel that is acting as the SSL or TLS server is defined with either the SSLCAUTH parameter set to REQUIRED or an SSLPEER parameter value set. For more information about connecting a queue manager anonymously, see “Connecting a client to a queue manager anonymously” on page 773.

## Using self-signed certificates for mutual authentication of a client and queue manager

Follow these sample instructions to implement mutual authentication between a client and a queue manager, by using self-signed SSL or TLS certificates.

### About this task

Scenario:

- You have a client, C1, and a queue manager, QM1, which need to communicate securely. You require mutual authentication to be carried out between C1 and QM1.
- You have decided to test your secure communication by using self-signed certificates.

DCM on IBM i does not support self-signed certificates, so this task is not applicable on IBM i systems.

The resulting configuration looks like this:

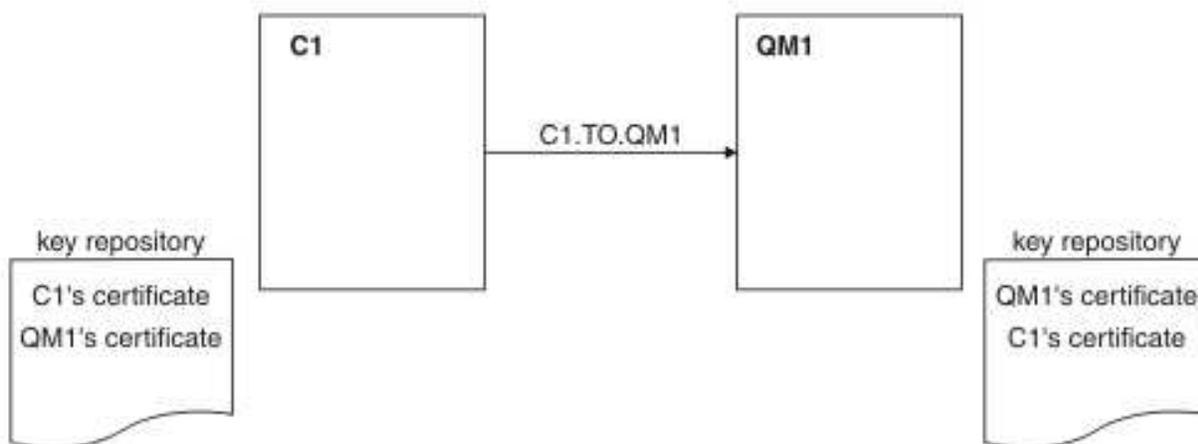


Figure 98. Configuration resulting from this task

In Figure 98, the key repository for QM1 contains the certificate for QM1 and the public certificate from C1. The key repository for C1 contains the certificate for C1 and the public certificate from QM1.

### Procedure

1. Prepare the key repository on the client and queue manager, according to operating system:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
2. Create self-signed certificates for the client and queue manager:
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
3. Extract a copy of each certificate:

- On UNIX, Linux, and Windows systems.
  - On z/OS systems.
4. Transfer the public part of the C1 certificate to the QM1 system and vice versa, using a utility such as FTP, as described in “Exchanging self-signed certificates” on page 683.
  5. Add the partner certificate to the key repository for the client and queue manager:
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems.
  6. Define a client-connection channel in either of the following ways:
    - Using the MQCONNX call with the MQSCO structure on C1, as described in  [Creating a client-connection channel on the WebSphere MQ MQI client \(WebSphere MQ V7.1 Installing Guide\)](#).
    - Using a client channel definition table, as described in  [Creating server-connection and client-connection definitions on the server \(WebSphere MQ V7.1 Installing Guide\)](#).
  7. On QM1, define a server-connection channel, by issuing a command like the following example:
 

```
DEFINE CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(RC4_MD5_US)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from C1 to QM1')
```

The channel must have the same name as the client-connection channel you defined in step 6, and use the same CipherSpec.

## Results

Key repositories and channels are created as illustrated in Figure 98 on page 792

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following example.

From queue manager QM1, enter the following command:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNAME(9.20.35.92) CURRENT
SSLCERTI("CN=QM1,OU=WebSphere MQ Development,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5E:02,CN=QM2,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
```

It is optional to set the SSLPEER filter attribute of the channel definitions. If the channel definition SSLPEER is set, its value must match the subject DN in the partner certificate that was created in Step 2. After a successful connection, the SSLPEER field in the DISPLAY CHSTATUS output shows the subject DN of the remote client certificate.

## Using CA-signed certificates for mutual authentication of a client and queue manager

Follow these sample instructions to implement mutual authentication between a client and a queue manager, by using CA-signed SSL or TLS certificates.

### About this task

Scenario:

- You have a client, C1, and a queue manager, QM1, which need to communicate securely. You require mutual authentication to be carried out between C1 and QM1.
- In the future you are planning to use this network in a production environment, and therefore you have decided to use CA-signed certificates from the beginning.

The resulting configuration looks like this:

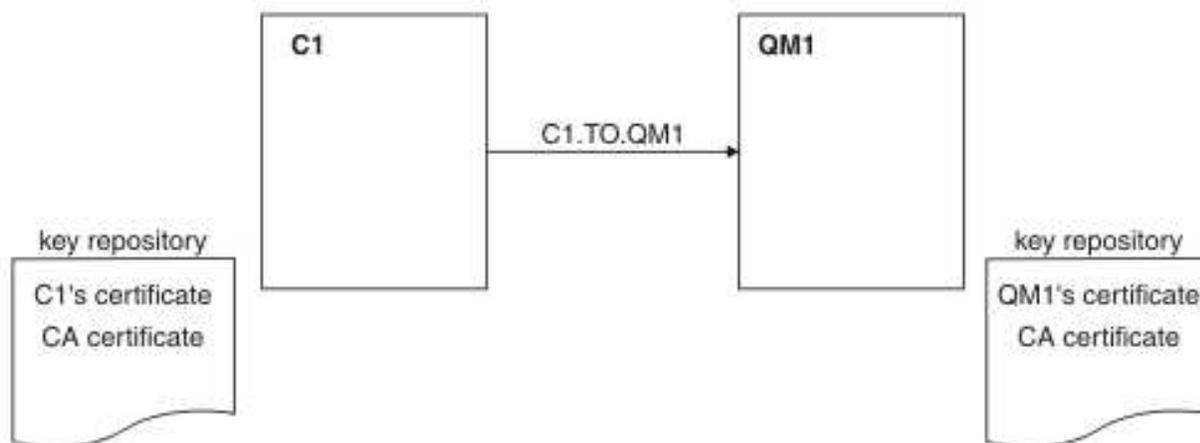


Figure 99. Configuration resulting from this task

In Figure 99, the key repository for C1 contains certificate for C1 and the CA certificate. The key repository for QM1 contains the certificate for QM1 and the CA certificate. In this example both C1's certificate and QM1's certificate were issued by the same CA. If C1's certificate and QM1's certificate were issued by different CAs then the key repositories for C1 and QM1 must contain both CA certificates.

### Procedure

1. Prepare the key repository on the client and queue manager, according to operating system:
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
2. Request a CA-signed certificate for the client and queue manager. You might use different CAs for the client and queue manager.
  - On IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
3. Add the certificate authority certificate to the key repository for the client and queue manager. If the client and queue manager are using different Certificate Authorities then the CA certificate for each Certificate Authority must be added to both key repositories.

- Do not perform this step on IBM i systems.
  - On UNIX, Linux, and Windows systems.
  - On z/OS systems (queue manager only).
4. Add the CA-signed certificate to the key repository for the client and queue manager:
    - On IBM i systems.
    - On UNIX, Linux, and Windows systems.
    - On z/OS systems (queue manager only).
  5. Define a client-connection channel in either of the following ways:
    - Using the MQCONN call with the MQSCO structure on C1, as described in  [Creating a client-connection channel on the WebSphere MQ MQI client \(WebSphere MQ V7.1 Installing Guide\)](#).
    - Using a client channel definition table, as described in  [Creating server-connection and client-connection definitions on the server \(WebSphere MQ V7.1 Installing Guide\)](#).
  6. On QM1, define a server-connection channel by issuing a command like the following example:
 

```
DEFINE CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(RC2_MD5_EXPORT)
SSLCAUTH(REQUIRED) DESCR('Receiver channel using SSL from C1 to QM1')
```

The channel must have the same name as the client-connection channel you defined in step 6, and use the same CipherSpec.

## Results

Key repositories and channels are created as illustrated in Figure 99 on page 794.

## What to do next

Check that the task has been completed successfully by using DISPLAY commands. If the task was successful, the resulting output is like that shown in the following example.

From the queue manager QM1, enter the following command:

```
DISPLAY CHSTATUS(TO.QMB) SSLPEER SSLCERTI
```

The resulting output is like the following example:

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNAME(9.20.35.92) CURRENT
SSLCERTI("CN=WebSphere MQ CA,OU=WebSphere MQ Devt,O=IBM,ST=Hampshire,C=UK")
SSLPEER("SERIALNUMBER=4C:D0:49:D5:02:5F:38,CN=QMA,OU=WebSphere MQ Development,
O=IBM,ST=Hampshire,C=UK")
STATUS(RUNNING) SUBSTATE(RECEIVE)
```

The SSLPEER field in the DISPLAY CHSTATUS output shows the subject DN of the remote client certificate that was created in Step 2. The issuer name matches the subject DN of the CA certificate that signed the personal certificate added in Step 4.

## Connecting a client to a queue manager anonymously

Follow these sample instructions to modify a system with mutual authentication to allow a queue manager to connect anonymously to another.

### About this task

Scenario:

- Your queue manager and client (QM1 and C1) have been set up as in “Using CA-signed certificates for mutual authentication of a client and queue manager” on page 771.
- You want to change C1 so that it connects anonymously to QM1.

The resulting configuration looks like this:

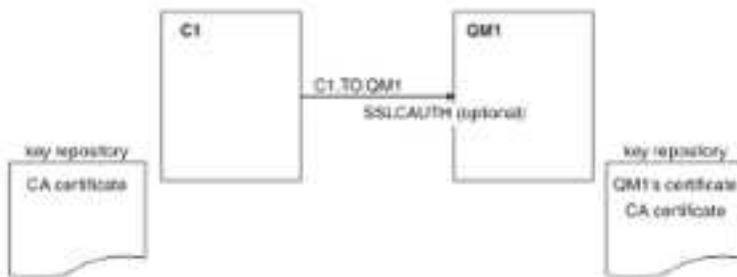


Figure 100. Client and queue manager allowing anonymous connection

### Procedure

1. Remove the personal certificate from key repository for C1, according to operating system:
  - On IBM i systems. The certificate is labeled as follows:
    - `ibmwebsphermq` followed by your logon user ID folded to lowercase, for example `ibmwebsphermqmyuserid`.
  - On UNIX, Linux, and Windows systems. The certificate is labeled as follows:
    - `ibmwebsphermq` followed by your logon user ID folded to lowercase, for example `ibmwebsphermqmyuserid`.
2. Restart the client application, or cause the client application to close and reopen all SSL or TLS connections.
3. Allow anonymous connections on the queue manager, by issuing the following command:  

```
ALTER CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN) SSLCAUTH(OPTIONAL)
```

### Results

Key repositories and channels are changed as illustrated in Figure 100

### What to do next

At the server end of the channel, the presence of the peer name parameter value on the channel status display indicates that a client certificate has flowed.

Verify that the task has been completed successfully by issuing some DISPLAY commands. If the task was successful, the resulting output is similar to that shown in the following example:

From queue manager QM1, enter the following command:  

```
DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
```

The resulting output will be similar to the following example:

```

DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
 5 : DISPLAY CHSTATUS(C1.TO.QM1) SSLPEER SSLCERTI
AMQ8417: Display Channel Status details.
CHANNEL(C1.TO.QM1) CHLTYPE(SVRCONN)
CONNAME(9.20.35.92) CURRENT
SSLCERTI() SSLPEER()
STATUS(RUNNING) SUBSTATE(RECEIVE)

```

The SSLCERTI and SSLPEER fields are empty, showing that C1 did not send a certificate.

## Specifying CipherSpecs

Specify a CipherSpec by using the **SSLCIPH** parameter in either the **DEFINE CHANNEL** MQSC command or the **ALTER CHANNEL** MQSC command.

Some of the CipherSpecs that you can use with IBM WebSphere MQ are FIPS compliant. Others, such as NULL\_MD5, are not. Similarly, some of the FIPS compliant CipherSpecs are also Suite B compliant although others, are not. All Suite B compliant CipherSpecs are also FIPS compliant. All Suite B compliant CipherSpecs fall into two groups: 128 bit (for example, ECDHE\_ECDSA\_AES\_128\_GCM\_SHA256) and 192 bit (for example, ECDHE\_ECDSA\_AES\_256\_GCM\_SHA384),

The following diagram illustrates the relationship between these subsets:



Cipher specifications that you can use with IBM WebSphere MQ SSL and TLS support are listed in the following table. When you request a personal certificate, you specify a key size for the public and private key pair. The key size that is used during the SSL handshake is the size stored in the certificate unless it is determined by the CipherSpec, as noted in the table.

| CipherSpec name                 | Protocol used | Data integrity | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|---------------------------------|---------------|----------------|----------------------|-----------------|-------------------|---------|
| AES_SHA_US                      | SSL 3.0       | SHA-1          | AES                  | 128             | No                | No      |
| DES_SHA_EXPORT <sup>2</sup>     | SSL 3.0       | SHA-1          | DES                  | 56              | No                | No      |
| DES_SHA_EXPORT1024 <sup>3</sup> | SSL 3.0       | SHA-1          | DES                  | 56              | No                | No      |
| FIPS_WITH_DES_CBC_SHA           | SSL 3.0       | SHA-1          | DES                  | 56              | No <sup>6</sup>   | No      |
| NULL_MD5 <sup>a</sup>           | SSL 3.0       | MD5            | None                 | 0               | No                | No      |
| NULL_SHA <sup>a</sup>           | SSL 3.0       | SHA-1          | None                 | 0               | No                | No      |
| RC2_MD5_EXPORT <sup>2 a</sup>   | SSL 3.0       | MD5            | RC2                  | 40              | No                | No      |
| RC4_MD5_EXPORT <sup>2 a</sup>   | SSL 3.0       | MD5            | RC4                  | 40              | No                | No      |
| RC4_MD5_US <sup>a</sup>         | SSL 3.0       | MD5            | RC4                  | 128             | No                | No      |
| RC4_SHA_US <sup>a</sup>         | SSL 3.0       | SHA-1          | RC4                  | 128             | No                | No      |

| Cipherspec name                               | Protocol used | Data integrity         | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|-----------------------------------------------|---------------|------------------------|----------------------|-----------------|-------------------|---------|
| RC4_56_SHA_EXPORT1024 <sup>3 b</sup>          | SSL 3.0       | SHA-1                  | RC4                  | 56              | No                | No      |
| TLS_RSA_EXPORT_WITH_RC2_40_MD5 <sup>c</sup>   | TLS 1.0       | MD5                    | RC2                  | 40              | No                | No      |
| TLS_RSA_EXPORT_WITH_RC4_40_MD5 <sup>2 c</sup> | TLS 1.0       | MD5                    | RC4                  | 40              | No                | No      |
| TLS_RSA_WITH_AES_128_CBC_SHA <sup>a</sup>     | TLS 1.0       | SHA-1                  | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_CBC_SHA <sup>4 a</sup>   | TLS 1.0       | SHA-1                  | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_DES_CBC_SHA <sup>a</sup>         | TLS 1.0       | SHA-1                  | DES                  | 56              | No <sup>5</sup>   | No      |
| TLS_RSA_WITH_NULL_MD5 <sup>c</sup>            | TLS 1.0       | MD5                    | None                 | 0               | No                | No      |
| TLS_RSA_WITH_NULL_SHA <sup>c</sup>            | TLS 1.0       | SHA-1                  | None                 | 0               | No                | No      |
| TLS_RSA_WITH_RC4_128_MD5 <sup>c</sup>         | TLS 1.0       | MD5                    | RC4                  | 128             | No                | No      |
| ECDHE_ECDSA_AES_128_CBC_SHA256 <sup>d</sup>   | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| ECDHE_ECDSA_AES_256_CBC_SHA384 <sup>d</sup>   | TLS 1.2       | SHA-384                | AES                  | 256             | Yes               | No      |
| ECDHE_ECDSA_AES_128_GCM_SHA256 <sup>b</sup>   | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | 128 bit |
| ECDHE_ECDSA_AES_256_GCM_SHA384 <sup>b</sup>   | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | 192 bit |
| ECDHE_ECDSA_NULL_SHA256 <sup>b</sup>          | TLS 1.2       | SHA-1                  | None                 | 0               | No                | No      |
| ECDHE_ECDSA_RC4_128_SHA256 <sup>b</sup>       | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |
| ECDHE_RSA_AES_128_CBC_SHA256 <sup>d</sup>     | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| ECDHE_RSA_AES_256_CBC_SHA384 <sup>d</sup>     | TLS 1.2       | SHA-384                | AES                  | 256             | Yes               | No      |
| ECDHE_RSA_AES_128_GCM_SHA256 <sup>b</sup>     | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | No      |
| ECDHE_RSA_AES_256_GCM_SHA384 <sup>b</sup>     | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | No      |
| ECDHE_RSA_NULL_SHA256 <sup>b</sup>            | TLS 1.2       | SHA-1                  | None                 | 0               | No                | No      |
| ECDHE_RSA_RC4_128_SHA256 <sup>b</sup>         | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |
| TLS_RSA_WITH_AES_128_CBC_SHA256 <sup>d</sup>  | TLS 1.2       | SHA-256                | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_CBC_SHA256 <sup>d</sup>  | TLS 1.2       | SHA-256                | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_AES_128_GCM_SHA256 <sup>b</sup>  | TLS 1.2       | AEAD<br>AES-128<br>GCM | AES                  | 128             | Yes               | No      |
| TLS_RSA_WITH_AES_256_GCM_SHA384 <sup>b</sup>  | TLS 1.2       | AEAD<br>AES-256<br>GCM | AES                  | 256             | Yes               | No      |
| TLS_RSA_WITH_NULL_NULL <sup>b</sup>           | TLS 1.2       | None                   | None                 | 0               | No                | No      |
| TLS_RSA_WITH_NULL_SHA256 <sup>d</sup>         | TLS 1.2       | SHA-256                | None                 | 0               | No                | No      |
| TLS_RSA_WITH_RC4_128_SHA256 <sup>b</sup>      | TLS 1.2       | SHA-1                  | RC4                  | 128             | No                | No      |

| CipherSpec name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Protocol used | Data integrity | Encryption algorithm | Encryption bits | FIPS <sup>1</sup> | Suite B |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------------|----------------------|-----------------|-------------------|---------|
| <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. Specifies whether the CipherSpec is FIPS-certified on a FIPS-certified platform. See “Federal Information Processing Standards (FIPS)” on page 393 for an explanation of FIPS.</li> <li>2. The maximum handshake key size is 512 bits. If either of the certificates exchanged during the SSL handshake has a key size greater than 512 bits, a temporary 512-bit key is generated for use during the handshake.</li> <li>3. The handshake key size is 1024 bits.</li> <li>4. This CipherSpec cannot be used to secure a connection from the WebSphere MQ Explorer to a queue manager unless the appropriate unrestricted policy files are applied to the JRE used by the Explorer.</li> <li>5. This CipherSpec was FIPS 140-2 certified before 19 May 2007.</li> <li>6. This CipherSpec was FIPS 140-2 certified before 19 May 2007. The name FIPS_WITH_DES_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.</li> <li>7. This CipherSpec can be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, either avoid using triple DES, or enable secret key reset when using this CipherSpec.</li> </ol> <p><b>Platform support:</b></p> <ul style="list-style-type: none"> <li>• a Available on all supported platforms.</li> <li>• b Available only on UNIX, Linux, and Windows platforms.</li> <li>• c Available only on IBM i platforms.</li> <li>• d Available on UNIX, Linux, and Windows platforms, z/OS, and IBM i V7R1M0 TR6 or later.</li> </ul> <p>For further information, on using these CipherSpecs on the IBM i platform, see  Using TLS Version 1.2</p> <p>To use these CipherSpecs on z/OS, you must be using z/OS V1R13, and install the IBM WebSphere MQ APAR  PM77341 in conjunction with the System SSL APAR  OA39422.</p> |               |                |                      |                 |                   |         |

### Related concepts:

“Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403

### Related reference:

 DEFINE CHANNEL (*WebSphere MQ V7.1 Reference*)

 ALTER CHANNEL (*WebSphere MQ V7.1 Reference*)

## Obtaining information about CipherSpecs using WebSphere MQ Explorer

You can use WebSphere MQ Explorer to display descriptions of CipherSpecs.

Use the following procedure to obtain information about the CipherSpecs in “Specifying CipherSpecs” on page 774:

1. Open **WebSphere MQ Explorer** and expand the **Queue Managers** folder.
2. Ensure that you have started your queue manager.
3. Select the queue manager you want to work with and click **Channels**.
4. Right-click the channel you want to work with and select **Properties**.
5. Select the **SSL** property page.
6. Select from the list the CipherSpec you want to work with. A description is displayed in the window below the list.

## Alternatives for specifying CipherSpecs

For those platforms where the operating system provides the SSL support, your system might support new CipherSpecs. You can specify a new CipherSpec with the SSLCIPH parameter, but the value you supply depends on your platform.

**Note:** This section does not apply to UNIX, Linux or Windows systems, because the CipherSpecs are provided with the WebSphere MQ product, so new CipherSpecs do not become available after shipment.

For those platforms where the operating system provides the SSL support, your system might support new CipherSpecs that are not included in “Specifying CipherSpecs” on page 774. You can specify a new CipherSpec with the SSLCIPH parameter, but the value you supply depends on your platform. In all cases the specification *must* correspond to an SSL CipherSpec that is both valid and supported by the version of SSL your system is running.

**IBM i** A two-character string representing a hexadecimal value.

For more information about the permitted values, refer to the appropriate product documentation (search on *cipher\_spec*). For all IBM i versions, start here: [🔗 IBM i](#)

You can use either the CHGMQMCHL or the CRTMQMCHL command to specify the value, for example:

```
CRTMQMCHL CHLNAME('channel name') SSLCIPH('hexadecimal value')
```

You can also use the ALTER QMGR MQSC command to set the SSLCIPH parameter.

**z/OS** A two-character string representing a hexadecimal value. The hexadecimal codes correspond to the values defined in the SSL protocol.

For more information, refer to the description of `gsk_environment_open()` in the API reference chapter of *z/OS Cryptographic Services System SSL Programming*, SC24-5901, where there is a list of all the supported SSL V3.0 and TLS V1.0 cipher specifications in the form of 2-digit hexadecimal codes.

## Considerations for WebSphere MQ clusters

With WebSphere MQ clusters it is safest to use the CipherSpec names in “Specifying CipherSpecs” on page 774. If you use an alternative specification, be aware that the specification might not be valid on other platforms. For more information, refer to “SSL and clusters” on page 806.

## Specifying a CipherSpec for a WebSphere MQ MQI client

You have three options for specifying a CipherSpec for a WebSphere MQ MQI client.

These options are as follows:

- Using a channel definition table
- Using the SSL configuration options structure, MQSCO, on an MQCONN call
- Using the Active Directory (on Windows systems with Active Directory support)

## Specifying a CipherSuite with WebSphere MQ classes for Java and WebSphere MQ classes for JMS

WebSphere MQ classes for Java and WebSphere MQ classes for JMS specify CipherSuites differently from other platforms.

For information about specifying a CipherSuite with WebSphere MQ classes for Java, see  [Secure Sockets Layer \(SSL\) support](#) (*WebSphere MQ V7.1 Programming Guide*)

For information about specifying a CipherSuite with WebSphere MQ classes for JMS, see  [Using Secure Sockets Layer \(SSL\) with WebSphere MQ classes for JMS](#) (*WebSphere MQ V7.1 Programming Guide*)

---

## Auditing

You can check for security intrusions, or attempted intrusions, by using event messages. You can also check the security of your system by using the IBM WebSphere MQ Explorer.

To detect attempts to perform unauthorized actions such as connecting to a queue manager or put a message on a queue, inspect the event messages produced by your queue managers, particularly authority event messages. For more information about queue manager event messages, see “Queue manager events” on page 834, and for more information about event monitoring in general, see “Event monitoring” on page 829.

---

## Keeping clusters secure

Authorize or prevent queue managers joining clusters or putting messages on cluster queues. Force a queue manager to leave a cluster. Take account of some additional considerations when configuring SSL for clusters.

## Stopping unauthorized queue managers sending messages

Prevent unauthorized queue managers sending messages to your queue manager using a channel security exit.

### Before you begin

Clustering has no effect on the way security exits work. You can restrict access to a queue manager in the same way as you would in a distributed queuing environment.

### About this task

Prevent selected queue managers from sending messages to your queue manager:

### Procedure

1. Define a channel security exit program on the CLUSRCVR channel definition.
2. Write a program that authenticates queue managers trying to send messages on your cluster-receiver channel and denies them access if they are not authorized.

### What to do next

Channel security exit programs are called at MCA initiation and termination.

### Related concepts:

“Security for queue manager clusters” on page 459

## Stopping unauthorized queue managers putting messages on your queues

Use the channel put authority attribute on the cluster-receiver channel to stop unauthorized queue managers putting messages on your queues. Authorize a remote queue manager by checking the user ID in the message using RACF on z/OS, or the OAM on other platforms.

### About this task

Use the security facilities of a platform and the access control mechanism in WebSphere MQ to control access to queues.

### Procedure

1. To prevent certain queue managers from putting messages on a queue, use the security facilities available on your platform.

For example:

- RACF or other external security managers on WebSphere MQ for z/OS
- The object authority manager (OAM) on other platforms.

2. Use the put authority, PUTAUT, attribute on the CLUSRCVR channel definition.

The PUTAUT attribute allows you to specify what user identifiers are to be used to establish authority to put a message to a queue.

The options on the PUTAUT attribute are:

**DEF** Use the default user ID. On z/OS, the check might involve using both the user ID received from the network and that derived from MCAUSER.

**CTX** Use the user ID in the context information associated with the message. On z/OS the check might involve using either the user ID received from the network, or that derived from MCAUSER, or both. Use this option if the link is trusted and authenticated.

**ONLYMCA (z/OS only)**

As for DEF, but any user ID received from the network is not used. Use this option if the link is not trusted. You want to allow only a specific set of actions on it, which are defined for the MCAUSER.

**ALTMCA (z/OS only)**

As for CTX, but any user ID received from the network is not used.

### Related concepts:

“Security for queue manager clusters” on page 459

## Authorizing putting messages on remote cluster queues

On z/OS set up authorization to put to a cluster queue using RACF. On other platforms, authorize access to connect to the queue manager and to put to the queue on that queue manager.

### About this task

The default behavior is to perform access control against the SYSTEM.CLUSTER.TRANSMIT.QUEUE. Note that this behavior applies, even if you are using multiple transmission queues.

The specific behavior described in this topic applies only when you have configured the

**ClusterQueueAccessControl** attribute in the `qm.ini` file to be *RQMName*, as described in the  Security stanza (*WebSphere MQ V7.1 Installing Guide*) topic, and restarted the queue manager.

## Procedure

- For z/OS, issue the following commands:  
RDEFINE MQQUEUE *QMgrName*.QUEUE.*QueueName* UACC(NONE)  
PERMIT *QMgrName*.QUEUE.*QueueName* CLASS(MQADMIN) ID(*GroupName*) ACCESS(UPDATE)
- For UNIX, Linux and Windows systems, issue the following commands:  
setmqaut -m *QMgrName* -t qmgr -g *GroupName* +connect  
setmqaut -m *QMgrName* -t queue -n *QueueName* -g *GroupName* -all +put
- For IBM i, issue the following commands:  
GRTMQMAUT OBJ(' *QMgrName* ') OBJTYPE(\*MQM) USER(*GroupName*) AUT(\*CONNECT)  
GRTMQMAUT OBJ(' *QueueName* ') OBJTYPE(\*Q) USER(*GroupName*) AUT(\*PUT)  
MQMNAME(' *QMgrName* ')

The user can put messages only to the specified cluster queue, and no other cluster queues.

The variable names have the following meanings:

### QMgrName

The name of the queue manager. On z/OS, this value can also be the name of a queue-sharing group.

### GroupName

The name of the group to be granted access.

### QueueName

Name of the queue or generic profile for which to change authorizations.

## What to do next

If you specify a reply-to queue when you put a message on a cluster queue, the consuming application must have authority to send the reply. Set this authority by following the instructions in “Granting authority to put messages to a remote cluster queue” on page 746.

### Related concepts:

“Security for queue manager clusters” on page 459

“Profiles for command security” on page 551

“Profiles for command resource security” on page 560

“Profiles to control queue-sharing group or queue manager level security” on page 525

### Related reference:

 Security stanza in qm.ini (*WebSphere MQ V7.1 Installing Guide*)

 setmqaut (*WebSphere MQ V7.1 Reference*)

 Security controls and options (*WebSphere MQ V7.1 Product Overview Guide*)

### Related information:

 Grant MQ Object Authority (GRTMQMAUT) (*WebSphere MQ V7.1 Reference*)

## Preventing queue managers joining a cluster

If a rogue queue manager joins a cluster it is difficult to prevent it receiving messages you do not want it to receive.

## Procedure

If you want to ensure that only certain authorized queue managers join a cluster you have a choice of three techniques:

- Using channel authentication records you can block the cluster channel connection based on: the remote IP address, the remote queue manager name, or the SSL/TLS Distinguished Name provided by the remote system.
- Write an exit program to prevent unauthorized queue managers from writing to SYSTEM.CLUSTER.COMMAND.QUEUE. Do not restrict access to SYSTEM.CLUSTER.COMMAND.QUEUE such that no queue manager can write to it, or you would prevent any queue manager from joining the cluster.
- A security exit program on the CLUSRCVR channel definition.

**Related concepts:**

“Security for queue manager clusters” on page 459

“Channel authentication records” on page 408

**Security exits on cluster channels**

Extra considerations when using security exits on cluster channels.

**About this task**

When a cluster-sender channel is first started, it uses attributes defined manually by a system administrator. When the channel is stopped and restarted, it picks up the attributes from the corresponding cluster-receiver channel definition. The original cluster-sender channel definition is overwritten with the new attributes, including the SecurityExit attribute.

**Procedure**

1. You must define a security exit on both the cluster-sender end and the cluster-receiver end of a channel.

The initial connection must be made with a security-exit handshake, even though the security exit name is sent over from the cluster-receiver definition.

2. Validate the PartnerName in the MQCXP structure in the security exit.  
The exit must allow the channel to start only if the partner queue manager is authorized
3. Design the security exit on the cluster-receiver definition to be receiver initiated.
4. If you design it as sender initiated, an unauthorized queue manager without a security exit can join the cluster because no security checks are performed.

Not until the channel is stopped and restarted can the SCYEXIT name be sent over from the cluster-receiver definition and full security checks made.

5. To view the cluster-sender channel definition that is currently in use, use the command:

```
DISPLAY CLUSQMGR(queue manager) ALL
```

The command displays the attributes that have been sent across from the cluster-receiver definition.

6. To view the original definition, use the command:

```
DISPLAY CHANNEL(channel name) ALL
```

7. You might need to define a channel auto-definition exit, CHADEXIT, on the cluster-sender queue manager, if the queue managers are on different platforms.

Use the channel auto-definition exit to set the SecurityExit attribute to an appropriate format for the target platform.

8. Deploy and configure the security-exit.

**z/OS** The security-exit load module must be in the data set specified in the CSQXLIB DD statement of the channel-initiator address-space procedure.

**Windows, UNIX and Linux systems**

- The security-exit dynamic link library must be in the path specified in the SCYEXIT attribute of the channel definition.

- The channel auto-definition exit dynamic link library must be in the path specified in the CHADEXIT attribute of the queue manager definition.

## Forcing unwanted queue managers to leave a cluster

Force an unwanted queue manager to leave a cluster by issuing the RESET CLUSTER command at a full repository queue manager.

### About this task

You can force an unwanted queue manager to leave a cluster. If for example, a queue manager is deleted but its cluster-receiver channels are still defined to the cluster. You might want to tidy up.

Only full repository queue managers are authorized to eject a queue manager from a cluster.

Follow this procedure to eject the queue manager OSLO from the cluster NORWAY:

### Procedure

1. On a full repository queue manager, issue the command:  
RESET CLUSTER(NORWAY) QMNAME(OSLO) ACTION(FORCEREMOVE)
2. Alternative use the QMID instead of QMNAME in the command:  
RESET CLUSTER(NORWAY) QMID(qmid) ACTION(FORCEREMOVE)

### Results

The queue manager that is force removed does not change: its local cluster definitions show it to be in the cluster. The definitions at all other queue managers do not show it in the cluster.

#### Related concepts:

“Security for queue manager clusters” on page 459

## Preventing queue managers receiving messages

You can prevent a cluster queue manager from receiving messages it is unauthorized to receive by using exit programs.

### About this task

It is difficult to stop a queue manager that is a member of a cluster from defining a queue. There is a danger that a rogue queue manager joins a cluster, and defines its own instance of one of the queues in the cluster. It can now receive messages that it is not authorized to receive. To prevent a queue manager receiving messages, use one of the following options given in the procedure.

### Procedure

- A channel exit program on each cluster-sender channel. The exit program uses the connection name to determine the suitability of the destination queue manager to be sent the messages.
- A cluster workload exit program, which uses the destination records to determine the suitability of the destination queue and queue manager to be sent the messages.

## Related concepts:

“Security for queue manager clusters” on page 459

## SSL and clusters

When configuring SSL for clusters, be aware a CLUSRCVR channel definition is propagated to other queue managers as an auto-defined CLUSSDR channel. If a CLUSRCVR channel uses SSL, you must configure SSL on all queue managers that communicate using the channel.

For more information about SSL, see WebSphere MQ support for SSL and TLS. The advice there is generally applicable to cluster channels, but you might want to give some special consideration to the following:

In a IBM WebSphere MQ cluster a particular CLUSRCVR channel definition is frequently propagated to many other queue managers where it is transformed into an auto-defined CLUSSDR. Subsequently the auto-defined CLUSSDR is used to start a channel to the CLUSRCVR. If the CLUSRCVR is configured for SSL connectivity the following considerations apply:

- All queue managers that want to communicate with this CLUSRCVR must have access to SSL support. This SSL provision must support the CipherSpec for the channel.
- The different queue managers to which the auto-defined cluster-sender channels have been propagated will each have a different distinguished name associated. If distinguished name peer checking is to be used on the CLUSRCVR it must be set up so all of the distinguished names that can be received are successfully matched.

For example, let us assume that all of the queue managers that will host cluster-sender channels which will connect to a particular CLUSRCVR, have certificates associated. Let us also assume that the distinguished names in all of these certificates define the country as UK, organization as IBM, the organization unit as IBM WebSphere MQ Development, and all have common names in the form DEVT.QMnnn, where nnn is numeric.

In this case an SSLPEER value of C=UK, O=IBM, OU=WebSphere MQ Development, CN=DEVT.QM\* on the CLUSRCVR will allow all the required cluster-sender channels to connect successfully, but will prevent unwanted cluster-sender channels from connecting.

- If custom CipherSpec strings are used, be aware that the custom string formats are not allowed on all platforms. An example of this is that the CipherSpec string RC4\_SHA\_US has a value of 05 on IBM i but is not a valid specification on UNIX, Linux or Windows systems. So if custom SSLCIPH parameters are used on a CLUSRCVR, all resulting auto-defined cluster-sender channels should reside on platforms on which the underlying SSL support implements this CipherSpec and on which it can be specified with the custom value. If you cannot select a value for the SSLCIPH parameter that will be understood throughout your cluster you will need a channel auto definition exit to change it into something the platforms being used will understand. Use the textual CipherSpec strings where possible (for example RC4\_MD5\_US).

An SSLCRLNL parameter applies to an individual queue manager and is not propagated to other queue managers within a cluster.

## Upgrading clustered queue managers and channels to SSL

Upgrade the cluster channels one at a time, changing all the CLUSRCVR channels before the CLUSSDR channels.

### Before you begin

Consider the following considerations, as these might affect your choice of CipherSpec for a cluster:

- Some CipherSpecs are not available on all platforms. Take care to choose a CipherSpec that is supported by all of the queue managers in the cluster.

- Some CipherSpecs might be new in the current WebSphere MQ release and not supported in older releases. A cluster containing queue managers running at different MQ releases is only be able to use the CipherSpecs supported by each release.

To use a new CipherSpec within a cluster, you must first migrate all of the cluster queue managers to the current release.

- Some CipherSpecs require a specific type of digital certificate to be used, notably those that use Elliptic Curve Cryptography.

Upgrade all queue managers in the cluster to WebSphere MQ V6 or higher, if they are not already at these levels. Distribute the certificates and keys so that SSL works from each of them.

## About this task

Change one CLUSRCVR at a time, and allow the changes to flow through the cluster before changing the next. Make sure that you do not change the reverse path until the changes for the current channel have been distributed throughout the cluster.

## Procedure

1. Switch the CLUSRCVR channels to SSL in any order you like. The changes flow in the opposite direction over channels which are not changed to SSL.
2. Switch all manual CLUSSDR channels to SSL. This does not have any effect on the operation of the cluster, unless you use the REFRESH CLUSTER command with the REPOS(YES) option.

**Note:** For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send

status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

## Related concepts:

“Specifying CipherSpecs” on page 774



Clustering: Using REFRESH CLUSTER best practices

“Digital certificates and CipherSpec compatibility in IBM WebSphere MQ” on page 403

## Disabling SSL on clustered queue managers and channels

To turn off SSL, set the SSLCIPH parameter to ' '. Disable SSL on the cluster channels individually, changing all the cluster receiver channels before the cluster sender channels.

## About this task

Change one cluster receiver channel at a time, and allow the changes to flow through the cluster before changing the next.

**Important:** Ensure that you do not change the reverse path until the changes for the current channel have been distributed throughout the cluster.

## Procedure

1. Set the value of the SSLCIPH parameter to ' ', an empty string in a single quotation mark , or \*NONE on IBM i. . You can turn off SSL on the cluster receiver channels in any order you like.  
Note that the changes flow in the opposite direction over channels on which you leave SSL active.
2. Check that the new value is reflected in all the other queue managers by using the command **DISPLAY CLUSQMGR(\*) ALL**.
3. Turn off SSL on all manual cluster sender channels. This does not have any effect on the operation of the cluster, unless you use the **REFRESH CLUSTER** command with the REPOS(YES) option.

For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at regular intervals thereafter, when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster for more information.

4. Stop and restart the cluster sender channels.

## Publish/subscribe security

The components and interactions that are involved in publish/subscribe are described as an introduction to the more detailed explanations and examples that follow.

There are a number of components involved in publishing and subscribing to a topic. Some of the security relationships between them are illustrated in Figure 101 and described in the following example.

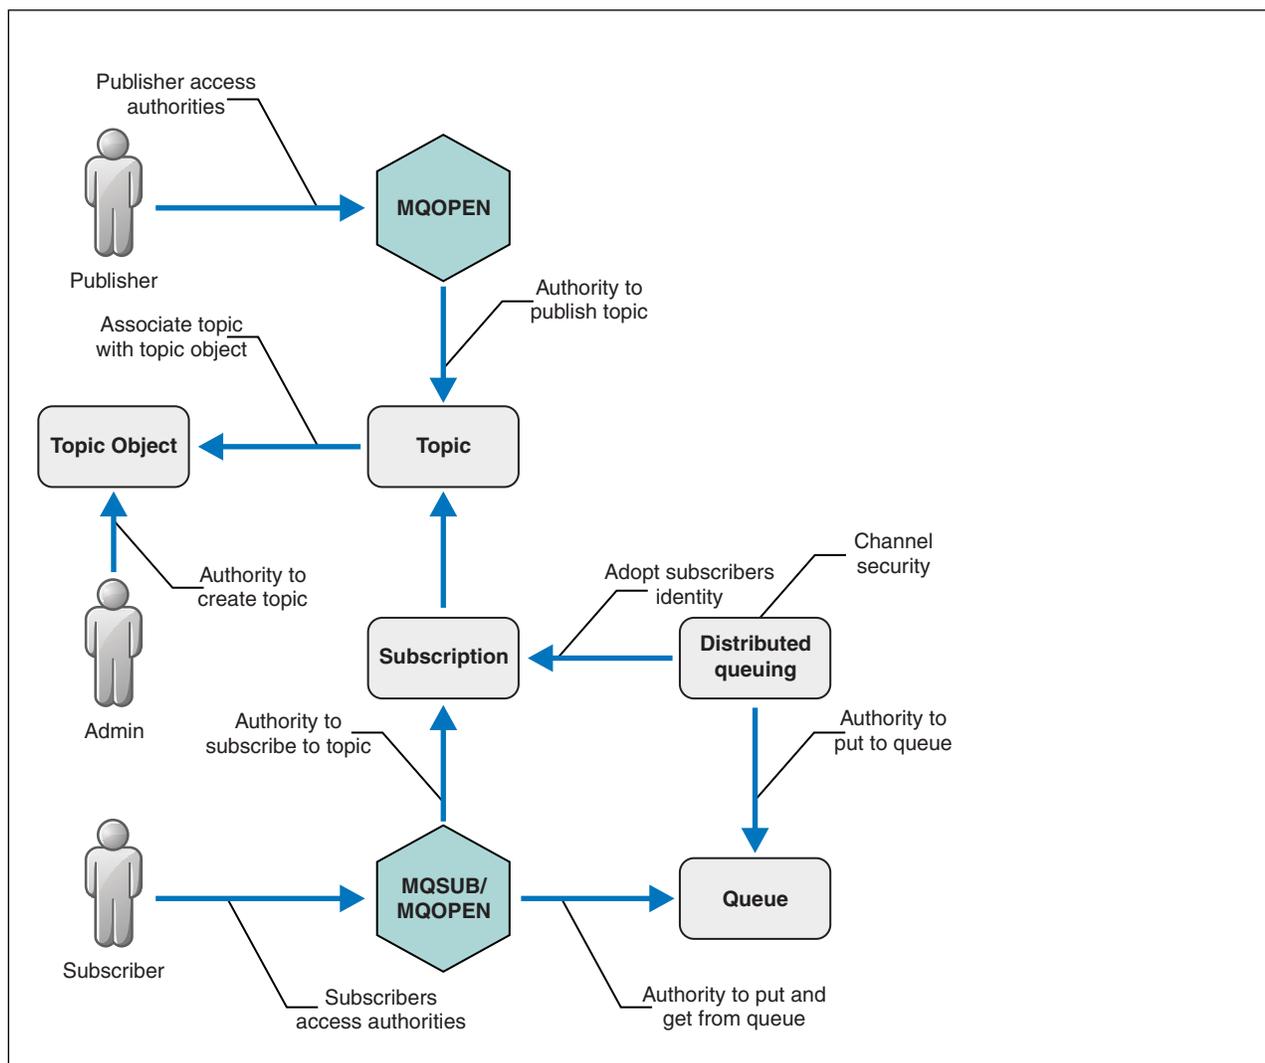


Figure 101. Publish/subscribe security relationships

### Topics (WebSphere MQ V7.1 Installing Guide)

Topics are identified by topic strings, and are typically organized into trees, see  Topic trees

(*WebSphere MQ V7.1 Installing Guide*). You need to associate a topic with a topic object to control access to the topic. “Topic security model” on page 810 explains how you secure topics using topic objects.

### **Administrative topic objects** (*WebSphere MQ V7.1 Installing Guide*)

You can control who has access to a topic, and for what purpose, by using the command **setmqaut** with a list of administrative topic objects. See the examples, “Grant access to a user to subscribe to a topic” on page 815 and “Grant access to a user to publish to a topic” on page 821. For controlling access to topic objects on z/OS, see Profiles for topic security.

### **Subscriptions** (*WebSphere MQ V7.1 Installing Guide*)

Subscribe to one or more topics by creating a subscription supplying a topic string, which can include wildcards, to match against the topic strings of publications. For further details, see:

#### **Subscribe using a topic object**

“Subscribing using the topic object name” on page 812

#### **Subscribe using a topic**

“Subscribing using a topic string where the topic node does not exist” on page 812

#### **Subscribe using a topic with wildcards**

“Subscribing using a topic string that contains wildcard characters” on page 813

A subscription contains information about the identity of the subscriber and the identity of the destination queue on to which the publications are to be placed. It also contains information about how the publication is to be placed on the destination queue.

As well as defining which subscribers have the authority to subscribe to certain topics, you can restrict subscriptions to being used by an individual subscriber. You can also control what information about the subscriber is used by the queue manager when publications are placed on to the destination queue. See “Subscription security” on page 826.

### **Queues** (*WebSphere MQ V7.1 Product Overview Guide*)

The destination queue is an important queue to secure. It is local to the subscriber, and publications that matched the subscription are placed onto it. You need to consider access to the destination queue from two perspectives:

1. Putting a publication on to the destination queue.
2. Getting the publication off the destination queue.

The queue manager puts a publication onto the destination queue using an identity provided by the subscriber. The subscriber, or a program that has been delegated the task of getting publications, takes messages off the queue. See “Authority to destination queues” on page 813.

There are no topic object aliases, but you can use an alias queue as the alias for a topic object. If you do so, as well as checking authority to use the topic for publish or subscribe, the queue manager checks authority to use the queue.

### **Publish/subscribe security between queue managers** (*WebSphere MQ V7.1 Installing Guide*)

Your permission to publish or subscribe to a topic is checked on the local queue manager using local identities and authorizations. Authorization does not depend on whether the topic is defined or not, nor where it is defined. Consequently, you need to perform topic authorization on every queue manager in a cluster when clustered topics are used.

**Note:** The security model for topics differs from the security model for queues. You can achieve the same result for queues by defining a queue alias locally for every clustered queue.

Queue managers exchange subscriptions in a cluster. In most WebSphere MQ cluster configurations, channels are configured with `PUTAUT=DEF` to place messages onto target queues

using the authority of the channel process. You can modify the channel configuration to use `PUTAUT=CTX` to require the subscribing user to have authority to propagate a subscription onto another queue manager in a cluster.

 Publish/subscribe security between queue managers (*WebSphere MQ V7.1 Installing Guide*) describes how to change your channel definitions to control who is allowed to propagate subscriptions onto other servers in the cluster.

#### **Authorization** (*WebSphere MQ V7.1 Reference*)

You can apply authorization to topic objects, just like queues and other objects. There are three authorization operations, `pub`, `sub`, and `resume` that you can apply only to topics. The details are described in  *Specifying authorities for different object types*.

#### **Function calls** (*WebSphere MQ V7.1 Reference*)

In publish and subscribe programs, like in queued programs, authorization checks are made when objects are opened, created, changed, or deleted. Checks are not made when `MQPUT` or `MQGET` `MQI` calls are made to put and get publications.

To publish a topic, perform an `MQOPEN` on the topic, which performs the authorization checks. Publish messages to the topic handle using the `MQPUT` command, which performs no authorization checks.

To subscribe to a topic, typically you perform an `MQSUB` command to create or resume the subscription, and also to open the destination queue to receive publications. Alternatively, perform a separate `MQOPEN` to open the destination queue, and then perform the `MQSUB` to create or resume the subscription.

Whichever calls you use, the queue manager checks that you can subscribe to the topic and get the resulting publications from the destination queue. If the destination queue is unmanaged, authorization checks are also made that the queue manager is able to place publications on the destination queue. It uses the identity it adopted from a matching subscription. It is assumed that the queue manager is always able to place publications onto managed destination queues.

## Roles

Users are involved in four roles in running publish/subscribe applications:

1. Publisher
2. Subscriber
3. Topic administrator
4. WebSphere MQ Administrator - member of group `mqm`

Define groups with appropriate authorizations corresponding to the publish, subscribe, and topic administration roles. You can then assign principals to these groups authorizing them to perform specific publish and subscribe tasks.

In addition, you need to extend the administrative operations authorizations to the administrator of the queues and channels responsible for moving publications and subscriptions.

## Topic security model

Only defined topic objects can have associated security attributes. For a description of topic objects, see

 Administrative topic objects. The security attributes specify whether a specified user ID, or security group, is permitted to perform a subscribe or a publish operation on each topic object.

The security attributes are associated with the appropriate administration node in the topic tree. When an authority check is made for a particular user ID during a subscribe or publish operation, the authority granted is based on the security attributes of the associated topic tree node.

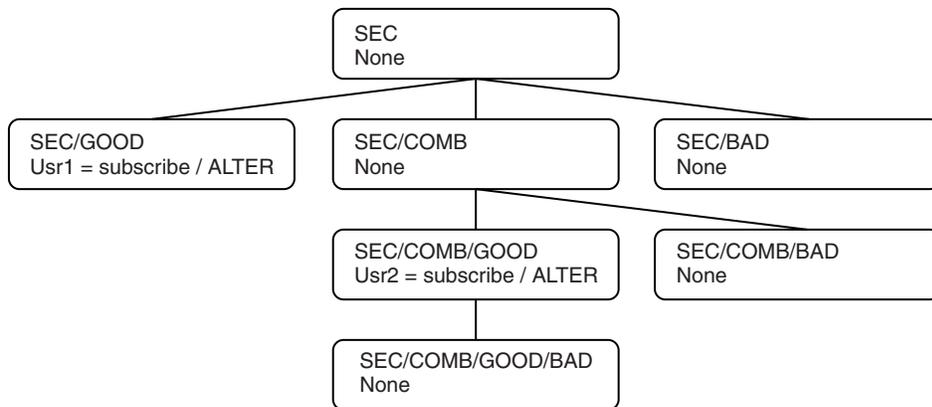
The security attributes are an access control list, indicating what authority a particular operating system user ID or security group has to the topic object.

Consider the following example where the topic objects have been defined with the security attributes, or authorities shown:

Table 84. Example topic object authorities

| Topic name | Topic string      | Authorities - not z/OS | z/OS authorities                |
|------------|-------------------|------------------------|---------------------------------|
| SECR00T    | SEC               | None                   | None                            |
| SECGOOD    | SEC/GOOD          | usr1+subscribe         | ALTER<br>HLQ.SUBSCRIBE.SECGOOD  |
| SECBAD     | SEC/BAD           | None                   | None<br>HLQ.SUBSCRIBE.SECBAD    |
| SECCOMB    | SEC/COMB          | None                   | None<br>HLQ.SUBSCRIBE.SECCOMB   |
| SECCOMBB   | SEC/COMB/GOOD/BAD | None                   | None<br>HLQ.SUBSCRIBE.SECCOMBB  |
| SECCOMBG   | SEC/COMB/GOOD     | usr2+subscribe         | ALTER<br>HLQ.SUBSCRIBE.SECCOMBG |
| SECCOMBN   | SEC/COMB/BAD      | None                   | None<br>HLQ.SUBSCRIBE.SECCOMBN  |

The topic tree with the associated security attributes at each node can be represented as follows:



The examples listed give the following authorizations:

- At the root node of the tree /SEC, no user has authority at that node.
- usr1 has been granted subscribe authority to the object /SEC/GOOD
- usr2 has been granted subscribe authority to the object /SEC/COMB/GOOD

## Subscribing using the topic object name

When subscribing to a topic object by specifying the MQCHAR48 name, the corresponding node in the topic tree is located. If the security attributes associated with the node indicate that the user has authority to subscribe, then access is granted.

If the user is not granted access, the parent node in the tree determines if the user has authority to subscribe at the parent node level. If so, then access is granted. If not, then the parent of that node is considered. The recursion continues until a node is located that grants subscribe authority to the user. The recursion stops when the root node is considered without authority having been granted. In the latter case, access is denied.

In short, if any node in the path grants authority to subscribe to that user or application, the subscriber is allowed to subscribe at that node, or anywhere below that node in the topic tree.

The root node in the example is SEC.

The user is granted subscribe authority if the access control list indicates that the user ID itself has authority, or that an operating system security group of which the user ID is a member has authority.

So, for example:

- If usr1 tries to subscribe, using a topic string of SEC/GOOD, the subscription would be allowed as the user ID has access to the node associated with that topic. However, if usr1 tried to subscribe using topic string SEC/COMB/GOOD the subscription would not be allowed as the user ID does not have access to the node associated with it.
- If usr2 tries to subscribe, using a topic string of SEC/COMB/GOOD the subscription would be allowed as the user ID has access to the node associated with the topic. However, if usr2 tried to subscribe to SEC/GOOD the subscription would not be allowed as the user ID does not have access to the node associated with it.
- If usr2 tries to subscribe using a topic string of SEC/COMB/GOOD/BAD the subscription would be allowed to because the user ID has access to the parent node SEC/COMB/GOOD.
- If usr1 or usr2 tries to subscribe using a topic string of /SEC/COMB/BAD, neither would be allowed as they do not have access to the topic node associated with it, or the parent nodes of that topic.

A subscribe operation specifying the name of a topic object that does not exist results in an MQRC\_UNKNOWN\_OBJECT\_NAME error.

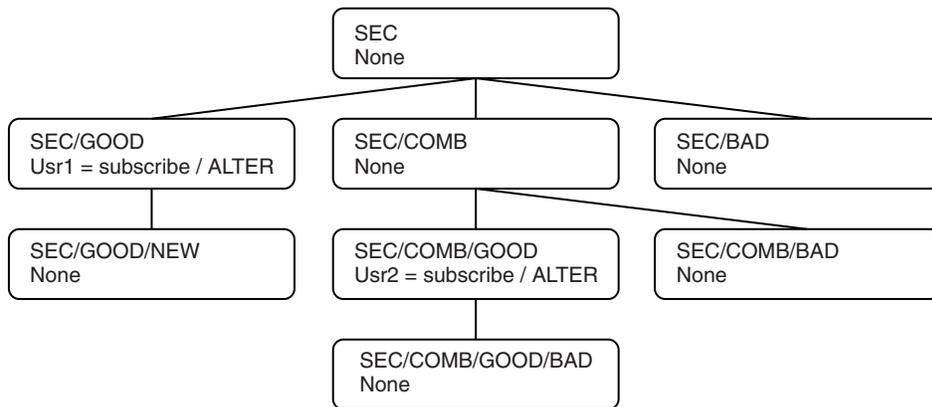
## Subscribing using a topic string where the topic node exists

The behavior is the same as when specifying the topic by the MQCHAR48 object name.

## Subscribing using a topic string where the topic node does not exist

Consider the case of an application subscribing, specifying a topic string representing a topic node that does not currently exist in the topic tree. The authority check is performed as outlined in the previous section. The check starts with the parent node of that which is represented by the topic string. If the authority is granted, a new node representing the topic string is created in the topic tree.

For example, usr1 tries to subscribe to a topic SEC/GOOD/NEW. Authority is granted as usr1 has access to the parent node SEC/GOOD. A new topic node is created in the tree as the following diagram shows. The new topic node is not a topic object it does not have any security attributes associated with it directly; the attributes are inherited from its parent.



## Subscribing using a topic string that contains wildcard characters

Consider the case of subscribing using a topic string that contains a wildcard character. The authority check is made against the node in the topic tree that matches the fully qualified part of the topic string.

So, if an application subscribes to `SEC/COMB/GOOD/*`, an authority check is carried out as outlined in the previous two sections on the node `SEC/COMB/GOOD` in the topic tree.

Similarly, if an application needs to subscribe to `SEC/COMB/*GOOD`, an authority check is carried out on the node `SEC/COMB`.

## Authority to destination queues

When subscribing to a topic, one of the parameters is the handle `hobj` of a queue that has been opened for output to receive the publications.

If `hobj` is not specified, but is blank, a managed queue is created if the following conditions apply:

- The `MQSO_MANAGED` option has been specified.
- The subscription does not exist.
- `Create` is specified.

If `hobj` is blank, and you are altering or resuming an existing subscription, the previously provided destination queue could be either managed or unmanaged.

The application or user making the `MQSUB` request must have the authority to put messages to the destination queue it has provided; in effect authority to have published messages put on that queue. The authority check follows the existing rules for queue security checking.

The security checking includes alternate user ID and context security checks where required. To be able to set any of the Identity context fields you must specify the `MQSO_SET_IDENTITY_CONTEXT` option as well as the `MQSO_CREATE` or `MQSO_ALTER` option. You cannot set any of the Identity context fields on an `MQSO_RESUME` request.

If the destination is a managed queue, no security checks are performed against the managed destination. If you are allowed to subscribe to a topic it is assumed that you can use managed destinations.

## **Publishing using the topic name or topic string where the topic node exists**

The security model for publishing is the same as that for subscribing, with the exception of wildcards. Publications do not contain wildcards; so there is no case of a topic string containing wildcards to consider.

The authorities to publish and subscribe are distinct. A user or group can have the authority to do one without necessarily being able to do the other.

When publishing to a topic object by specifying either the MQCHAR48 name or the topic string, the corresponding node in the topic tree is located. If the security attributes associated with the topic node indicates that the user has authority to publish, then access is granted.

If access is not granted, the parent node in the tree determines if the user has authority to publish at that level. If so, then access is granted. If not, the recursion continues until a node is located which grants publish authority to the user. The recursion stops when the root node is considered without authority having been granted. In the latter case, access is denied.

In short, if any node in the path grants authority to publish to that user or application, the publisher is allowed to publish at that node or anywhere below that node in the topic tree.

## **Publishing using the topic name or topic string where the topic node does not exist**

As with the subscribe operation, when an application publishes, specifying a topic string representing a topic node that does not currently exist in the topic tree, the authority check is performed starting with the parent of the node represented by the topic string. If the authority is granted, a new node representing the topic string is created in the topic tree.

## **Publishing using an alias queue that resolves to a topic object**

If you publish using an alias queue that resolves to a topic object then security checking occurs on both the alias queue and the underlying topic to which it resolves.

The security check on the alias queue verifies that the user has authority to put messages on that alias queue and the security check on the topic verifies that the user can publish to that topic. When an alias queue resolves to another queue, checks are *not* made on the underlying queue. Authority checking is performed differently for topics and queues.

## **Closing a subscription**

There is additional security checking if you close a subscription using the MQCO\_REMOVE\_SUB option if you did not create the subscription under this handle.

A security check is performed to ensure that you have the correct authority to do this as the action results in the removal of the subscription. If the security attributes associated with the topic node indicate that the user has authority, then access is granted. If not, then the parent node in the tree is considered to determine if the user has authority to close the subscription. The recursion continues until either authority is granted or the root node is reached.

## **Defining, altering, and deleting a subscription**

No subscribe security checks are performed when a subscription is created administratively, rather than using an MQSUB API request. The administrator has already been given this authority through the command.

Security checks are performed to ensure that publications can be put on the destination queue associated with the subscription. The checks are performed in the same way as for an MQSUB request.

The user ID that is used for these security checks depends upon the command being issued. If the **SUBUSER** parameter is specified it affects the way the check is performed, as shown in Table 85:

Table 85. User IDs used for security checks for commands

| Command       | SUBUSER specified and blank | SUBUSER specified and completed             | SUBUSER not specified                          |
|---------------|-----------------------------|---------------------------------------------|------------------------------------------------|
| <b>DEFINE</b> | Use the administrator ID    | Use the user ID specified in <b>SUBUSER</b> | Use the administrator ID                       |
| <b>ALTER</b>  | Use the administrator ID    | Use the user ID specified in <b>SUBUSER</b> | Use the user ID from the existing subscription |

The only security check performed when deleting subscriptions using the DELETE SUB command is the command security check.

## Example publish/subscribe security setup

This section describes a scenario that has access control setup on topics in a way that allows the security control to be applied as required.

### Grant access to a user to subscribe to a topic

This topic is the first one in a list of tasks that tells you how to grant access to topics by more than one user.

#### About this task

This task assumes that no administrative topic objects exist, nor have any profiles been defined for subscription or publication. The applications are creating new subscriptions, rather than resuming existing ones, and are doing so using the topic string only.

An application can make a subscription by providing a topic object, or a topic string, or a combination of both. Whichever way the application selects, the effect is to make a subscription at a certain point in the topic tree. If this point in the topic tree is represented by an administrative topic object, a security profile is checked based on the name of that topic object.

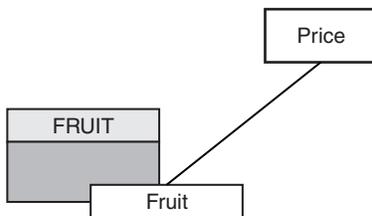


Figure 102. Topic object access example

Table 86. Example topic object access

| Topic       | Subscribe access required | Topic object |
|-------------|---------------------------|--------------|
| Price       | No user                   | None         |
| Price/Fruit | USER1                     | FRUIT        |

Define a new topic object as follows:

## Procedure

1. Issue the MQSC command `DEF TOPIC(FRUIT) TOPICSTR('Price/Fruit')`.

2. Grant access as follows:

- **z/OS:**

Grant access to USER1 to subscribe to topic "Price/Fruit" by granting the user access to the `hlq.SUBSCRIBE.FRUIT` profile. Do this, using the following RACF commands:

```
RDEFINE MXTOPIC hlq.SUBSCRIBE.FRUIT UACC(NONE)
PERMIT hlq.SUBSCRIBE.FRUIT CLASS(MXTOPIC) ID(USER1) ACCESS(ALTER)
```

- **Other platforms:**

Grant access to USER1 to subscribe to topic "Price/Fruit" by granting the user access to the FRUIT object. Do this, using the authorization command for the platform:

**Windows, UNIX and Linux systems**

```
setmqaut -t topic -n FRUIT -p USER1 +sub
```

**IBM i**

```
GRTMQAUT OBJ(FRUIT) OBJTYPE(*TOPIC) USER(USER1) AUT(*SUB)
```

## Results

When USER1 attempts to subscribe to topic "Price/Fruit" the result is success.

When USER2 attempts to subscribe to topic "Price/Fruit" the result is failure with an `MQRC_NOT_AUTHORIZED` message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```
ICH408I USER(USER2) ...
 hlq.SUBSCRIBE.FRUIT ...
```

```
ICH408I USER(USER2) ...
 hlq.SUBSCRIBE.SYSTEM.BASE.TOPIC ...
```

- On other platforms, the following authorization event:

```
MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_SUB_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames FRUIT, SYSTEM.BASE.TOPIC
TopicString "Price/Fruit"
```

Note that this is an illustration of what you see; not all the fields.

## Grant access to a user to subscribe to a topic deeper within the tree

This topic is the second in a list of tasks that tells you how to grant access to topics by more than one user.

### Before you begin

This topic uses the setup described in "Grant access to a user to subscribe to a topic" on page 815.

### About this task

If the point in the topic tree where the application makes the subscription is not represented by an administrative topic object, move up the tree until the closest parent administrative topic object is located. The security profile is checked, based on the name of that topic object.

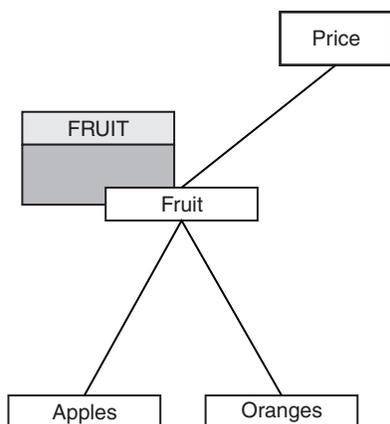


Figure 103. Example of granting access to a topic within a topic tree

Table 87. Access requirements for example topics and topic objects

| Topic               | Subscribe access required | Topic object |
|---------------------|---------------------------|--------------|
| Price               | No user                   | None         |
| Price/Fruit         | USER1                     | FRUIT        |
| Price/Fruit/Apples  | USER1                     |              |
| Price/Fruit/Oranges | USER1                     |              |

In the previous task USER1 was granted access to subscribe to topic "Price/Fruit" by granting it access to the hlq.SUBSCRIBE.FRUIT profile on z/OS and subscribe access to the FRUIT profile on other platforms. This single profile also grants USER1 access to subscribe to "Price/Fruit/Apples", "Price/Fruit/Oranges" and "Price/Fruit/#".

When USER1 attempts to subscribe to topic "Price/Fruit/Apples" the result is success.

When USER2 attempts to subscribe to topic "Price/Fruit/Apples" the result is failure with an MQRC\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```

ICH408I USER(USER2) ...
 hlq.SUBSCRIBE.FRUIT ...

```

```

ICH408I USER(USER2) ...
 hlq.SUBSCRIBE.SYSTEM.BASE.TOPIC ...

```

- On other platforms, the following authorization event:

```

MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_SUB_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames FRUIT, SYSTEM.BASE.TOPIC
TopicString "Price/Fruit/Apples"

```

Note the following:

- The messages you receive on z/OS are identical to those received in the previous task as the same topic objects and profiles are controlling the access.

- The event message you receive on other platforms is similar to the one received in the previous task, but the actual topic string is different.

## Grant another user access to subscribe to only the topic deeper within the tree

This topic is the third in a list of tasks that tells you how to grant access to subscribe to topics by more than one user.

### Before you begin

This topic uses the setup described in “Grant access to a user to subscribe to a topic deeper within the tree” on page 816.

### About this task

In the previous task USER2 was refused access to topic "Price/Fruit/Apples". This topic tells you how to grant access to that topic, but not to any other topics.

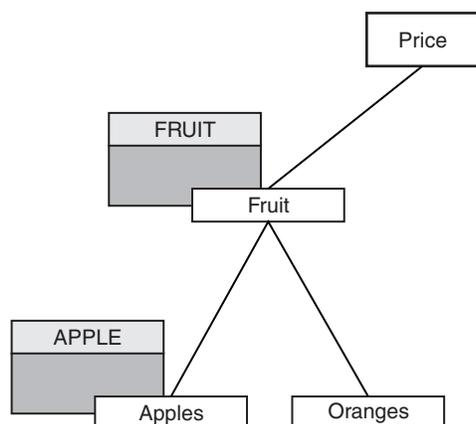


Figure 104. Granting access to specific topics within a topic tree

Table 88. Access requirements for example topics and topic objects

| Topic               | Subscribe access required | Topic object |
|---------------------|---------------------------|--------------|
| Price               | No user                   | None         |
| Price/Fruit         | USER1                     | FRUIT        |
| Price/Fruit/Apples  | USER1 and USER2           | APPLE        |
| Price/Fruit/Oranges | USER1                     |              |

Define a new topic object as follows:

### Procedure

1. Issue the MQSC command `DEF TOPIC(APPLE) TOPICSTR('Price/Fruit/Apples')`.
2. Grant access as follows:
  - **z/OS:**

In the previous task USER1 was granted access to subscribe to topic "Price/Fruit/Apples" by granting the user access to the `hlq.SUBSCRIBE.FRUIT` profile.

This single profile also granted USER1 access to subscribe to "Price/Fruit/Oranges" "Price/Fruit/#" and this access remains even with the addition of the new topic object and the profiles associated with it.

Grant access to USER2 to subscribe to topic "Price/Fruit/Apples" by granting the user access to the hlq.SUBSCRIBE.APPLE profile. Do this, using the following RACF commands:

```
RDEFINE MXTOPIC hlq.SUBSCRIBE.APPLE UACC(NONE)
PERMIT hlq.SUBSCRIBE.FRUIT APPLE(MXTOPIC) ID(USER2) ACCESS(ALTER)
```

- Other platforms:

In the previous task USER1 was granted access to subscribe to topic "Price/Fruit/Apples" by granting the user subscribe access to the FRUIT profile.

This single profile also granted USER1 access to subscribe to "Price/Fruit/Oranges" and "Price/Fruit/#", and this access remains even with the addition of the new topic object and the profiles associated with it.

Grant access to USER2 to subscribe to topic "Price/Fruit/Apples" by granting the user subscribe access to the APPLE profile. Do this, using the authorization command for the platform:

#### Windows, UNIX and Linux systems

```
setmqaut -t topic -n APPLE -p USER2 +sub
```

#### IBM i

```
GRTMQAUT OBJ(APPLE) OBJTYPE(*TOPIC) USER(USER2) AUT(*SUB)
```

## Results

On z/OS, when USER1 attempts to subscribe to topic "Price/Fruit/Apples" the first security check on the hlq.SUBSCRIBE.APPLE profile fails, but on moving up the tree the hlq.SUBSCRIBE.FRUIT profile allows USER1 to subscribe, so the subscription succeeds and no return code is sent to the MQSUB call. However, a RACF ICH message is generated for the first check:

```
ICH408I USER(USER1) ...
hlq.SUBSCRIBE.APPLE ...
```

When USER2 attempts to subscribe to topic "Price/Fruit/Apples" the result is success because the security check passes on the first profile.

When USER2 attempts to subscribe to topic "Price/Fruit/Oranges" the result is failure with an MQRC\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```
ICH408I USER(USER2) ...
hlq.SUBSCRIBE.FRUIT ...
```

```
ICH408I USER(USER2) ...
hlq.SUBSCRIBE.SYSTEM.BASE.TOPIC ...
```

- On other platforms, the following authorization event:

```
MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_SUB_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames FRUIT, SYSTEM.BASE.TOPIC
TopicString "Price/Fruit/Oranges"
```

The disadvantage of this setup is that, on z/OS, you receive additional ICH messages on the console. You can avoid this if you secure the topic tree in a different manner.

## Change access control to avoid additional messages

This topic is the fourth in a list of tasks that tells you how to grant access to subscribe to topics by more than one user and to avoid additional RACF ICH408I messages on z/OS.

### Before you begin

This topic enhances the setup described in “Grant another user access to subscribe to only the topic deeper within the tree” on page 818 so that you avoid additional error messages.

### About this task

This topic tells you how to grant access to topics deeper in the tree, and how to remove access to the topic lower down the tree when no user requires it.

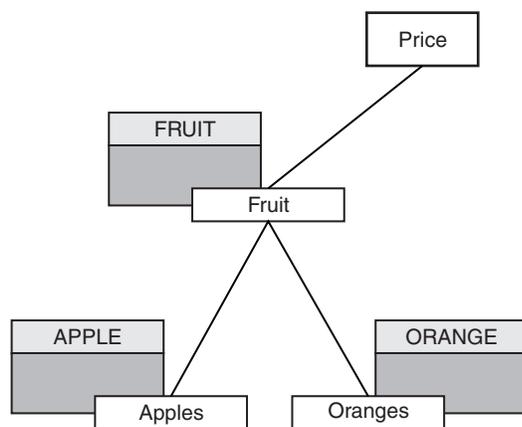


Figure 105. Example of granting access control to avoid additional messages.

Define a new topic object as follows:

### Procedure

1. Issue the MQSC command `DEF TOPIC(ORANGE) TOPICSTR('Price/Fruit/Oranges')`.
2. Grant access as follows:

- **z/OS:**

Define a new profile and add access to that profile, and the existing profiles. Do this, using the following RACF commands:

```
RDEFINE MXTOPIC h1q.SUBSCRIBE.ORANGE UACC(NONE)
PERMIT h1q.SUBSCRIBE.ORANGE CLASS(MXTOPIC) ID(USER1) ACCESS(ALTER)
PERMIT h1q.SUBSCRIBE.APPLE CLASS(MXTOPIC) ID(USER1) ACCESS(ALTER)
```

- **Other platforms:**

Setup the equivalent access by using the authorization commands for the platform:

#### **Windows, UNIX and Linux systems**

```
setmqaut -t topic -n ORANGE -p USER1 +sub
setmqaut -t topic -n APPLE -p USER1 +sub
```

#### **IBM i**

```
GRTMQAUT OBJ(ORANGE) OBJTYPE(*TOPIC) USER(USER1) AUT(*SUB)
GRTMQAUT OBJ(APPLE) OBJTYPE(*TOPIC) USER(USER1) AUT(*SUB)
```

## Results

On z/OS, when USER1 attempts to subscribe to topic "Price/Fruit/Apples" the first security check on the hlq.SUBSCRIBE.APPLE profile succeeds.

Similarly, when USER2 attempts to subscribe to topic "Price/Fruit/Apples" the result is success because the security check passes on the first profile.

When USER2 attempts to subscribe to topic "Price/Fruit/Oranges" the result is failure with an MQRC\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```
ICH408I USER(USER2) ...
hlq.SUBSCRIBE.ORANGE ...
```

```
ICH408I USER(USER2) ...
hlq.SUBSCRIBE.FRUIT ...
```

```
ICH408I USER(USER2) ...
hlq.SUBSCRIBE.SYSTEM.BASE.TOPIC ...
```

- On other platforms, the following authorization event:

```
MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_SUB_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames ORANGE, FRUIT, SYSTEM.BASE.TOPIC
TopicString "Price/Fruit/Oranges"
```

## Grant access to a user to publish to a topic

This topic is the first one in a list of tasks that tells you how to grant access to publish topics by more than one user.

### About this task

This task assumes that no administrative topic objects exist on the right hand side of the topic tree, nor have any profiles been defined for publication. The assumption used is that publishers are using the topic string only.

An application can publish to a topic by providing a topic object, or a topic string, or a combination of both. Whichever way the application selects, the effect is to publish at a certain point in the topic tree. If this point in the topic tree is represented by an administrative topic object, a security profile is checked based on the name of that topic object. For example:

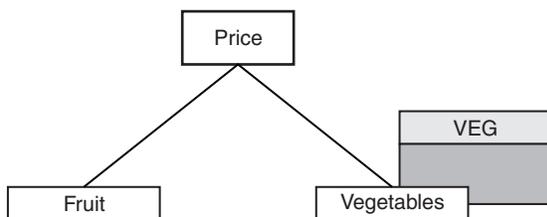


Figure 106. Granting publish access to a topic

Table 89. Example publish access requirements

| Topic            | Publish access required | Topic object |
|------------------|-------------------------|--------------|
| Price            | No user                 | None         |
| Price/Vegetables | USER1                   | VEG          |

Define a new topic object as follows:

### Procedure

1. Issue the MQSC command `DEF TOPIC(VEG) TOPICSTR('Price/Vegetables')`.
2. Grant access as follows:

- **z/OS:**

Grant access to USER1 to publish to topic "Price/Vegetables" by granting the user access to the h1q.PUBLISH.VEG profile. Do this, using the following RACF commands:

```
RDEFINE MXTOPIC h1q.PUBLISH.VEG UACC(NONE)
PERMIT h1q.PUBLISH.VEG CLASS(MXTOPIC) ID(USER1) ACCESS(UPDATE)
```

- **Other platforms:**

Grant access to USER1 to publish to topic "Price/Vegetables" by granting the user access to the VEG profile. Do this, using the authorization command for the platform:

**Windows, UNIX and Linux systems**

```
setmqaut -t topic -n VEG -p USER1 +pub
```

**IBM i**

```
GRTMQAUT OBJ(VEG) OBJTYPE(*TOPIC) USER(USER1) AUT(*PUB)
```

### Results

When USER1 attempts to publish to topic "Price/Vegetables" the result is success; that is, the MQOPEN call succeeds.

When USER2 attempts to publish to topic "Price/Vegetables" the MQOPEN call fails with an MQRC\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```
ICH408I USER(USER2) ...
h1q.PUBLISH.VEG ...
```

```
ICH408I USER(USER2) ...
h1q.PUBLISH.SYSTEM.BASE.TOPIC ...
```

- On other platforms, the following authorization event:

```
MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_OPEN_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames VEG, SYSTEM.BASE.TOPIC
TopicString "Price/Vegetables"
```

Note that this is an illustration of what you see; not all the fields.

## Grant access to a user to publish to a topic deeper within the tree

This topic is the second in a list of tasks that tells you how to grant access to publish to topics by more than one user.

### Before you begin

This topic uses the setup described in “Grant access to a user to publish to a topic” on page 821.

### About this task

If the point in the topic tree where the application publishes is not represented by an administrative topic object, move up the tree until the closest parent administrative topic object is located. The security profile is checked, based on the name of that topic object.

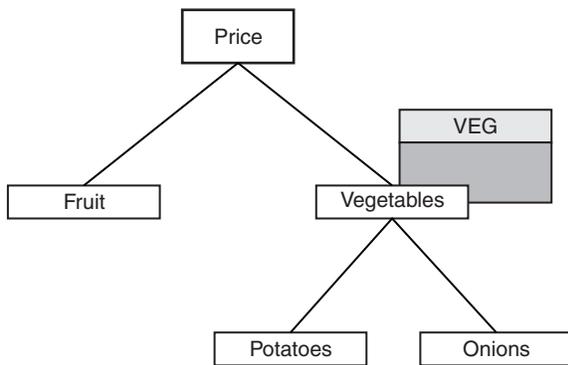


Figure 107. Granting publish access to a topic within a topic tree

Table 90. Example publish access requirements

| Topic                     | Subscribe access required | Topic object |
|---------------------------|---------------------------|--------------|
| Price                     | No user                   | None         |
| Price/Vegetables          | USER1                     | VEG          |
| Price/Vegetables/Potatoes | USER1                     |              |
| Price/Vegetables/Onions   | USER1                     |              |

In the previous task USER1 was granted access to publish topic "Price/Vegetables/Potatoes" by granting it access to the hlq.PUBLISH.VEG profile on z/OS or publish access to the VEG profile on other platforms. This single profile also grants USER1 access to publish at "Price/Vegetables/Onions".

When USER1 attempts to publish at topic "Price/Vegetables/Potatoes" the result is success; that is the MQOPEN call succeeds.

When USER2 attempts to subscribe to topic "Price/Vegetables/Potatoes" the result is failure; that is, the MQOPEN call fails with an MQRN\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```

ICH408I USER(USER2) ...
hlq.PUBLISH.VEG ...

```

```

ICH408I USER(USER2) ...
hlq.PUBLISH.SYSTEM.BASE.TOPIC ...

```

- On other platforms, the following authorization event:

```

MQRC_NOT_AUTHORIZED
ReasonQualifier MQRQ_OPEN_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames VEG, SYSTEM.BASE.TOPIC
TopicString "Price/Vegetables/Potatoes"

```

Note the following:

- The messages you receive on z/OS are identical to those received in the previous task as the same topic objects and profiles are controlling the access.
- The event message you receive on other platforms is similar to the one received in the previous task, but the actual topic string is different.

## Grant access for publish and subscribe

This topic is the last in a list of tasks that tells you how to grant access to publish and subscribe to topics by more than one user.

### Before you begin

This topic uses the setup described in “Grant access to a user to publish to a topic deeper within the tree” on page 823.

### About this task

In a previous task USER1 was given access to subscribe to the topic "Price/Fruit". This topic tells you how to grant access to that user to publish to that topic.

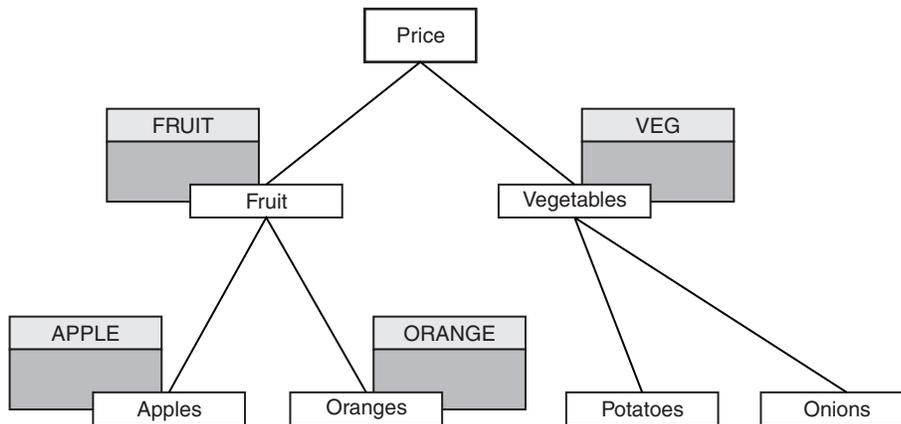


Figure 108. Granting access for publishing and subscribing

Table 91. Example publishing and subscribing access requirements

| Topic               | Subscribe access required | Publish access required | Topic object |
|---------------------|---------------------------|-------------------------|--------------|
| Price               | No user                   | No user                 | None         |
| Price/Fruit         | USER1                     | USER1                   | FRUIT        |
| Price/Fruit/Apples  | USER1 and USER2           |                         | APPLE        |
| Price/Fruit/Oranges | USER1                     |                         | ORANGE       |

## Procedure

Grant access as follows:

- **z/OS:**

In an earlier task USER1 was granted access to subscribe to topic "Price/Fruit" by granting the user access to the hlq.SUBSCRIBE.FRUIT profile.

In order to publish to the "Price/Fruit" topic, grant access to USER1 to the hlq.PUBLISH.FRUIT profile. Do this, using the following RACF commands:

```
RDEFINE MXTOPIC hlq.PUBLISH.FRUIT UACC(NONE)
PERMIT hlq.PUBLISH.FRUIT CLASS(MXTOPIC) ID(USER1) ACCESS(ALTER)
```

- **Other platforms:**

Grant access to USER1 to publish to topic "Price/Fruit" by granting the user publish access to the FRUIT profile. Do this, using the authorization command for the platform:

**Windows, UNIX and Linux systems**

```
setmqaut -t topic -n FRUIT -p USER1 +pub
```

**IBM i**

```
GRTMQAUT OBJ(FRUIT) OBJTYPE(*TOPIC) USER(USER1) AUT(*PUB)
```

## Results

On z/OS, when USER1 attempts to publish to topic "Price/Fruit" the security check on the MQOPEN call passes.

When USER2 attempts to publish at topic "Price/Fruit" the result is failure with an MQRQ\_NOT\_AUTHORIZED message, together with:

- On z/OS, the following messages seen on the console that show the full security path through the topic tree that has been attempted:

```
ICH408I USER(USER2) ...
hlq.PUBLISH.FRUIT ...
```

```
ICH408I USER(USER2) ...
hlq.PUBLISH.SYSTEM.BASE.TOPIC ...
```

- On other platforms, the following authorization event:

```
MQRQ_NOT_AUTHORIZED
ReasonQualifier MQRQ_OPEN_NOT_AUTHORIZED
UserIdentifier USER2
AdminTopicNames FRUIT, SYSTEM.BASE.TOPIC
TopicString "Price/Fruit"
```

Following the complete set of these tasks, gives USER1 and USER2 the following access authorities for publish and subscribe to the topics listed:

*Table 92. Complete list of access authorities resulting from security examples*

| Topic                     | Subscribe access required | Publish access required | Topic object |
|---------------------------|---------------------------|-------------------------|--------------|
| Price                     | No user                   | No user                 | None         |
| Price/Fruit               | USER1                     | USER1                   | FRUIT        |
| Price/Fruit/Apples        | USER1 and USER2           |                         | APPLE        |
| Price/Fruit/Oranges       | USER1                     |                         | ORANGE       |
| Price/Vegetables          |                           | USER1                   | VEG          |
| Price/Vegetables/Potatoes |                           |                         |              |

Table 92. Complete list of access authorities resulting from security examples (continued)

| Topic                   | Subscribe access required | Publish access required | Topic object |
|-------------------------|---------------------------|-------------------------|--------------|
| Price/Vegetables/Onions |                           |                         |              |

Where you have different requirements for security access at different levels within the topic tree, careful planning ensures that you do not receive extraneous security warnings on the z/OS console log. Setting up security at the correct level within the tree avoids misleading security messages.

## Subscription security

### MQSO\_ALTERNATE\_USER\_AUTHORITY

The AlternateUserId field contains a user identifier to use to validate this MQSUB call. The call can succeed only if this AlternateUserId is authorized to subscribe to the topic with the specified access options, regardless of whether the user identifier under which the application is running is authorized to do so.

### MQSO\_SET\_IDENTITY\_CONTEXT

The subscription is to use the accounting token and application identity data supplied in the PubAccountingToken and PubApplIdentityData fields.

If this option is specified, the same authorization check is carried out as if the destination queue was accessed using an MQOPEN call with MQOO\_SET\_IDENTITY\_CONTEXT, except in the case where the MQSO\_MANAGED option is also used in which case there is no authorization check on the destination queue.

If this option is not specified, the publications sent to this subscriber have default context information associated with them as follows:

Table 93. Default publication context information

| Field in MQMD    | Value used                                                                                                                |
|------------------|---------------------------------------------------------------------------------------------------------------------------|
| UserIdentifier   | The user ID associated with the subscription (see SUBUSER field on DISPLAY SBSTATUS) at the time the publication is made. |
| AccountingToken  | Determined from the environment if possible; set to MQACT_NONE otherwise.                                                 |
| ApplIdentityData | Set to blanks.                                                                                                            |

This option is only valid with MQSO\_CREATE and MQSO\_ALTER. If used with MQSO\_RESUME, the PubAccountingToken and PubApplIdentityData fields are ignored, so this option has no effect.

If a subscription is altered without using this option where previously the subscription had supplied identity context information, default context information is generated for the altered subscription.

If a subscription allowing different user IDs to use it with option MQSO\_ANY\_USERID, is resumed by a different user ID, default identity context is generated for the new user ID now owning the subscription and any subsequent publications are delivered containing the new identity context.

### AlternateSecurityId

This is a security identifier that is passed with the AlternateUserId to the authorization service to allow appropriate authorization checks to be performed. AlternateSecurityId is used only if MQSO\_ALTERNATE\_USER\_AUTHORITY is specified, and the AlternateUserId field is not entirely blank

up to the first null character or the end of the field.

### **MQSO\_ANY\_USERID subscription option**

When MQSO\_ANY\_USERID is specified, the identity of the subscriber is not restricted to a single user ID. This allows any user to alter or resume the subscription when they have suitable authority. Only a single user may have the subscription at any one time. An attempt to resume use of a subscription currently in use by another application will cause the call to fail with MQRC\_SUBSCRIPTION\_IN\_USE.

To add this option to an existing subscription the MQSUB call (using MQSO\_ALTER) must come from the same user ID as the original subscription.

If an MQSUB call refers to an existing subscription with MQSO\_ANY\_USERID set, and the user ID differs from the original subscription, the call succeeds only if the new user ID has authority to subscribe to the topic. After successful completion, future publications to this subscriber are put to the subscriber's queue with the new user ID set in the publication.

### **MQSO\_FIXED\_USERID**

When MQSO\_FIXED\_USERID is specified, the subscription can only be altered or resumed by a single owning user ID. This user ID is the last user ID to alter the subscription that set this option, thereby removing the MQSO\_ANY\_USERID option, or if no alters have taken place, it is the user ID that created the subscription.

If an MQSUB verb refers to an existing subscription with MQSO\_ANY\_USERID set and alters the subscription (using MQSO\_ALTER) to use option MQSO\_FIXED\_USERID, the user ID of the subscription is now fixed at this new user ID. The call succeeds only if the new user ID has authority to subscribe to the topic.

If a user ID other than the one recorded as owning a subscription tries to resume or alter an MQSO\_FIXED\_USERID subscription, the call will fail with MQRC\_IDENTITY\_MISMATCH. The owning user ID of a subscription can be viewed using the DISPLAY SBSTATUS command.

If neither MQSO\_ANY\_USERID or MQSO\_FIXED\_USERID is specified, the default is MQSO\_FIXED\_USERID.



---

## Monitoring and performance

A number of monitoring techniques are available in IBM WebSphere MQ to obtain statistics and other specific information about how your queue manager network is running. Use the monitoring information and guidance in this section to help improve the performance of your queue manager network.

Depending on the size and complexity of your queue manager network, you can obtain a range of information from monitoring your queue manager network. The following list provides examples of reasons for monitoring your queue manager network:

- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Improve the efficiency of your queue manager network.
- Familiarize yourself with the running of your queue manager network.
- Confirm that your queue manager network is running correctly.
- Generate messages when certain events occur.
- Record message activity.
- Determine the last known location of a message.
- Check various statistics of a queue manager network in real time.
- Generate an audit trail.
- Account for application resource usage.
- Capacity planning.

### Related concepts:



Configuring (*WebSphere MQ V7.1 Installing Guide*)

“Administering IBM WebSphere MQ” on page 1

“Event monitoring”

“Message monitoring” on page 885

“Accounting and statistics messages” on page 966

“Real-time monitoring” on page 1102

“Monitoring performance and resource usage” on page 1114



Configuring (*WebSphere MQ V7.1 Installing Guide*)

“Administering IBM WebSphere MQ” on page 1

“Event monitoring”

“Message monitoring” on page 885

“Accounting and statistics messages” on page 966

“Real-time monitoring” on page 1102

“Monitoring performance and resource usage” on page 1114

---

## Event monitoring

Event monitoring is the process of detecting occurrences of *instrumentation events* in a queue manager network. An instrumentation event is a logical combination of events that is detected by a queue manager or channel instance. Such an event causes the queue manager or channel instance to put a special message, called an *event message*, on an event queue.

IBM WebSphere MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. Use these events to monitor the operation of the queue managers in your queue manager network to achieve the following goals:

- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Generate an audit trail.
- React to queue manager state changes

**Related concepts:**

“Instrumentation events”

“Performance events” on page 845

“Configuration events” on page 863

“Command events” on page 867

“Logger events” on page 869

**Related reference:**



Event message reference (*WebSphere MQ V7.1 Reference*)

“Event types” on page 832



Event message reference (*WebSphere MQ V7.1 Reference*)



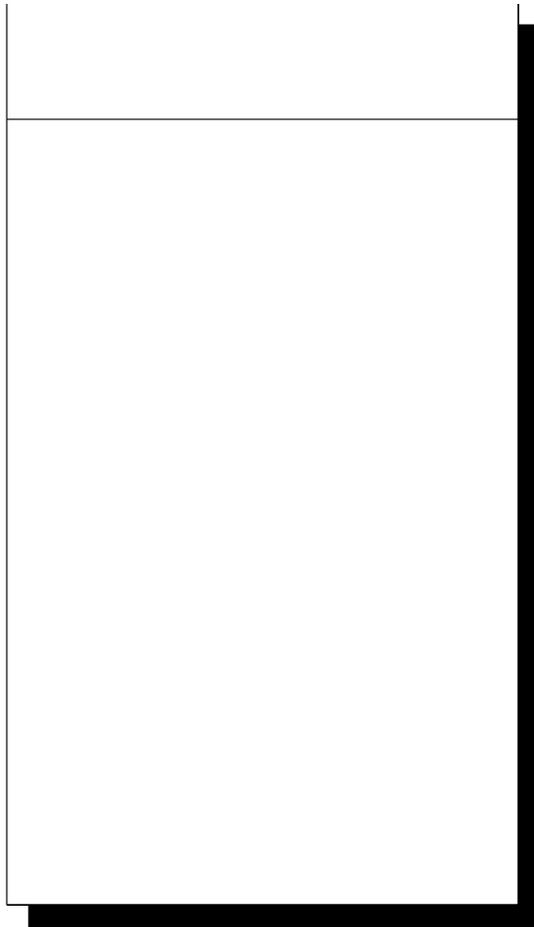
Event message format (*WebSphere MQ V7.1 Reference*)

## Instrumentation events

An instrumentation event is a logical combination of conditions that a queue manager or channel instance detects and puts a special message, called an *event message*, on an event queue.

IBM WebSphere MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. You can use these events to monitor the operation of queue managers (with other methods such as Tivoli NetView for z/OS).

Figure 109 on page 831 illustrates the concept of instrumentation events.



*Figure 109. Understanding instrumentation events*

## **Event monitoring applications**

Applications that use events to monitor queue managers must include the following provisions:

1. Set up channels between the queue managers in your network.
2. Implement the required data conversions. The normal rules of data conversion apply. For example, if you are monitoring events on a UNIX system queue manager from a z/OS queue manager, ensure that you convert EBCDIC to ASCII.

## **Event notification through event queues**

When an event occurs, the queue manager puts an event message on the appropriate event queue, if defined. The event message contains information about the event that you can retrieve by writing a suitable MQI application program that performs the following steps:

- Get the message from the queue.
- Process the message to extract the event data.

The related information describes the format of event messages.

## Conditions that cause events

The following list gives examples of conditions that can cause instrumentation events:

- A threshold limit for the number of messages on a queue is reached.
- A channel instance is started or stopped.
- A queue manager becomes active, or is requested to stop.
- An application tries to open a queue specifying a user ID that is not authorized on IBM WebSphere MQ for IBM i, Windows, UNIX and Linux systems.
- Objects are created, deleted, changed, or refreshed.
- An MQSC or PCF command runs successfully.
- A queue manager starts writing to a new log extent.
- Putting a message on the dead-letter queue, if the event conditions are met.

### Related concepts:

“Controlling events” on page 840

“Event queues” on page 843

“Performance events” on page 845

“Sample program to monitor instrumentation events” on page 877

“Controlling events” on page 840

“Event queues” on page 843

“Performance events” on page 845

“Sample program to monitor instrumentation events” on page 877

### Related reference:

“Event types”

“Event types”

## Event types

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

IBM WebSphere MQ instrumentation events have the following types:

- Queue manager events
- Channel and bridge events
- Performance events
- Configuration events
- Command events
- Logger events
- Local events

For each queue manager, each category of event has its own event queue. All events in that category result in an event message being put onto the same queue.

**This event queue:**

SYSTEM.ADMIN.QMGR.EVENT  
SYSTEM.ADMIN.CHANNEL.EVENT  
SYSTEM.ADMIN.PERFM.EVENT  
SYSTEM.ADMIN.CONFIG.EVENT  
SYSTEM.ADMIN.COMMAND.EVENT  
SYSTEM.ADMIN.LOGGER.EVENT  
SYSTEM.ADMIN.PUBSUB.EVENT

**Contains messages from:**

Queue manager events  
Channel events  
Performance events  
Configuration events  
Command events  
Logger events  
Gets events related to Publish/Subscribe. Only used with Multicast. For more information see, "Multicast application monitoring" on page 163.

By incorporating instrumentation events into your own system management application, you can monitor the activities across many queue managers, across many different nodes, and for multiple IBM WebSphere MQ applications. In particular, you can monitor all the nodes in your system from a single node (for those nodes that support IBM WebSphere MQ events) as shown in Figure 110.

Instrumentation events can be reported through a user-written reporting mechanism to an administration application that can present the events to an operator.

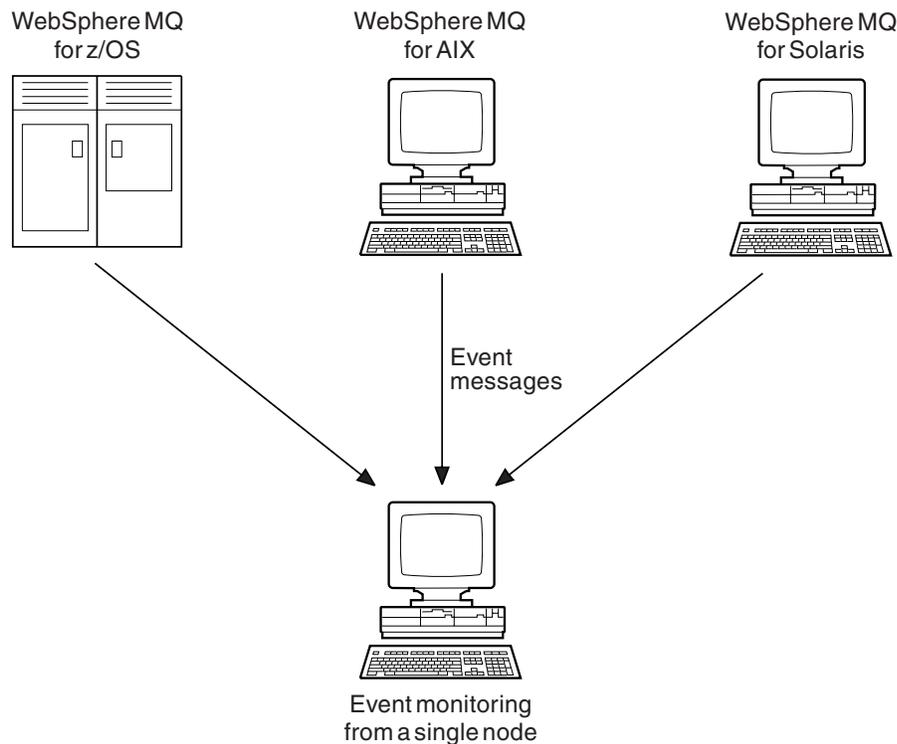


Figure 110. Monitoring queue managers across different platforms, on a single node

Instrumentation events also enable applications acting as agents for other administration networks, for example Tivoli NetView for z/OS, to monitor reports and create the appropriate alerts.

**Related concepts:**

“Queue manager events”

“Channel and bridge events” on page 836

“Performance events” on page 837

“Configuration events” on page 837

“Command events” on page 838

“Logger events” on page 838

“Queue manager events”

“Channel and bridge events” on page 836

“Performance events” on page 837

“Configuration events” on page 837

“Command events” on page 838

“Logger events” on page 838

**Related reference:**

“Event message data summary” on page 838

“Event message data summary” on page 838

**Queue manager events:**

Queue manager events are related to the use of resources within queue managers. For example, a queue manager event is generated if an application tries to put a message on a queue that does not exist.

The following examples are conditions that can cause a queue manager event:

- An application issues an MQI call that fails. The reason code from the call is the same as the reason code in the event message.

A similar condition can occur during the internal operation of a queue manager; for example, when generating a report message. The reason code in an event message might match an MQI reason code, even though it is not associated with any application. Do not assume that, because an event message reason code looks like an MQI reason code, the event was necessarily caused by an unsuccessful MQI call from an application.

- A command is issued to a queue manager and processing this command causes an event. For example:
  - A queue manager is stopped or started.
  - A command is issued where the associated user ID is not authorized for that command.

WebSphere MQ puts messages for queue manager events on the SYSTEM.ADMIN.QMGR.EVENT queue, and supports the following queue manager event types:

**Authority (on Windows, HP OpenVMS, and UNIX systems only)**

Authority events report an authorization, such as an application trying to open a queue for which it does not have the required authority, or a command being issued from a user ID that does not have the required authority. The authority event message can contain the following event data:

-  Not Authorized (type 1) (*WebSphere MQ V7.1 Reference*)
-  Not Authorized (type 2) (*WebSphere MQ V7.1 Reference*)
-  Not Authorized (type 3) (*WebSphere MQ V7.1 Reference*)
-  Not Authorized (type 4) (*WebSphere MQ V7.1 Reference*)
-  Not Authorized (type 5) (*WebSphere MQ V7.1 Reference*)
-  Not Authorized (type 6) (*WebSphere MQ V7.1 Reference*)

All authority events are valid on HP OpenVMS, Windows, and UNIX systems only.

### Inhibit

Inhibit events indicate that an MQPUT or MQGET operation has been attempted against a queue where the queue is inhibited for puts or gets, or against a topic where the topic is inhibited for publishes. The inhibit event message can contain the following event data:

-  Get Inhibited (*WebSphere MQ V7.1 Reference*)
-  Put Inhibited (*WebSphere MQ V7.1 Reference*)

**Local** Local events indicate that an application (or the queue manager) has not been able to access a local queue or other local object. For example, an application might try to access an object that has not been defined. The local event message can contain the following event data:

-  Alias Base Queue Type Error (*WebSphere MQ V7.1 Reference*)
-  Unknown Alias Base Queue (*WebSphere MQ V7.1 Reference*)
-  Unknown Object Name (*WebSphere MQ V7.1 Reference*)

### Remote

Remote events indicate that an application or the queue manager cannot access a remote queue on another queue manager. For example, the transmission queue to be used might not be correctly defined. The remote event message can contain the following event data:

-  Default Transmission Queue Type Error (*WebSphere MQ V7.1 Reference*)
-  Default Transmission Queue Usage Error (*WebSphere MQ V7.1 Reference*)
-  Queue Type Error (*WebSphere MQ V7.1 Reference*)
-  Remote Queue Name Error (*WebSphere MQ V7.1 Reference*)
-  Transmission Queue Type Error (*WebSphere MQ V7.1 Reference*)
-  Transmission Queue Usage Error (*WebSphere MQ V7.1 Reference*)
-  Unknown Default Transmission Queue (*WebSphere MQ V7.1 Reference*)
-  Unknown Remote Queue Manager (*WebSphere MQ V7.1 Reference*)
-  Unknown Transmission Queue (*WebSphere MQ V7.1 Reference*)

### Start and stop

Start and stop events indicate that a queue manager has been started or has been requested to stop or quiesce.

z/OS supports only start events.

Stop events are not recorded unless the default message-persistence of the SYSTEM.ADMIN.QMGR.EVENT queue is defined as persistent. The start and stop event message can contain the following event data:

-  Queue Manager Active (*WebSphere MQ V7.1 Reference*)
-  Queue Manager Not Active (*WebSphere MQ V7.1 Reference*)

For each event type in this list, you can set a queue manager attribute to enable or disable the event type.

## Channel and bridge events:

Channels report these events as a result of conditions detected during their operation. For example, when a channel instance is stopped.

Channel events are generated in the following circumstances:

- When a command starts or stops a channel.
- When a channel instance starts or stops.
- When a channel receives a conversion error warning when getting a message.
- When an attempt is made to create a channel automatically; the event is generated whether the attempt succeeds or fails.

**Note:** Client connections do not cause Channel Started or Channel Stopped events.

When a command is used to start a channel, an event is generated. Another event is generated when the channel instance starts. However, starting a channel by a listener, the `runmqchl` command, or a queue manager trigger message does not generate an event. In these cases, an event is generated only when the channel instance starts.

A successful start or stop channel command generates at least two events. These events are generated for both queue managers connected by the channel (providing they support events).

If a channel event is put on an event queue, an error condition causes the queue manager to create an event.

The event messages for channel and bridge events are put on the `SYSTEM.ADMIN.CHANNEL.EVENT` queue.

The channel event messages can contain the following event data:

-  Channel Activated (*WebSphere MQ V7.1 Reference*)
-  Channel Auto-definition Error (*WebSphere MQ V7.1 Reference*)
-  Channel Auto-definition OK (*WebSphere MQ V7.1 Reference*)
-  Channel Conversion Error (*WebSphere MQ V7.1 Reference*)
-  Channel Not Activated (*WebSphere MQ V7.1 Reference*)
-  Channel Started (*WebSphere MQ V7.1 Reference*)
-  Channel Stopped (*WebSphere MQ V7.1 Reference*)
-  Channel Stopped By User (*WebSphere MQ V7.1 Reference*)
-  Channel Blocked (*WebSphere MQ V7.1 Reference*)

## IMS bridge events (z/OS only)

These events are reported when an IMS bridge starts or stops.

The IMS bridge event messages can contain the following event data:

-  Bridge Started (*WebSphere MQ V7.1 Reference*)
-  Bridge Stopped (*WebSphere MQ V7.1 Reference*)

## SSL events

The only Secure Sockets Layer (SSL or TLS) event is the Channel SSL Error event. This event is reported when a channel using SSL or TLS fails to establish an SSL connection.

The SSL event messages can contain the following event data:

-  Channel SSL Error (*WebSphere MQ V7.1 Reference*)
-  Channel SSL Warning (*WebSphere MQ V7.1 Reference*)

## Performance events:

Performance events are notifications that a resource has reached a threshold condition. For example, a queue depth limit has been reached.

Performance events relate to conditions that can affect the performance of applications that use a specified queue. They are not generated for the event queues themselves.

The event type is returned in the command identifier field in the message data.

If a queue manager tries to put a queue manager event or performance event message on an event queue and an error that would typically create an event is detected, another event is not created and no action is taken.

MQGET and MQPUT calls within a unit of work can generate performance events regardless of whether the unit of work is committed or backed out.

The event messages for performance events are put on the SYSTEM.ADMIN.PERFM.EVENT queue.

There are two types of performance event:

### Queue depth events

Queue depth events relate to the number of messages on a queue; that is, how full or empty the queue is. These events are supported for shared queues. The queue depth event messages can contain the following event data:

-  Queue Depth High (*WebSphere MQ V7.1 Reference*)
-  Queue Depth Low (*WebSphere MQ V7.1 Reference*)
-  Queue Full (*WebSphere MQ V7.1 Reference*)

### Queue service interval events

Queue service interval events relate to whether messages are processed within a user-specified time interval. These events are not supported for shared queues.

WebSphere MQ for z/OS supports queue depth events for QSGDISP (SHARED) queues, but not service interval events. Queue manager and channel events remain unaffected by shared queues. The queue service event messages can contain the following event data:

-  Queue Service Interval High (*WebSphere MQ V7.1 Reference*)
-  Queue Service Interval OK (*WebSphere MQ V7.1 Reference*)

## Configuration events:

Configuration events are generated when a configuration event is requested explicitly, or automatically when an object is created, modified, or deleted.

A configuration event message contains information about the attributes of an object. For example, a configuration event message is generated if a namelist object is created, and contains information about the attributes of the namelist object.

The event messages for configuration events are put on the SYSTEM.ADMIN.CONFIG.EVENT queue.

There are four types of configuration event:

#### **Create object events**

Create object events are generated when an object is created. The event message contains the following event data:  Create object (*WebSphere MQ V7.1 Reference*).

#### **Change object events**

Change object events are generated when an object is changed. The event message contains the following event data:  Change object (*WebSphere MQ V7.1 Reference*).

#### **Delete object events**

Delete object events are generated when an object is deleted. The event message contains the following event data:  Delete object (*WebSphere MQ V7.1 Reference*).

#### **Refresh object events**

Refresh object events are generated by an explicit request to refresh. The event message contains the following event data:  Refresh object (*WebSphere MQ V7.1 Reference*).

#### **Command events:**

Command events are reported when an MQSC or PCF command runs successfully.

A command event message contains information about the origin, context, and content of a command. For example, a command event message is generated with such information if the MQSC command, ALTER QLOCAL, runs successfully.

The event messages for command events are put on the SYSTEM.ADMIN.COMMAND.EVENT queue.

Command events contain the following event data:  Command (*WebSphere MQ V7.1 Reference*).

#### **Logger events:**

Logger events are reported when a queue manager that uses linear logging starts writing log records to a new log extent or, on IBM i, to a new journal receiver. Logger events are not available with WebSphere MQ for z/OS.

A logger event message contains information specifying the log extents required by the queue manager to restart the queue manager, or for media recovery.

The event messages for logger events are put on the SYSTEM.ADMIN.LOGGER.EVENT queue.

The logger event message contains the following event data:  Logger (*WebSphere MQ V7.1 Reference*).

#### **Event message data summary:**

Use this summary to obtain information about the event data that each type of event message can contain.

| Event type           | See these topics                                                                                                                                       |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Authority events     |  Not Authorized (type 1) ( <i>WebSphere MQ V7.1 Reference</i> )       |
|                      |  Not Authorized (type 2) ( <i>WebSphere MQ V7.1 Reference</i> )       |
|                      |  Not Authorized (type 3) ( <i>WebSphere MQ V7.1 Reference</i> )       |
|                      |  Not Authorized (type 4) ( <i>WebSphere MQ V7.1 Reference</i> )       |
|                      |  Not Authorized (type 5) ( <i>WebSphere MQ V7.1 Reference</i> )       |
|                      |  Not Authorized (type 6) ( <i>WebSphere MQ V7.1 Reference</i> )       |
| Channel events       |  Channel Activated ( <i>WebSphere MQ V7.1 Reference</i> )             |
|                      |  Channel Auto-definition Error ( <i>WebSphere MQ V7.1 Reference</i> ) |
|                      |  Channel Auto-definition OK ( <i>WebSphere MQ V7.1 Reference</i> )    |
|                      |  Channel Blocked ( <i>WebSphere MQ V7.1 Reference</i> )               |
|                      |  Channel Conversion Error ( <i>WebSphere MQ V7.1 Reference</i> )      |
|                      |  Channel Not Activated ( <i>WebSphere MQ V7.1 Reference</i> )         |
|                      |  Channel Started ( <i>WebSphere MQ V7.1 Reference</i> )              |
|                      |  Channel Stopped ( <i>WebSphere MQ V7.1 Reference</i> )             |
|                      |  Channel Stopped By User ( <i>WebSphere MQ V7.1 Reference</i> )     |
| Command events       |  Command ( <i>WebSphere MQ V7.1 Reference</i> )                     |
| Configuration events |  Create object ( <i>WebSphere MQ V7.1 Reference</i> )               |
|                      |  Change object ( <i>WebSphere MQ V7.1 Reference</i> )               |
|                      |  Delete object ( <i>WebSphere MQ V7.1 Reference</i> )               |
|                      |  Refresh object ( <i>WebSphere MQ V7.1 Reference</i> )              |
| IMS Bridge events    |  Bridge Started ( <i>WebSphere MQ V7.1 Reference</i> )              |
|                      |  Bridge Stopped ( <i>WebSphere MQ V7.1 Reference</i> )              |
| Inhibit events       |  Get Inhibited ( <i>WebSphere MQ V7.1 Reference</i> )               |
|                      |  Put Inhibited ( <i>WebSphere MQ V7.1 Reference</i> )               |
| Local events         |  Alias Base Queue Type Error ( <i>WebSphere MQ V7.1 Reference</i> ) |
|                      |  Unknown Alias Base Queue ( <i>WebSphere MQ V7.1 Reference</i> )    |
|                      |  Unknown Object Name ( <i>WebSphere MQ V7.1 Reference</i> )         |

| Event type            | See these topics                                                                                                                                                       |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Logger events         |  <a href="#">Logger (WebSphere MQ V7.1 Reference)</a>                                 |
| Performance events    |  <a href="#">Queue Depth High (WebSphere MQ V7.1 Reference)</a>                       |
|                       |  <a href="#">Queue Depth Low (WebSphere MQ V7.1 Reference)</a>                        |
|                       |  <a href="#">Queue Full (WebSphere MQ V7.1 Reference)</a>                             |
|                       |  <a href="#">Queue Service Interval High (WebSphere MQ V7.1 Reference)</a>            |
|                       |  <a href="#">Queue Service Interval OK (WebSphere MQ V7.1 Reference)</a>              |
| Remote events         |  <a href="#">Default Transmission Queue Type Error (WebSphere MQ V7.1 Reference)</a>  |
|                       |  <a href="#">Default Transmission Queue Usage Error (WebSphere MQ V7.1 Reference)</a> |
|                       |  <a href="#">Queue Type Error (WebSphere MQ V7.1 Reference)</a>                       |
|                       |  <a href="#">Remote Queue Name Error (WebSphere MQ V7.1 Reference)</a>                |
|                       |  <a href="#">Transmission Queue Type Error (WebSphere MQ V7.1 Reference)</a>          |
|                       |  <a href="#">Transmission Queue Usage Error (WebSphere MQ V7.1 Reference)</a>         |
|                       |  <a href="#">Unknown Default Transmission Queue (WebSphere MQ V7.1 Reference)</a>    |
|                       |  <a href="#">Unknown Remote Queue Manager (WebSphere MQ V7.1 Reference)</a>         |
|                       |  <a href="#">Unknown Transmission Queue (WebSphere MQ V7.1 Reference)</a>           |
| SSL events            |  <a href="#">Channel SSL Error (WebSphere MQ V7.1 Reference)</a>                    |
| Start and stop events |  <a href="#">Queue Manager Active (WebSphere MQ V7.1 Reference)</a>                 |
|                       |  <a href="#">Queue Manager Not Active (WebSphere MQ V7.1 Reference)</a>             |

## Controlling events

You enable and disable events by specifying the appropriate values for queue manager, queue attributes, or both, depending on the type of event.

You must enable each instrumentation event that you want to be generated. For example, the conditions causing a Queue Full event are:

- Queue Full events are enabled for a specified queue, and
- An application issues an MQPUT request to put a message on that queue, but the request fails because the queue is full.

Enable and disable events by using any of the following techniques:

- IBM WebSphere MQ script commands (MQSC).
- The corresponding IBM WebSphere MQ PCF commands.
- The operations and control panels for queue managers on z/OS.
- The IBM WebSphere MQ Explorer.

**Note:** You can set attributes related to events for both queues and queue managers only by command. The MQI call MQSET does not support attributes related to events.

**Related concepts:**

- “Instrumentation events” on page 830
- “Automating administration tasks” on page 5
- “Using Programmable Command Formats” on page 7
- “Introducing the operations and control panels” on page 260

**Related reference:**

“Event types” on page 832



The MQSC commands (*WebSphere MQ V7.1 Reference*)

**Controlling queue manager events:**

You control queue manager events by using queue manager attributes. To enable queue manager events, set the appropriate queue manager attribute to ENABLED. To disable queue manager events, set the appropriate queue manager attribute to DISABLED.

To enable or disable queue manager events, use the MQSC command ALTER QMGR, specifying the appropriate queue manager attribute. Table 94 summarizes how to enable queue manager events. To disable a queue manager event, set the appropriate parameter to DISABLED.

*Table 94. Enabling queue manager events using MQSC commands*

| Event          | ALTER QMGR parameter |
|----------------|----------------------|
| Authority      | AUTHOREV (ENABLED)   |
| Inhibit        | INHIBTEV (ENABLED)   |
| Local          | LOCALEV (ENABLED)    |
| Remote         | REMOTEEV (ENABLED)   |
| Start and Stop | STRSTPEV (ENABLED)   |

**Controlling channel and bridge events:**

You control channel events by using queue manager attributes. To enable channel events, set the appropriate queue manager attribute to ENABLED. To disable channel events, set the appropriate queue manager attribute to DISABLED.

To enable or disable channels events use the MQSC command ALTER QMGR, specifying the appropriate queue manager attribute. Table 95 on page 842 summarizes how you enable channel and bridge events. To disable a queue manager event, set the appropriate parameter to DISABLED.

**Restriction:** Channel auto-definition events are not available on WebSphere MQ for z/OS.

Table 95. Enabling channel and bridge events using MQSC commands

| Event                          | ALTER QMGR parameter |
|--------------------------------|----------------------|
| Channel                        | CHLEV (ENABLED)      |
| Related to channel errors only | CHLEV (EXCEPTION)    |
| IMS Bridge                     | BRIDGEEV (ENABLED)   |
| SSL                            | SSLEV (ENABLED)      |
| Channel auto-definition        | CHADEV(ENABLED)      |

### Controlling performance events:

You control performance events using the PERFMEV queue manager attribute. To enable performance events, set PERFMEV to ENABLED. To disable performance events, set the PERFMEV queue manager attribute to DISABLED.

To set the PERFMEV queue manager attribute to ENABLED, use the following MQSC command:

```
ALTER QMGR PERFMEV (ENABLED)
```

To enable specific performance events, set the appropriate queue attribute. Also, specify the conditions that cause the event.

### Queue depth events

By default, all queue depth events are disabled. To configure a queue for any of the queue depth events:

1. Enable performance events on the queue manager.
2. Enable the event on the required queue.
3. Set the limits, if required, to the appropriate levels, expressed as a percentage of the maximum queue depth.

### Queue service interval events

To configure a queue for queue service interval events you must:

1. Enable performance events on the queue manager.
2. Set the control attribute for a Queue Service Interval High or OK event on the queue as required.
3. Specify the service interval time by setting the QSVCINT attribute for the queue to the appropriate length of time.

**Note:** When enabled, a queue service interval event can be generated at any appropriate time, not necessarily waiting until an MQI call for the queue is issued. However, if an MQI call is used on a queue to put or remove a message, any applicable performance event is generated at that time. The event is *not* generated when the elapsed time becomes equal to the service interval time.

### Controlling configuration, command, and logger events:

You control configuration, command, and logger events by using the queue manager attributes CONFIGEV, CMDEV, and LOGGEREV. To enable these events, set the appropriate queue manager attribute to ENABLED. To disable these events, set the appropriate queue manager attribute to DISABLED.

### Configuration events

To enable configuration events, set CONFIGEV to ENABLED. To disable configuration events, set CONFIGEV to DISABLED. For example, you can enable configuration events by using the following MQSC command:

```
ALTER QMGR CONFIGEV (ENABLED)
```

## Command events

To enable command events, set CMDEV to ENABLED. To enable command events for commands except DISPLAY MQSC commands and Inquire PCF commands, set the CMDEV to NODISPLAY. To disable command events, set CMDEV to DISABLED. For example, you can enable command events by using the following MQSC command:

```
ALTER QMGR CMDEV (ENABLED)
```

## Logger events

To enable logger events, set LOGGEREV to ENABLED. To disable logger events, set LOGGEREV to DISABLED. For example, you can enable logger events by using the following MQSC command:

```
ALTER QMGR LOGGEREV(ENABLED)
```

## Event queues

When an event occurs, the queue manager puts an event message on the defined event queue. The event message contains information about the event.

You can define event queues either as local queues, alias queues, or as local definitions of remote queues. If you define all your event queues as local definitions of the same remote queue on one queue manager, you can centralize your monitoring activities.

You must not define event queues as transmission queues, because event messages have formats that are incompatible with the message format that is required for transmission queues.

Shared event queues are local queues defined with the QSGDISP(SHARED) value.

For more information about defining shared queues on z/OS, see  Application programming with shared queues (*WebSphere MQ V7.1 Programming Guide*).

## When an event queue is unavailable

If an event occurs when the event queue is not available, the event message is lost. For example, if you do not define an event queue for a category of event, all event messages for that category are lost. The event messages are not, for example, saved on the dead-letter (undelivered-message) queue.

However, you can define the event queue as a remote queue. Then, if there is a problem on the remote system putting messages to the resolved queue, the event message arrives on the dead-letter queue of the remote system.

An event queue might be unavailable for many different reasons including:

- The queue has not been defined.
- The queue has been deleted.
- The queue is full.
- The queue has been put-inhibited.

The absence of an event queue does not prevent the event from occurring. For example, after a performance event, the queue manager changes the queue attributes and resets the queue statistics. This change happens whether the event message is put on the performance event queue or not. The same is true in the case of configuration and command events.

## Using triggered event queues

You can set up the event queues with triggers so that when an event is generated, the event message being put onto the event queue starts a user-written monitoring application. This application can process

the event messages and take appropriate action. For example, certain events might require an operator to be informed, other events might start an application that performs some administration tasks automatically.

Event queues can have trigger actions associated with them and can create trigger messages. However, if these trigger messages in turn cause conditions that would normally generate an event, no event is generated. Not generating an event in this instance ensures that looping does not occur.

**Related concepts:**

“Controlling events” on page 840

“Format of event messages”

 Application programming with shared queues (*WebSphere MQ V7.1 Programming Guide*)

 Conditions for a trigger event (*WebSphere MQ V7.1 Programming Guide*)

**Related reference:**

 QSGDisp (MQLONG) (*WebSphere MQ V7.1 Reference*)

## Format of event messages

Event messages contain information about an event and its cause. Like other WebSphere MQ messages, an event message has two parts: a message descriptor and the message data.

- The message descriptor is based on the MQMD structure.
- The message data consists of an *event header* and the *event data*. The event header contains the reason code that identifies the event type. Putting the event message, and any subsequent action, does not affect the reason code returned by the MQI call that caused the event. The event data provides further information about the event.

Typically, you process event messages with a system management application tailored to meet the requirements of the enterprise at which it runs.

When the queue managers in a queue sharing group detect the conditions for generating an event message, several queue managers can generate an event message for the shared queue, resulting in several event messages. To ensure that a system can correlate multiple event messages from different queue managers, these event messages have a unique correlation identifier (*CorrelId*) set in the message descriptor (MQMD).

**Related concepts:**

 Event message descriptions (*WebSphere MQ V7.1 Reference*)

**Related reference:**

“Activity report MQMD (message descriptor)” on page 928

“Activity report MQEPH (Embedded PCF header)” on page 933

“Activity report MQCFH (PCF header)” on page 934

 Event message reference (*WebSphere MQ V7.1 Reference*)

 Event message format (*WebSphere MQ V7.1 Reference*)

 Event message MQMD (message descriptor) (*WebSphere MQ V7.1 Reference*)

 Event message MQCFH (PCF header) (*WebSphere MQ V7.1 Reference*)

## Performance events

Performance events relate to conditions that can affect the performance of applications that use a specified queue. The scope of performance events is the queue. **MQPUT** calls and **MQGET** calls on one queue do not affect the generation of performance events on another queue.

Performance event messages can be generated at any appropriate time, not necessarily waiting until an MQI call for the queue is issued. However, if you use an MQI call on a queue to put or remove a message, any appropriate performance events are generated at that time.

Every performance event message that is generated is placed on the queue, `SYSTEM.ADMIN.PERFM.EVENT`.

The event data contains a reason code that identifies the cause of the event, a set of performance event statistics, and other data. The types of event data that can be returned in performance event messages are described in the following list:

-  Queue Depth High (*WebSphere MQ V7.1 Reference*)
-  Queue Depth Low (*WebSphere MQ V7.1 Reference*)
-  Queue Full (*WebSphere MQ V7.1 Reference*)
-  Queue Service Interval High (*WebSphere MQ V7.1 Reference*)
-  Queue Service Interval OK (*WebSphere MQ V7.1 Reference*)

Examples that illustrate the use of performance events assume that you set queue attributes by using the appropriate IBM WebSphere MQ commands (MQSC). On z/OS, you can also set queue attributes using the operations and controls panels for queue managers.

### Related concepts:

- “Queue service interval events” on page 846
- “Queue service interval events examples” on page 850
- “Queue depth events” on page 855
- “Queue depth events examples” on page 859
- “Queue service interval events” on page 846
- “Queue service interval events examples” on page 850
- “Queue depth events” on page 855
- “Queue depth events examples” on page 859

### Related reference:

- “Performance event statistics”
- “Event types” on page 832
- “Performance event statistics”
- “Event types” on page 832

## Performance event statistics

The performance event data in the event message contains statistics about the event. Use the statistics to analyze the behavior of a specified queue.

The event data in the event message contains information about the event for system management programs. For all performance events, the event data contains the names of the queue manager and the queue associated with the event. The event data also contains statistics related to the event. Table 96 on page 846 summarizes the event statistics that you can use to analyze the behavior of a queue. All the statistics refer to what has happened since the last time the statistics were reset.

Table 96. Performance event statistics

| Parameter      | Description                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------|
| TimeSinceReset | The elapsed time since the statistics were last reset.                                                          |
| HighQDepth     | The maximum number of messages on the queue since the statistics were last reset.                               |
| MsgEnqCount    | The number of messages enqueued (the number of MQPUT calls to the queue), since the statistics were last reset. |
| MsgDeqCount    | The number of messages dequeued (the number of MQGET calls to the queue), since the statistics were last reset. |

Performance event statistics are reset when any of the following changes occur:

- A performance event occurs (statistics are reset on all active queue managers).
- A queue manager stops and restarts.
- The PCF command, Reset Queue Statistics, is issued from an application program.
- On z/OS only, the RESET QSTATS command is issued at the console.

**Related concepts:**

“Performance events” on page 845

“The service timer” on page 848

“Rules for queue service interval events” on page 848

**Related tasks:**

“Enabling queue service interval events” on page 849

**Related reference:**

 Queue Depth High (*WebSphere MQ V7.1 Reference*)

 Reset Queue Statistics (*WebSphere MQ V7.1 Reference*)

 RESET QSTATS (*WebSphere MQ V7.1 Reference*)

**Queue service interval events**

Queue service interval events indicate whether an operation was performed on a queue within a user-defined time interval called the *service interval*. Depending on your installation, you can use queue service interval events to monitor whether messages are being taken off queues quickly enough.

Queue service interval events are *not* supported on shared queues.

The following types of queue service interval events can occur, where the term *get operation* refers to an MQGET call or an activity that removes a messages from a queue, such as using the **CLEAR QLOCAL** command:

**Queue Service Interval OK**

Indicates that after one of the following operations:

- An MQPUT call
- A get operation that leaves a non-empty queue

a get operation was performed within a user-defined time period, known as the *service interval*.

Only a get operation can cause the Queue Service Interval OK event message. Queue Service Interval OK events are sometimes described as OK events.

**Queue Service Interval High**

Indicates that after one of the following operations:

- An MQPUT call
  - A get operation that leaves a non-empty queue
- a get operation was **not** performed within a user-defined service interval.

Either a get operation or an MQPUT call can cause the Queue Service Interval High event message. Queue Service Interval High events are sometimes described as High events.

To enable both Queue Service Interval OK and Queue Service Interval High events, set the `QServiceIntervalEvent` control attribute to High. Queue Service Interval OK events are automatically enabled when a Queue Service Interval High event is generated. You do not need to enable Queue Service Interval OK events independently.

OK and High events are mutually exclusive, so if one is enabled the other is disabled. However, both events can be simultaneously disabled.

Figure 111 shows a graph of queue depth against time. At time P1, an application issues an MQPUT, to put a message on the queue. At time G1, another application issues an MQGET to remove the message from the queue.

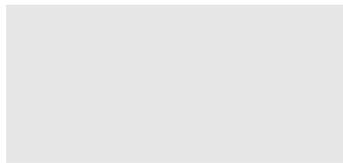


Figure 111. Understanding queue service interval events

The possible outcomes of queue service interval events are as follows:

- If the elapsed time between the put and the get is less than or equal to the service interval:
  - A *Queue Service Interval OK* event is generated at time G1, if queue service interval events are enabled
- If the elapsed time between the put and get is greater than the service interval:
  - A *Queue Service Interval High* event is generated at time G1, if queue service interval events are enabled.

The algorithm for starting the service timer and generating events is described in “Rules for queue service interval events” on page 848.

**Related reference:**

-  Queue Service Interval OK (*WebSphere MQ V7.1 Reference*)
-  Queue Service Interval High (*WebSphere MQ V7.1 Reference*)
-  QServiceIntervalEvent (MQLONG) (*WebSphere MQ V7.1 Reference*)
-  QServiceIntervalEvent (10-digit signed integer) (*WebSphere MQ V7.1 Reference*)
-  ServiceIntervalEvent property (*WebSphere MQ V7.1 Programming Guide*)

## The service timer:

Queue service interval events use an internal timer, called the *service timer*, which is controlled by the queue manager. The service timer is used only if a queue service interval event is enabled.

### What precisely does the service timer measure?

The service timer measures the elapsed time between an MQPUT call to an empty queue or a get operation, and the next put or get, provided the queue depth is nonzero between these two operations.

### When is the service timer active?

The service timer is always active (running), if the queue has messages on it (depth is nonzero) and a queue service interval event is enabled. If the queue becomes empty (queue depth zero), the timer is put into an OFF state, to be restarted on the next put.

### When is the service timer reset?

The service timer is always reset after a get operation. It is also reset by an MQPUT call to an empty queue. However, it is not necessarily reset on a queue service interval event.

### How is the service timer used?

Following a get operation or an MQPUT call, the queue manager compares the elapsed time as measured by the service timer, with the user-defined service interval. The result of this comparison is that:

- An OK event is generated if there is a get operation and the elapsed time is less than or equal to the service interval, AND this event is enabled.
- A high event is generated if the elapsed time is greater than the service interval, AND this event is enabled.

### Can applications read the service timer?

No, the service timer is an internal timer that is not available to applications.

### What about the *TimeSinceReset* parameter?

The *TimeSinceReset* parameter is returned as part of the event statistics in the event data. It specifies the time between successive queue service interval events, unless the event statistics are reset.

## Rules for queue service interval events:

Formal rules control when the service timer is set and queue service interval events are generated.

### Rules for the service timer

The service timer is reset to zero and restarted as follows:

- After an MQPUT call to an empty queue.
- After an MQGET call, if the queue is not empty after the MQGET call.

The resetting of the timer does not depend on whether an event has been generated.

At queue manager startup the service timer is set to startup time if the queue depth is greater than zero.

If the queue is empty following a get operation, the timer is put into an OFF state.

### Queue Service Interval High events

The Queue Service Interval event must be enabled (set to HIGH).

Queue Service Interval High events are automatically enabled when a Queue Service Interval OK event is generated.

If the service time is greater than the service interval, an event is generated on, or before, the next MQPUT or get operation.

### Queue Service Interval OK events

Queue Service Interval OK events are automatically enabled when a Queue Service Interval High event is generated.

If the service time (elapsed time) is less than or equal to the service interval, an event is generated on, or before, the next get operation.

#### Related information:

“Enabling queue service interval events”

#### Enabling queue service interval events:

To configure a queue for queue service interval events you set the appropriate queue manager and queue attributes.

#### About this task

The high and OK events are mutually exclusive; that is, when one is enabled, the other is automatically disabled:

- When a high event is generated on a queue, the queue manager automatically disables high events and enables OK events for that queue.
- When an OK event is generated on a queue, the queue manager automatically disables OK events and enables high events for that queue.

Table 97. Enabling queue service interval events using MQSC

| Queue service interval event                                                                 | Queue attributes                                                                     |
|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Queue Service Interval High<br>Queue Service Interval OK<br>No queue service interval events | QSVCI EV (HIGH)<br>QSVCI EV (OK)<br>QSVCI EV (NONE)                                  |
| Service interval                                                                             | QSVCI NT ( <i>tt</i> ) where <i>tt</i> is the service interval time in milliseconds. |

Perform the following steps to enable queue service interval events:

#### Procedure

1. Set the queue manager attribute PERFMEV to ENABLED. Performance events are enabled on the queue manager.
2. Set the control attribute, QSVCI EV, for a Queue Service Interval High or OK event on the queue, as required.
3. Set the QSVCI NT attribute for the queue to specify the appropriate service interval time.

#### Example

To enable Queue Service Interval High events with a service interval time of 10 seconds (10 000 milliseconds) use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QSVCI(10000) QSVCI(10000) QSVCI(HIGH)
```

## Queue service interval events examples

Use these examples to understand the information that you can obtain from queue service interval events

The three examples provide progressively more complex illustrations of the use of queue service interval events.

The figures accompanying the examples have the same structure:

- Figure 1 is a graph of queue depth against time, showing individual MQGET calls and MQPUT calls.
- The Commentary section shows a comparison of the time constraints. There are three time periods that you must consider:
  - The user-defined service interval.
  - The time measured by the service timer.
  - The time since event statistics were last reset (TimeSinceReset in the event data).
- The Event statistics summary section shows which events are enabled at any instant and what events are generated.

The examples illustrate the following aspects of queue service interval events:

- How the queue depth varies over time.
- How the elapsed time as measured by the service timer compares with the service interval.
- Which event is enabled.
- Which events are generated.

**Remember:** Example 1 shows a simple case where the messages are intermittent and each message is removed from the queue before the next one arrives. From the event data, you know that the maximum number of messages on the queue was one. You can, therefore, work out how long each message was on the queue.

However, in the general case, where there is more than one message on the queue and the sequence of MQGET calls and MQPUT calls is not predictable, you cannot use queue service interval events to calculate how long an individual message remains on a queue. The TimeSinceReset parameter, which is returned in the event data, can include a proportion of time when there are no messages on the queue. Therefore any results you derive from these statistics are implicitly averaged to include these times.

### Related concepts:

“Queue service interval events: example 1”

“Queue service interval events: example 2” on page 852

“Queue service interval events: example 3” on page 854

“Queue service interval events” on page 846

“The service timer” on page 848

### Queue service interval events: example 1:

A basic sequence of MQGET calls and MQPUT calls, where the queue depth is always one or zero.

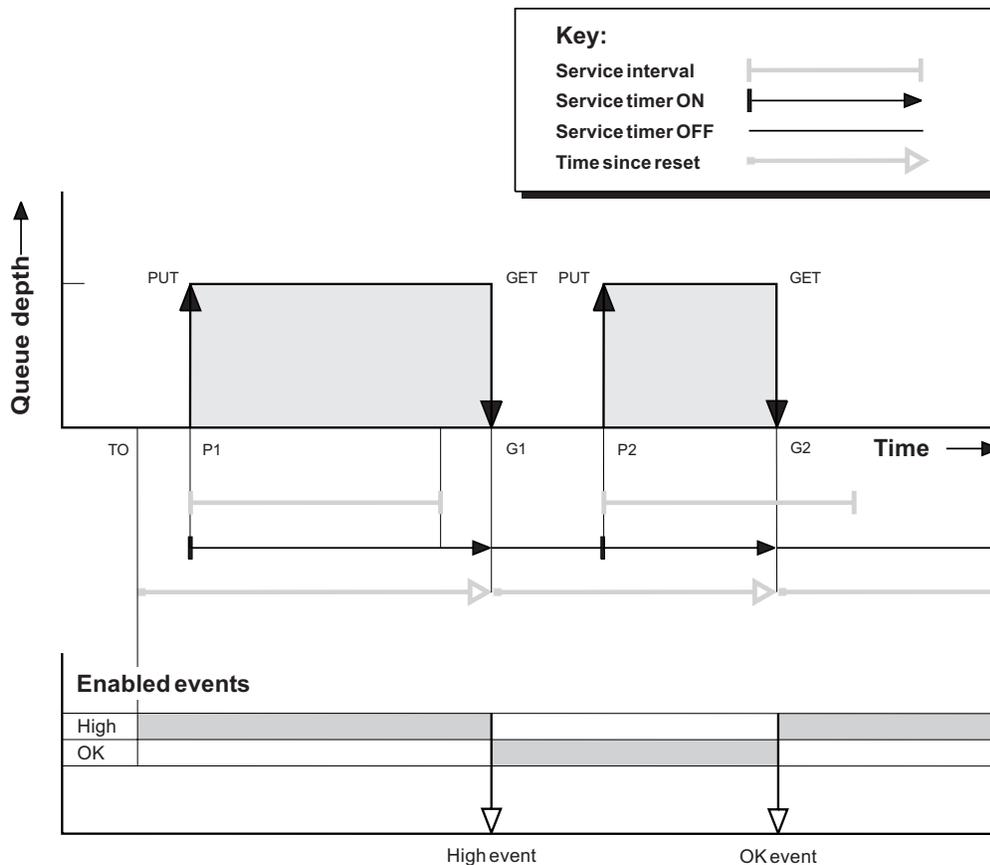


Figure 112. Queue service interval events - example 1

### Commentary

- At P1, an application puts a message onto an empty queue. This starts the service timer.  
Note that T0 might be queue manager startup time.
- At G1, another application gets the message from the queue. Because the elapsed time between P1 and G1 is greater than the service interval, a Queue Service Interval High event is generated on the MQGET call at G1. When the high event is generated, the queue manager resets the event control attribute so that:
  - The OK event is automatically enabled.
  - The high event is disabled.
Because the queue is now empty, the service timer is switched to an OFF state.
- At P2, a second message is put onto the queue. This restarts the service timer.
- At G2, the message is removed from the queue. However, because the elapsed time between P2 and G2 is less than the service interval, a Queue Service Interval OK event is generated on the MQGET call at G2. When the OK event is generated, the queue manager resets the control attribute so that:
  - The high event is automatically enabled.
  - The OK event is disabled.
Because the queue is empty, the service timer is again switched to an OFF state.

### Event statistics summary

Table 98 on page 852 summarizes the event statistics for this example.

Table 98. Event statistics summary for example 1

|                | Event 1      | Event 2       |
|----------------|--------------|---------------|
| Time of event  | T(G1)        | T(G2)         |
| Type of event  | High         | OK            |
| TimeSinceReset | T(G1) - T(0) | T(G2) - T(G1) |
| HighQDepth     | 1            | 1             |
| MsgEnqCount    | 1            | 1             |
| MsgDeqCount    | 1            | 1             |

The middle part of Figure 112 on page 851 shows the elapsed time as measured by the service timer compared to the service interval for that queue. To see whether a queue service interval event might occur, compare the length of the horizontal line representing the service timer (with arrow) to that of the line representing the service interval. If the service timer line is longer, and the Queue Service Interval High event is enabled, a Queue Service Interval High event occurs on the next get. If the timer line is shorter, and the Queue Service Interval OK event is enabled, a Queue Service Interval OK event occurs on the next get.

**Queue service interval events: example 2:**

A sequence of MQPUT calls and MQGET calls, where the queue depth is not always one or zero.

This example also shows instances of the timer being reset without events being generated, for example, at time P2.

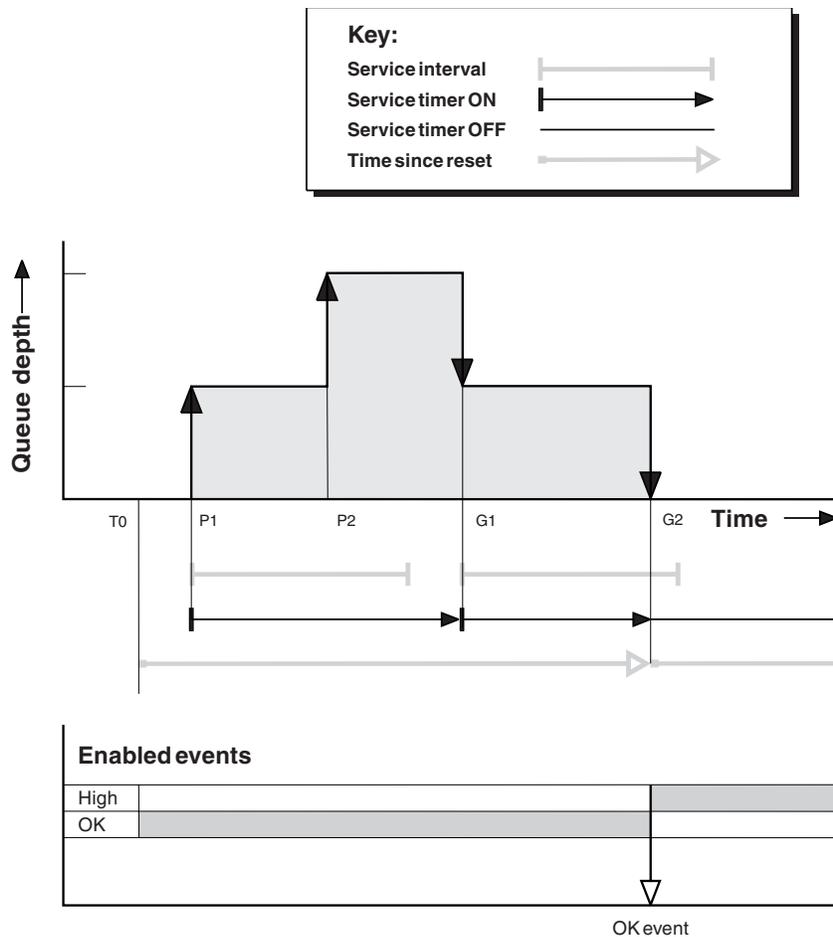


Figure 113. Queue service interval events - example 2

### Commentary

In this example, OK events are enabled initially and queue statistics were reset at time T0.

1. At P1, the first put starts the service timer.
2. At P2, the second put does not generate an event because a put cannot cause an OK event.
3. At G1, the service interval has now been exceeded and therefore an OK event is not generated. However, the MQGET call causes the service timer to be reset.
4. At G2, the second get occurs within the service interval and this time an OK event is generated. The queue manager resets the event control attribute so that:
  - a. The high event is automatically enabled.
  - b. The OK event is disabled.

Because the queue is now empty, the service timer is switched to an OFF state.

### Event statistics summary

Table 99 on page 854 summarizes the event statistics for this example.

Table 99. Event statistics summary for example 2

|                |                |
|----------------|----------------|
|                | <b>Event 2</b> |
| Time of event  | T(G2)          |
| Type of event  | OK             |
| TimeSinceReset | T(G2) - T(0)   |
| HighQDepth     | 2              |
| MsgEnqCount    | 2              |
| MsgDeqCount    | 2              |

**Queue service interval events: example 3:**

A sequence of MQGET calls and MQPUT calls that is more sporadic than the previous examples.

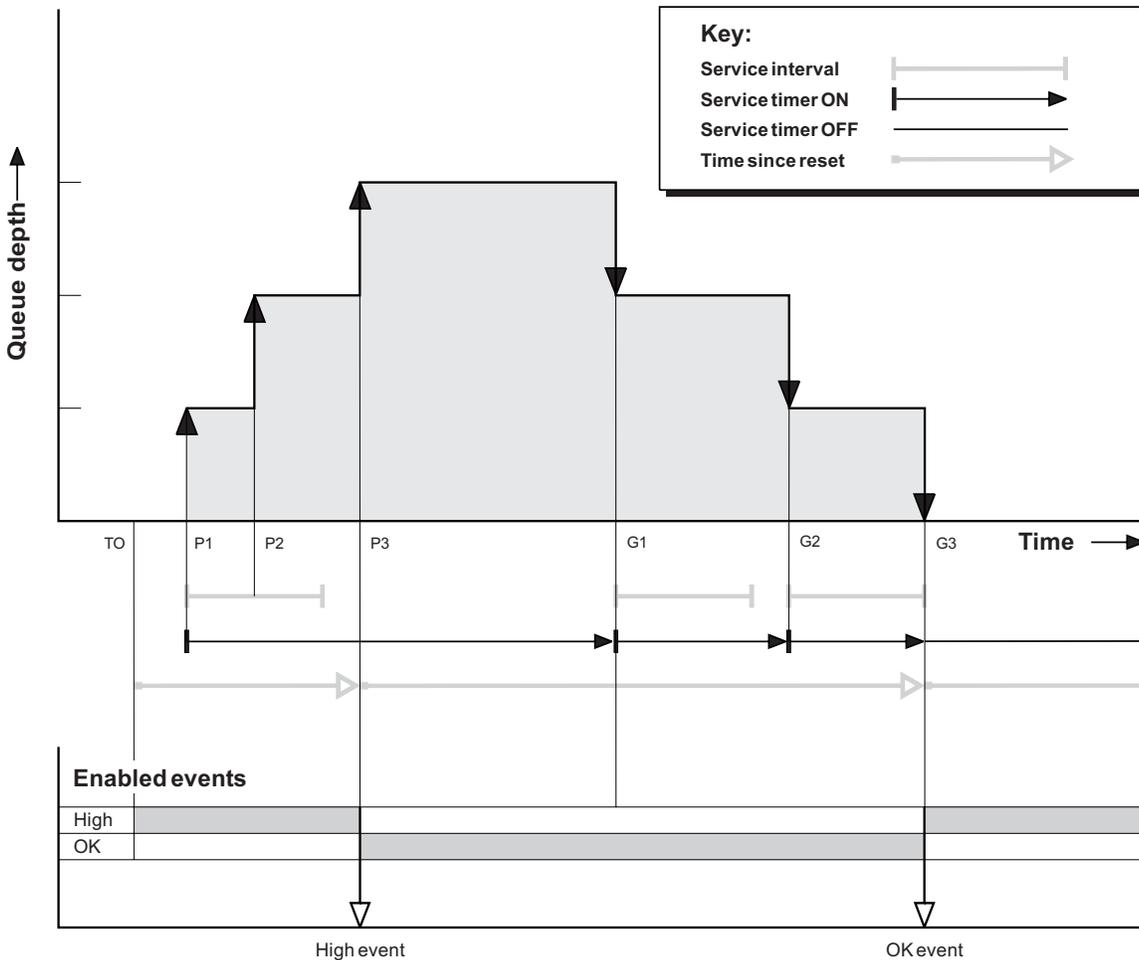


Figure 114. Queue service interval events - example 3

**Commentary**

1. At time T(0), the queue statistics are reset and Queue Service Interval High events are enabled.
2. At P1, the first put starts the service timer.
3. At P2, the second put increases the queue depth to two. A high event is not generated here because the service interval time has not been exceeded.

4. At P3, the third put causes a high event to be generated. (The timer has exceeded the service interval.) The timer is not reset because the queue depth was not zero before the put. However, OK events are enabled.
5. At G1, the MQGET call does not generate an event because the service interval has been exceeded and OK events are enabled. The MQGET call does, however, reset the service timer.
6. At G2, the MQGET call does not generate an event because the service interval has been exceeded and OK events are enabled. Again, the MQGET call resets the service timer.
7. At G3, the third get empties the queue and the service timer is *equal* to the service interval. Therefore an OK event is generated. The service timer is reset and high events are enabled. The MQGET call empties the queue, and this puts the timer in the OFF state.

### Event statistics summary

Table 100 summarizes the event statistics for this example.

*Table 100. Event statistics summary for example 3*

|                | Event 1      | Event 2       |
|----------------|--------------|---------------|
| Time of event  | T(P3)        | T(G3)         |
| Type of event  | High         | OK            |
| TimeSinceReset | T(P3) - T(0) | T(G3) - T(P3) |
| HighQDepth     | 3            | 3             |
| MsgEnqCount    | 3            | 0             |
| MsgDeqCount    | 0            | 3             |

### Queue depth events

Queue depth events are related to the queue depth, that is, the number of messages on the queue.

In WebSphere MQ applications, queues must not become full. If they do, applications can no longer put messages on the queue that they specify. Although the message is not lost if this occurs, a full queue can cause considerable inconvenience. The number of messages can build up on a queue if the messages are being put onto the queue faster than the applications that process them can take them off.

The solution to this problem depends on the particular circumstances, but might involve:

- Diverting some messages to another queue.
- Starting new applications to take more messages off the queue.
- Stopping nonessential message traffic.
- Increasing the queue depth to overcome a transient maximum.

Advance warning that problems might be on their way makes it easier to take preventive action. For this purpose, WebSphere MQ provides the following queue depth events:

#### Queue Depth High events

Indicate that the queue depth has increased to a predefined threshold called the Queue Depth High limit.

#### Queue Depth Low events

Indicate that the queue depth has decreased to a predefined threshold called the Queue Depth Low limit.

#### Queue Full events

Indicate that the queue has reached its maximum depth, that is, the queue is full.

A Queue Full Event is generated when an application attempts to put a message on a queue that has reached its maximum depth. Queue Depth High events give advance warning that a queue is filling up. This means that having received this event, the system administrator needs to take some preventive action. You can configure the queue manager such that, if the preventive action is successful and the queue depth drops to a safer level, the queue manager generates a Queue Depth Low event.

The first queue depth event example illustrates the effect of presumed action preventing the queue becoming full.

**Related concepts:**

“Shared queues and queue depth events (WebSphere MQ for z/OS)” on page 858

“Queue depth events examples” on page 859

**Related tasks:**

“Enabling queue depth events”

**Related reference:**

 Queue Full (*WebSphere MQ V7.1 Reference*)

 Queue Depth High (*WebSphere MQ V7.1 Reference*)

 Queue Depth Low (*WebSphere MQ V7.1 Reference*)

**Enabling queue depth events:**

To configure a queue for any of the queue depth events you set the appropriate queue manager and queue attributes.

**About this task**

By default, all queue depth events are disabled. When enabled, queue depth events are generated as follows:

- A Queue Depth High event is generated when a message is put on the queue, causing the queue depth to be greater than or equal to the value determined by the Queue Depth High limit.
  - A Queue Depth High event is automatically enabled by a Queue Depth Low event on the same queue.
  - A Queue Depth High event automatically enables both a Queue Depth Low and a Queue Full event on the same queue.
- A Queue Depth Low event is generated when a message is removed from a queue by a get operation causing the queue depth to be less than or equal to the value determined by the Queue Depth Low limit.
  - A Queue Depth Low event is automatically enabled by a Queue Depth High event or a Queue Full event on the same queue.
  - A Queue Depth Low event automatically enables both a Queue Depth High and a Queue Full event on the same queue.
- A Queue Full event is generated when an application is unable to put a message onto a queue because the queue is full.
  - A Queue Full event is automatically enabled by a Queue Depth High or a Queue Depth Low event on the same queue.
  - A Queue Full event automatically enables a Queue Depth Low event on the same queue.

Perform the following steps to configure a queue for any of the queue depth events:

## Procedure

1. Enable performance events on the queue manager, using the queue manager attribute PERFMEEV.
2. Set one of the following attributes to enable the event on the required queue:
  - *QDepthHighEvent* (QDPHIEV in MQSC)
  - *QDepthLowEvent* (QDPLOEV in MQSC)
  - *QDepthMaxEvent* (QDPMAXEV in MQSC)
3. Optional: To set the limits, assign the following attributes, as a percentage of the maximum queue depth:
  - *QDepthHighLimit* (QDEPTHHI in MQSC)
  - *QDepthLowLimit* (QDEPTHLO in MQSC)

**Restriction:** QDEPTHHI must not be less than QDEPTHLO.

If QDEPTHHI equals QDEPTHLO an event message is generated every time the queue depth passes the value in either direction, because the high threshold is enabled when the queue depth is below the value and the low threshold is enabled when the depth is above the value.

## Results

### Note:

A Queue Depth Low event is not generated when expired messages are removed from a queue by a get operation causing the queue depth to be less than, or equal to, the value determined by the Queue Depth Low limit.

WebSphere MQ generates the low event message only during a successful get operation. Therefore, when the expired messages are removed from the queue, no queue depth low event message is generated.

Additionally, after the removal of these expired messages from the queue, queue depth high event and queue depth low event are not reset.

### Example

To enable Queue Depth High events on the queue MYQUEUE with a limit set at 80%, use the following MQSC commands:

```
ALTER QMGR PERFMEEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDEPTHHI(80) QDPHIEV(ENABLED)
```

To enable Queue Depth Low events on the queue MYQUEUE with a limit set at 20%, use the following MQSC commands:

```
ALTER QMGR PERFMEEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDEPTHLO(20) QDPLOEV(ENABLED)
```

To enable Queue Full events on the queue MYQUEUE, use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDPMAXEV(ENABLED)
```

### Shared queues and queue depth events (WebSphere MQ for z/OS):

Event monitoring is more straightforward for an application that uses shared queues if all the queue managers in the queue-sharing group have the same setting for the PERFMEV attribute.

When a queue depth event occurs on a shared queue, and the queue manager attribute PERFMEV is set to ENABLED, the queue managers in the queue-sharing group produce an event message. If PERFMEV is set to DISABLED on some of the queue managers, event messages are not produced by those queue managers, making event monitoring from an application more difficult. For more straightforward monitoring, give each queue manager the same setting for the PERFMEV attribute.

This event message that each queue manager generates represents its individual usage of the shared queue. If a queue manager performs no activity on the shared queue, various values in the event message are null or zero. You can use null event messages as follows:

- Ensure that each active queue manager in a queue-sharing group generates one event message
- Highlight cases of no activity on a shared queue for the queue manager that produced the event message

### Coordinating queue manager

When a queue manager issues a queue depth event, it updates the shared queue object definition to toggle the active performance event attributes. For example, depending on the definition of the queue attributes, a Queue Depth High event enables a Queue Depth Low and a Queue Full event. After updating the shared queue object successfully, the queue manager that detected the performance event initially becomes the *coordinating queue manager*.

If enabled for performance events, the coordinating queue manager performs the following actions:

1. Issues an event message that captures all shared queue performance data it has gathered since the last time an event message was created, or since the queue statistics were last reset. The message descriptor (MQMD) of this message contains a unique correlation identifier (*CorrelId*) created by the coordinating queue manager.
2. Broadcasts to all other *active* queue managers in the same queue-sharing group to request the production of an event message for the shared queue. The broadcast contains the correlation identifier created by the coordinating queue manager for the set of event messages.

Having received a request from the coordinating queue manager, if there is an active queue manager in the queue-sharing group that is enabled for performance events, that active queue manager issues an event message for the shared queue. The event message that is issued contains information about all the operations performed by the receiving (active) queue manager since the last time an event message was created, or since the statistics were last reset. The message descriptor (MQMD) of this event message contains the unique correlation identifier (*CorrelId*) specified by the coordinating queue manager.

When performance events occur on a shared queue,  $n$  event messages are produced, where  $n$  is a number from 1 to the number of active queue managers in the queue-sharing group. Each event message contains data that relates to the shared queue activity for the queue manager that generated the event message.

You can view event message data for a shared queue using the following views:

#### Queue-sharing view

Collects all data from event messages with the same correlation identifier

## Queue manager view

Each event message shows how much it has been used by its originating queue manager

## Differences between shared and nonshared queues

Enabling queue depth events on shared queues differs from enabling them on nonshared queues. A key difference is that events are switched on for shared queues even if PERFMEV is DISABLED on the queue manager. This is not the case for nonshared queues.

Consider the following example, which illustrates this difference:

- QM1 is a queue manager with *PerformanceEvent* (PERFMEV in MQSC) set to DISABLED.
- SQ1 is a shared queue with QSGDISP set to (SHARED) QLOCAL in MQSC.
- LQ1 is a nonshared queue with QSGDISP set to (QMGR) QLOCAL in MQSC.

Both queues have the following attributes set on their definitions:

- QDPHIEV (ENABLED)
- QDPLOEV (DISABLED)
- QDPMAXEV (DISABLED)

If messages are placed on both queues so that the depth meets or exceeds the QDEPTHHI threshold, the QDPHIEV value on SQ1 switches to DISABLED. Also, QDPLOEV and QDPMAXEV are switched to ENABLED. SQ1's attributes are automatically switched for each performance event at the time the event criteria are met.

In contrast the attributes for LQ1 remain unchanged until PERFMEV on the queue manager is ENABLED. This means that if the queue manager's PERFMEV attribute is ENABLED, DISABLED and then re-ENABLED for example, the performance event settings on shared queues might not be consistent with those of nonshared queues, even though they might have initially been the same.

## Queue depth events examples

Use these examples to understand the information that you can obtain from queue depth events

The first example provides a basic illustration of queue depth events. The second example is more extensive, but the principles are the same as for the first example. Both examples use the same queue definition, as follows:

The queue, MYQUEUE1, has a maximum depth of 1000 messages. The high queue depth limit is 80% and the low queue depth limit is 20%. Initially, Queue Depth High events are enabled, while the other queue depth events are disabled.

The WebSphere MQ commands (MQSC) to configure this queue are:

```
ALTER QMGR PERFMEV(ENABLED)
```

```
DEFINE QLOCAL('MYQUEUE1') MAXDEPTH(1000) QDPMAXEV(DISABLED) QDEPTHHI(80)
QDPHIEV(ENABLED) QDEPTHLO(20) QDPLOEV(DISABLED)
```

**Related concepts:**

“Queue depth events: example 1”

“Queue depth events: example 2” on page 861

“Queue depth events” on page 855

**Related tasks:**

“Enabling queue depth events” on page 856

**Related reference:**

 The MQSC commands (*WebSphere MQ V7.1 Reference*)

**Queue depth events: example 1:**

A basic sequence of queue depth events.

Figure 115 shows the variation of queue depth over time.

*Figure 115. Queue depth events (1)*

**Commentary**

1. At T(1), the queue depth is increasing (more MQPUT calls than MQGET calls) and crosses the Queue Depth Low limit. No event is generated at this time.
2. The queue depth continues to increase until T(2), when the depth high limit (80%) is reached and a Queue Depth High event is generated.  
This enables both Queue Full and Queue Depth Low events.
3. The (presumed) preventive actions instigated by the event prevent the queue from becoming full. By time T(3), the Queue Depth High limit has been reached again, this time from above. No event is generated at this time.

4. The queue depth continues to fall until T(4), when it reaches the depth low limit (20%) and a Queue Depth Low event is generated.

This enables both Queue Full and Queue Depth High events.

### Event statistics summary

Table 101 summarizes the queue event statistics and Table 102 summarizes which events are enabled.

*Table 101. Event statistics summary for queue depth events (example 1)*

|                                              | Event 2          | Event 4         |
|----------------------------------------------|------------------|-----------------|
| Time of event                                | T(2)             | T(4)            |
| Type of event                                | Queue Depth High | Queue Depth Low |
| TimeSinceReset                               | T(2) - T(0)      | T(4) - T(2)     |
| HighQDepth (Maximum queue depth since reset) | 800              | 900             |
| MsgEnqCount                                  | 1157             | 1220            |
| MsgDeqCount                                  | 357              | 1820            |

*Table 102. Summary showing which events are enabled*

| Time period  | Queue Depth High event | Queue Depth Low event | Queue Full event |
|--------------|------------------------|-----------------------|------------------|
| Before T(1)  | ENABLED                | -                     | -                |
| T(1) to T(2) | ENABLED                | -                     | -                |
| T(2) to T(3) | -                      | ENABLED               | ENABLED          |
| T(3) to T(4) | -                      | ENABLED               | ENABLED          |
| After T(4)   | ENABLED                | -                     | ENABLED          |

### Queue depth events: example 2:

A more extensive sequence of queue depth events.

Figure 116 on page 862 shows the variation of queue depth over time.

*Figure 116. Queue depth events (2)*

### **Commentary**

1. No Queue Depth Low event is generated at the following times:
  - T(1) (Queue depth increasing, and not enabled)
  - T(2) (Not enabled)
  - T(3) (Queue depth increasing, and not enabled)
2. At T(4) a Queue Depth High event occurs. This enables both Queue Full and Queue Depth Low events.
3. At T(9) a Queue Full event occurs **after** the first message that cannot be put on the queue because the queue is full.
4. At T(12) a Queue Depth Low event occurs.

### **Event statistics summary**

Table 103 on page 863 summarizes the queue event statistics and Table 104 on page 863 summarizes which events are enabled at different times for this example.

Table 103. Event statistics summary for queue depth events (example 2)

|                | Event 4          | Event 6         | Event 8          | Event 9     | Event 12        |
|----------------|------------------|-----------------|------------------|-------------|-----------------|
| Time of event  | T(4)             | T(6)            | T(8)             | T(9)        | T(12)           |
| Type of event  | Queue Depth High | Queue Depth Low | Queue Depth High | Queue Full  | Queue Depth Low |
| TimeSinceReset | T(4) - T(0)      | T(6) - T(4)     | T(8) - T(6)      | T(9) - T(8) | T(12) - T(9)    |
| HighQDepth     | 800              | 855             | 800              | 1000        | 1000            |
| MsgEnqCount    | 1645             | 311             | 1377             | 324         | 221             |
| MsgDeqCount    | 845              | 911             | 777              | 124         | 1021            |

Table 104. Summary showing which events are enabled

| Time period   | Queue Depth High event | Queue Depth Low event | Queue Full event |
|---------------|------------------------|-----------------------|------------------|
| T(0) to T(4)  | ENABLED                | -                     | -                |
| T(4) to T(6)  | -                      | ENABLED               | ENABLED          |
| T(6) to T(8)  | ENABLED                | -                     | ENABLED          |
| T(8) to T(9)  | -                      | ENABLED               | ENABLED          |
| T(9) to T(12) | -                      | ENABLED               | -                |
| After T(12)   | ENABLED                | -                     | ENABLED          |

**Note:** Events are out of syncpoint. Therefore you could have an empty queue, then fill it up causing an event, then roll back all of the messages under the control of a syncpoint manager. However, event enabling has been automatically set, so that the next time the queue fills up, no event is generated.

## Configuration events

Configuration events are notifications that are generated when an object is created, changed, or deleted, and can also be generated by explicit requests.

Configuration events notify you about changes to the attributes of an object. There are four types of configuration events:

- Create object events
- Change object events
- Delete object events
- Refresh object events

The event data contains the following information:

### Origin information

comprises the queue manager from where the change was made, the ID of the user that made the change, and how the change came about, for example by a console command.

### Context information

a replica of the context information in the message data from the command message.

Context information is included in the event data only when the command was entered as a message on the SYSTEM.COMMAND.INPUT queue.

### Object identity

comprises the name, type and disposition of the object.

### Object attributes

comprises the values of all the attributes in the object.

In the case of change object events, two messages are generated, one with the information before the change, the other with the information after.

Every configuration event message that is generated is placed on the queue SYSTEM.ADMIN.CONFIG.EVENT.

**Related concepts:**

“Configuration event usage” on page 865

“Refresh Object configuration event” on page 866

“Configuration events” on page 837

“Configuration event usage” on page 865

“Refresh Object configuration event” on page 866

“Configuration events” on page 837

**Related reference:**

“Configuration event generation”



Create object (*WebSphere MQ V7.1 Reference*)



Change object (*WebSphere MQ V7.1 Reference*)



Delete object (*WebSphere MQ V7.1 Reference*)



Refresh object (*WebSphere MQ V7.1 Reference*)

“Event types” on page 832

“Configuration event generation”



Create object (*WebSphere MQ V7.1 Reference*)



Change object (*WebSphere MQ V7.1 Reference*)



Delete object (*WebSphere MQ V7.1 Reference*)



Refresh object (*WebSphere MQ V7.1 Reference*)

“Event types” on page 832

## Configuration event generation

Use this page to view the commands that cause configuration events to be generated and to understand the circumstances in which configuration events are not generated

A configuration event message is put to the configuration event queue when the CONFIGEV queue manager attribute is ENABLED and

- any of the following commands, or their PCF equivalent, are issued:
  - DELETE AUTHINFO
  - DELETE CFSTRUCT
  - DELETE CHANNEL
  - DELETE NAMELIST
  - DELETE PROCESS
  - DELETE QMODEL/QALIAS/QREMOTE
  - DELETE STGCLASS
  - DELETE TOPIC
  - REFRESH QMGR
- any of the following commands, or their PCF equivalent, are issued even if there is no change to the object:

- DEFINE/ALTER AUTHINFO
- DEFINE/ALTER CFSTRUCT
- DEFINE/ALTER CHANNEL
- DEFINE/ALTER NAMELIST
- DEFINE/ALTER PROCESS
- DEFINE/ALTER QMODEL/QALIAS/QREMOTE
- DEFINE/ALTER STGCLASS
- DEFINE/ALTER TOPIC
- DEFINE MAXSMSGS
- SET CHLAUTH
- ALTER QMGR, unless the CONFIGEV attribute is DISABLED and is not changed to ENABLED
- any of the following commands, or their PCF equivalent, are issued for a local queue that is not temporary dynamic, even if there is no change to the queue.
  - DELETE QLOCAL
  - DEFINE/ALTER QLOCAL
- an MQSET call is issued, other than for a temporary dynamic queue, even if there is no change to the object.

### When configuration events are not generated

Configuration events messages are not generated in the following circumstances:

- When a command or an MQSET call fails
- When a queue manager encounters an error trying to put a configuration event on the event queue, in which case the command or MQSET call completes, but no event message is generated
- For a temporary dynamic queue
- When internal changes are made to the TRIGGER queue attribute
- For the configuration event queue SYSTEM.ADMIN.CONFIG.EVENT, except by the REFRESH QMGR command
- For REFRESH/RESET CLUSTER and RESUME/SUSPEND QMGR commands that cause clustering changes
- When Creating or deleting a queue manager

#### Related concepts:

“Introduction to Programmable Command Formats” on page 6

“Configuration events” on page 863

#### Related reference:



The MQSC commands (*WebSphere MQ V7.1 Reference*)



MQSET – Set object attributes (*WebSphere MQ V7.1 Reference*)



MQSET - Set object attributes (*WebSphere MQ V7.1 Reference*)

### Configuration event usage

Use this page to view how you can use configuration events to obtain information about your system, and to understand the factors, such as CMDSCOPE, that can affect your use of configuration events.

You can use configuration events for the following purposes:

1. To produce and maintain a central configuration repository, from which reports can be produced and information about the structure of the system can be generated.

2. To generate an audit trail. For example, if an object is changed unexpectedly, information regarding who made the alteration and when it was done can be stored.

This can be particularly useful when command events are also enabled. If an MQSC or PCF command causes a configuration event and a command event to be generated, both event messages will share the same correlation identifier in their message descriptor.

For an MQSET call or any of the following commands:

- DEFINE object
- ALTER object
- DELETE object

if the queue manager attribute CONFIGEV is enabled, but the configuration event message cannot be put on the configuration event queue, for example the event queue has not been defined, the command or MQSET call is executed regardless.

## Effects of CMDSCOPE

For commands where CMDSCOPE is used, the configuration event message or messages will be generated on the queue manager or queue managers where the command is executed, not where the command is entered. However, all the origin and context information in the event data will relate to the original command as entered, even where the command using CMDSCOPE is one that has been generated by the source queue manager.

Where a queue sharing group includes queue managers that are not at the current version, events will be generated for any command that is executed by means of CMDSCOPE on a queue manager that is at the current version, but not on those that are at a previous version. This happens even if the queue manager where the command is entered is at the previous version, although in such a case no context information is included in the event data.

### Related concepts:

“Introduction to Programmable Command Formats” on page 6

“Configuration events” on page 863

### Related reference:



MQSET – Set object attributes (*WebSphere MQ V7.1 Reference*)



MQSET - Set object attributes (*WebSphere MQ V7.1 Reference*)

## Refresh Object configuration event

The Refresh Object configuration event is different from the other configuration events, because it occurs only when explicitly requested.

The create, change, and delete events are generated by an MQSET call or by a command to change an object but the refresh object event occurs only when explicitly requested by the MQSC command, REFRESH QMGR, or its PCF equivalent.

The REFRESH QMGR command is different from all the other commands that generate configuration events. All the other commands apply to a particular object and generate a single configuration event for that object. The REFRESH QMGR command can produce many configuration event messages potentially representing every object definition stored by a queue manager. One event message is generated for each object that is selected.

The REFRESH QMGR command uses a combination of three selection criteria to filter the number of objects involved:

- Object Name

- Object Type
- Refresh Interval

If you specify none of the selection criteria on the REFRESH QMGR command, the default values are used for each selection criteria and a refresh configuration event message is generated for every object definition stored by the queue manager. This might cause unacceptable processing times and event message generation. Consider specifying some selection criteria.

The REFRESH QMGR command that generates the refresh events can be used in the following situations:

- When configuration data is wanted about all or some of the objects in a system regardless of whether the objects have been recently manipulated, for example, when configuration events are first enabled. Consider using several commands, each with a different selection of objects, but such that all are included.
- If there has been an error in the SYSTEM.ADMIN.CONFIG.EVENT queue. In this circumstance, no configuration event messages are generated for Create, Change, or Delete events. When the error on the queue has been corrected, the Refresh Queue Manager command can be used to request the generation of event messages, which were lost while there was an error in the queue. In this situation consider setting the refresh interval to the time for which the queue was unavailable.

**Related concepts:**

“Configuration events” on page 863

**Related reference:**



REFRESH QMGR (*WebSphere MQ V7.1 Reference*)



Refresh Queue Manager (*WebSphere MQ V7.1 Reference*)

## Command events

Command events are notifications that an MQSC, or PCF command has run successfully.

The event data contains the following information:

**Origin information**

comprises the queue manager from where the command was issued, the ID of the user that issued the command, and how the command was issued, for example by a console command.

**Context information**

a replica of the context information in the message data from the command message. If a command is not entered using a message, context information is omitted.

Context information is included in the event data only when the command was entered as a message on the SYSTEM.COMMAND.INPUT queue.

**Command information**

the type of command that was issued.

**Command data**

- for PCF commands, a replica of the command data
- for MQSC commands, the command text

The command data format does not necessarily match the format of the original command. For example, on distributed platforms the command data format is always in PCF format, even if the original request was an MQSC command.

Every command event message that is generated is placed on the command event queue, SYSTEM.ADMIN.COMMAND.EVENT.

**Related concepts:**

“Command event generation”

“Command event usage” on page 869

“Command event generation”

“Command event usage” on page 869

**Related reference:**

Command (*WebSphere MQ V7.1 Reference*)

“Event types” on page 832



Command (*WebSphere MQ V7.1 Reference*)

“Event types” on page 832

**Command event generation**

Use this page to view the situations that cause command events to be generated and to understand the circumstances in which command events are not generated

A command event message is generated in the following situations:

- When the CMDEV queue manager attribute is specified as ENABLED and an MQSC or PCF command runs successfully.
- When the CMDEV queue manager attribute is specified as NODISPLAY and any command runs successfully, with the exception of DISPLAY commands (MQSC), and Inquire commands (PCF).
- When you run the MQSC command, ALTER QMGR, or the PCF command, Change Queue Manager, and the CMDEV queue manager attribute meets either of the following conditions:
  - CMDEV is not specified as DISABLED after the change
  - CMDEV was not specified as DISABLED before the change

If a command runs against the command event queue, SYSTEM.ADMIN.COMMAND.EVENT, a command event is generated if the queue still exists and it is not put-inhibited.

**When command events are not generated**

A command event message is not generated in the following circumstances:

- When a command fails
- When a queue manager encounters an error trying to put a command event on the event queue, in which case the command runs regardless, but no event message is generated
- For the MQSC command REFRESH QMGR TYPE (EARLY)
- For the MQSC command START QMGR MQSC
- For the MQSC command SUSPEND QMGR, if the parameter LOG is specified
- For the MQSC command RESUME QMGR, if the parameter LOG is specified

**Related concepts:**

“Command events” on page 867

**Related reference:**

REFRESH QMGR (*WebSphere MQ V7.1 Reference*)



START QMGR (*WebSphere MQ V7.1 Reference*)



SUSPEND QMGR (*WebSphere MQ V7.1 Reference*)



RESUME QMGR (*WebSphere MQ V7.1 Reference*)



SUSPEND QMGR, RESUME QMGR and clusters (*WebSphere MQ V7.1 Reference*)

**Command event usage**

Use this page to view how you can use command events to generate an audit trail of the commands that have run

For example, if an object is changed unexpectedly, information regarding who made the alteration and when it was done can be stored. This can be particularly useful when configuration events are also enabled. If an MQSC or PCF command causes a command event and a configuration event to be generated, both event messages will share the same correlation identifier in their message descriptor.

If a command event message is generated, but cannot be put on the command event queue, for example if the command event queue has not been defined, the command for which the command event was generated still runs regardless.

**Effects of CMDSCOPE**

For commands where CMDSCOPE is used, the command event message or messages will be generated on the queue manager or queue managers where the command runs, not where the command is entered. However, all the origin and context information in the event data will relate to the original command as entered, even where the command using CMDSCOPE is one that has been generated by the source queue manager.

**Related concepts:**

“Command events” on page 867

“Command event generation” on page 868

**Related reference:**

The MQSC commands (*WebSphere MQ V7.1 Reference*)



PCF commands and responses in groups (*WebSphere MQ V7.1 Reference*)

**Logger events**

Logger events are notifications that a queue manager has started writing to a new log extent or, on IBM i, a journal receiver. Logger event messages are not available with IBM WebSphere MQ for z/OS.

The event data contains the following information:

- The name of the current log extent.
- The name of the earliest log extent needed for restart recovery.
- The name of the earliest log extent needed for media recovery.
- The directory in which the log extents are located.

Every logger event message that is generated is placed on the logger event queue, SYSTEM.ADMIN.LOGGER.EVENT.

**Related concepts:**

“Logger event usage”

“Logger event usage”

**Related reference:**

“Logger event generation”

“Sample program to monitor the logger event queue” on page 871

 [Logger \(WebSphere MQ V7.1 Reference\)](#)

“Event types” on page 832

“Logger event generation”

“Sample program to monitor the logger event queue” on page 871

 [Logger \(WebSphere MQ V7.1 Reference\)](#)

“Event types” on page 832

## Logger event generation

Use this page to view the situations that cause logger events to be generated and to understand the circumstances in which logger events are not generated

A logger event message is generated in the following situations:

- When the LOGGEREV queue manager attribute is specified as ENABLED and the queue manager starts writing to a new log extent or, on IBM i, a journal receiver.
- When the LOGGEREV queue manager attribute is specified as ENABLED and the queue manager starts.
- When the LOGGEREV queue manager attribute is changed from DISABLED to ENABLED.

**Tip:** You can use the RESET QMGR MQSC command to request a queue manager to start writing to a new log extent.

## When logger events are not generated

A logger event message is not generated in the following circumstances:

- When a queue manager is configured to use circular logging.  
In this case, the LOGGEREV queue manager attribute is set as DISABLED and cannot be altered.
- When a queue manager encounters an error trying to put a logger event on the event queue, in which case the action that caused the event completes, but no event message is generated.

**Related concepts:**

“Logger events” on page 869

**Related reference:**

 [LoggerEvent \(MQLONG\) \(WebSphere MQ V7.1 Reference\)](#)

 [LoggerEvent \(10-digit signed integer\) \(WebSphere MQ V7.1 Reference\)](#)

 [RESET QMGR \(WebSphere MQ V7.1 Reference\)](#)

## Logger event usage

Use this page to view how you can use logger events to determine the log extents that are no longer required for queue manager restart, or media recovery.

You can archive superfluous log extents to a medium such as tape for disaster recovery before removing them from the active log directory. Regular removal of superfluous log extents keeps disk space usage to a minimum.

If the LOGGEREV queue manager attribute is enabled, but a logger event message cannot be put on the logger event queue, for example because the event queue has not been defined, the action that caused the event continues regardless.

**Related concepts:**

“Logger events” on page 869

**Related reference:**

 LoggerEvent (MQLONG) (*WebSphere MQ V7.1 Reference*)

 LoggerEvent (10-digit signed integer) (*WebSphere MQ V7.1 Reference*)

“Logger event generation” on page 870

**Sample program to monitor the logger event queue**

Use this page to view a sample C program that monitors the logger event queue for new event messages, reads those messages, and puts the contents of the message to stdout.

```

/*****
/*
/* Program name: AMQSLOG0.C
/*
/* Description: Sample C program to monitor the logger event queue and output*
/* a message to stdout when a logger event occurs
/*
/* <N_OCO_COPYRIGHT>
/* Licensed Materials - Property of IBM
/*
/* 63H9336
/* (c) Copyright IBM Corp. 2005, 2019 All Rights Reserved.
/*
/* US Government Users Restricted Rights - Use, duplication or
/* disclosure restricted by GSA ADP Schedule Contract with
/* IBM Corp.
/* <NOC_COPYRIGHT>
/*****
/*
/* Function: AMQSLOG is a sample program which monitors the logger event
/* queue for new event messages, reads those messages, and puts the contents
/* of the message to stdout.
/*
/*****
/*
/* AMQSLOG has 1 parameter - the queue manager name (optional, if not
/* specified then the default queue manager is implied)
/*
/*****

/*****
/* Includes
/*****
#include <stdlib.h>
#include <stdio.h>
#include <string.h>

#include <cmqc.h> /* MQI constants*/
#include <cmqfc.h> /* PCF constants*/

```

```

/*****
/* Constants
*****/

#define MAX_MESSAGE_LENGTH 8000

typedef struct _ParmTableEntry
{
 MQLONG ConstVal;
 PMQCHAR Desc;
} ParmTableEntry;

ParmTableEntry ParmTable[] =
{
 0, "",
 MQCA_Q_MGR_NAME, "Queue Manager Name",
 MQCMD_LOGGER_EVENT, "Logger Event Command",
 MQRC_LOGGER_STATUS, "Logger Status",
 MQCACF_CURRENT_LOG_EXTENT_NAME, "Current Log Extent",
 MQCACF_RESTART_LOG_EXTENT_NAME, "Restart Log Extent",
 MQCACF_MEDIA_LOG_EXTENT_NAME, "Media Log Extent",
 MQCACF_LOG_PATH, "Log Path"};

/*****
/* Function prototypes
*****/

static void ProcessPCF(MQHCONN hConn,
 MQHOBJ hEventQueue,
 PMQCHAR pBuffer);

static PMQCHAR ParmToString(MQLONG Parameter);

/*****
/* Function: main
*****/
int main(int argc, char * argv[])
{
 MQLONG CompCode;
 MQLONG Reason;
 MQHCONN hConn = MQHC_UNUSABLE_HCONN;
 MQOD ObjDesc = { MQOD_DEFAULT };
 MQCHAR QMName[MQ_Q_MGR_NAME_LENGTH+1] = "";
 MQCHAR LogEvQ[MQ_Q_NAME_LENGTH] = "SYSTEM.ADMIN.LOGGER.EVENT";
 MQHOBJ hEventQueue;
 PMQCHAR pBuffer = NULL;

 printf("\n/*****/\n");
 printf("/* Sample Logger Event Monitor start */\n");
 printf("/*****/\n");

 /*****/
 /* Parse any command line options
 */
 /*****/

 if (argc > 1)
 strncpy(QMName, argv[1], (size_t)MQ_Q_MGR_NAME_LENGTH);

 pBuffer = (char *)malloc(MAX_MESSAGE_LENGTH);
 if (!pBuffer)

```

```

{
 printf("Can't allocate %d bytes\n",MAX_MESSAGE_LENGTH);
 goto MOD_EXIT;
}

/*****
/* Connect to the specified (or default) queue manager */
*****/

MQCONN(QMName,
 &hConn,
 &CompCode,
 &Reason);

if (Reason != MQCC_OK)
{
 printf("Error in call to MQCONN, Reason %d, CompCode %d\n", Reason,
 CompCode);
 goto MOD_EXIT;
}

/* Open the logger event queue for input */

strncpy(ObjDesc.ObjectQMgrName,QMName, MQ_Q_MGR_NAME_LENGTH);
strncpy(ObjDesc.ObjectName, LogEvQ, MQ_Q_NAME_LENGTH);

MQOPEN(hConn,
 &ObjDesc,
 MQOO_INPUT_EXCLUSIVE,
 &hEventQueue,
 &CompCode,
 &Reason);
if (Reason)
{
 printf("MQOPEN failed for queue manager %.48s Queue %.48s Reason: %d\n",
 ObjDesc.ObjectQMgrName,
 ObjDesc.ObjectName,
 Reason);

 goto MOD_EXIT;
}
else
{
 ProcessPCF(hConn, hEventQueue, pBuffer);
}

MOD_EXIT:

if (pBuffer != NULL) {
 free(pBuffer);
}

/*****
/* Disconnect */
*****/
if (hConn != MQHC_UNUSABLE_HCONN) {
 MQDISC(&hConn, &CompCode, &Reason);
}

return 0;
}

```

```

/*****/
/* Function: ProcessPCF */
/*****/
/*
/* Input Parameters: Handle to queue manager connection */
/* Handle to the opened logger event queue object */
/* Pointer to a memory buffer to store the incoming PCF msg*/
/*
/* Output Parameters: None */
/*
/* Logic: Wait for messages to appear on the logger event queue and display */
/* their contents. */
/*
/*****/

static void ProcessPCF(MQHCONN hConn,
 MQHOBJ hEventQueue,
 PMQCHAR pBuffer)
{
 MQCFH * pCfh;
 MQCFST * pCfst;
 MQGMO Gmo = { MQGMO_DEFAULT };
 MQMD Mqmd = { MQMD_DEFAULT };
 PMQCHAR pPCFCmd;
 MQLONG Reason = 0;
 MQLONG CompCode;
 MQLONG MsgLen;
 PMQCHAR Parm = NULL;

 /* Set timeout value */
 Gmo.Options |= MQGMO_WAIT;
 Gmo.Options |= MQGMO_CONVERT;
 Gmo.WaitInterval = MQWI_UNLIMITED;
/*****/
/* Process response Queue */
/*****/
while (Reason == MQCC_OK)
{
 memcpy(&Mqmd.MsgId; , MQMI_NONE, sizeof(Mqmd.MsgId));
 memset(&Mqmd.CorrelId, 0, sizeof(Mqmd.CorrelId));

 MQGET(hConn,
 hEventQueue,
 &Mqmd,
 &Gmo,
 MAX_MESSAGE_LENGTH,
 pBuffer,
 &MsgLen,
 &CompCode,
 &Reason);
 if (Reason != MQCC_OK)
 {
 switch(Reason)
 {
 case MQRC_NO_MSG_AVAILABLE:
 printf("Timed out");
 break;

 default:
 printf("MQGET failed RC(%d)\n", Reason);
 }
 }
}
}

```

```

 break;
 }
 goto MOD_EXIT;
}

/*****
/* Only expect PCF event messages on this queue */
/*****
if (memcmp(Mqmd.Format, MQFMT_EVENT, sizeof(Mqmd.Format)))
{
 printf("Unexpected message format '%8.8s' received\n",Mqmd.Format);
 continue;
}

/*****
/* Build the output by parsing the received PCF message, first the */
/* header, then each of the parameters */
/*****

pCfh = (MQCFH *)pBuffer;

if (pCfh -> Reason)
{
 printf("-----\n");
 printf("Event Message Received\n");

 Parm = ParmToString(pCfh->Command);
 if (Parm != NULL) {
 printf("Command :%s \n",Parm);
 }
 else
 {
 printf("Command :%d \n",pCfh->Command);
 }

 printf("CompCode :%d\n" ,pCfh->CompCode);

 Parm = ParmToString(pCfh->Reason);
 if (Parm != NULL) {
 printf("Reason :%s \n",Parm);
 }
 else
 {
 printf("Reason :%d \n",pCfh->Reason);
 }
}

pPCFCmd = (char *) (pCfh+1);
printf("-----\n");
while(pCfh -> ParameterCount--)
{
 pCfst = (MQCFST *) pPCFCmd;
 switch(pCfst -> Type)
 {
 case MQCFT_STRING:
 Parm = ParmToString(pCfst -> Parameter);
 if (Parm != NULL) {
 printf("%-32s",Parm);
 }
 }
}

```

```

 else
 {
 printf("%-32d",pCfst -> Parameter);
 }

 fwrite(pCfst -> String, pCfst -> StringLength, 1, stdout);
 pPCFCmd += pCfst -> StruLength;
 break;
 default:
 printf("Unrecognised datatype %d returned\n",pCfst->Type);
 goto MOD_EXIT;
 }
 putchar('\n');
}
printf("-----\n");
}
MOD_EXIT:

return;
}

/*****
/* Function: ParmToString
/*****
/*
/* Input Parameters: Parameter for which to get string description
/*
/* Output Parameters: None
/*
/* Logic: Takes a parameter as input and returns a pointer to a string
/* description for that parameter, or NULL if the parameter does not
/* have an associated string description
/*****

static PMQCHAR ParmToString(MQLONG Parameter){
 long i;
 for (i=0 ; i< sizeof(ParmTable)/sizeof(ParmTableEntry); i++)
 {
 if (ParmTable[i].ConstVal == Parameter ParmTable[i].Desc)
 return ParmTable[i].Desc;
 }
 return NULL;
}

```

## Sample output

This application produces the following form of output:

```

/*****
/* Sample Logger Event Monitor start */
/*****

Event Message Received
Command :Logger Event Command
CompCode :0
Reason :Logger Status

Queue Manager Name CSIM

```

```

Current Log Extent AMQA000001
Restart Log Extent AMQA000001
Media Log Extent AMQA000001
Log Path QMCSIM

```

**Related concepts:**

“Logger event usage” on page 870

“Command event usage” on page 869

**Related reference:**

“Logger event generation” on page 870

## Sample program to monitor instrumentation events

Use this page to view a sample C program for monitoring instrumentation events

This sample program is not part of any IBM WebSphere MQ product and is therefore not supplied as an actual physical item. The example is incomplete in that it does not enumerate all the possible outcomes of specified actions. However, you can use this sample as a basis for your own programs that use events, in particular, the PCF formats used in event messages. However, you need to modify this program before running it on your own systems.

```

/*****/
/*
/* Program name: EVMON
/*
/* Description: C program that acts as an event monitor
/*
/*
/*
/*****/
/*
/* Function:
/*
/*
/* EVMON is a C program that acts as an event monitor - reads an
/* event queue and tells you if anything appears on it
/*
/*
/* Its first parameter is the queue manager name, the second is
/* the event queue name. If these are not supplied it uses the
/* defaults.
/*
/*
/*****/
#include <time.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#ifndef min
#define min(a,b) ((a) < (b)) ? (a) : (b)
#endif

/*****/
/* includes for MQI
/*****/
#include <cmqc.h>
#include <cmqcfst.h>
void printfmqcfst(MQCFST* pmqcfst);
void printfmqcfst(MQCFIN* pmqcfst);
void printreas(MQLONG reason);

#define PRINTREAS(param) \

```

```

 case param:
 printf("Reason = %s\n",#param);
 break;

/*****
/* global variable */
/*****
MQCFH *evtmsg; /* evtmsg message buffer */

int main(int argc, char **argv)
{
 /*****
 /* declare variables */
 /*****
 int i; /* auxiliary counter */
 /*****
 /* Declare MQI structures needed */
 /*****
 MQOD od = {MQOD_DEFAULT}; /* Object Descriptor */
 MQMD md = {MQMD_DEFAULT}; /* Message Descriptor */
 MQGMO gmo = {MQGMO_DEFAULT}; /* get message options */
 /*****
 /* note, uses defaults where it can */
 /*****

 MQHCONN Hcon; /* connection handle */
 MQHOBJ Hobj; /* object handle */
 MQLONG O_options; /* MQOPEN options */
 MQLONG C_options; /* MQCLOSE options */
 MQLONG CompCode; /* completion code */
 MQLONG OpenCode; /* MQOPEN completion code */
 MQLONG Reason; /* reason code */
 MQLONG CReason; /* reason code for MQCONN */
 MQLONG buflen; /* buffer length */
 MQLONG evtmsglen; /* message length received */
 MQCHAR command[1100]; /* call command string ... */
 MQCHAR p1[600]; /* ApplId insert */
 MQCHAR p2[900]; /* evtmsg insert */
 MQCHAR p3[600]; /* Environment insert */
 MQLONG mytype; /* saved application type */
 char QMName[50]; /* queue manager name */
 MQCFST *paras; /* the parameters */
 int counter; /* loop counter */
 time_t ltime;

 /*****
 /* Connect to queue manager */
 /*****
 /*****
 QMName[0] = 0; /* default queue manager */
 if (argc > 1)
 strcpy(QMName, argv[1]);
 MQCONN(QMName,
 &Hcon, /* connection handle */
 &CompCode, /* completion code */
 &CReason); /* reason code */

```

```

/*****/
/* Initialize object descriptor for subject queue */
/*****/
strcpy(od.ObjectName, "SYSTEM.ADMIN.QMGR.EVENT");
if (argc > 2)
 strcpy(od.ObjectName, argv[2]);

/*****/
/* Open the event queue for input; exclusive or shared. Use of */
/* the queue is controlled by the queue definition here */
/*****/

O_options = MQOO_INPUT_AS_Q_DEF /* open queue for input */
 + MQOO_FAIL_IF QUIESCING /* but not if qmgr stopping */
 + MQOO_BROWSE;
MQOPEN(Hcon, /* connection handle */
 &od, /* object descriptor for queue*/
 O_options, /* open options */
 &Hobj, /* object handle */
 &CompCode, /* completion code */
 &Reason); /* reason code */

/*****/
/* Get messages from the message queue */
/*****/
while (CompCode != MQCC_FAILED)
{
 /*****/
 /* I don't know how big this message is so just get the */
 /* descriptor first */
 /*****/
 gmo.Options = MQGMO_WAIT + MQGMO_LOCK
 + MQGMO_BROWSE_FIRST + MQGMO_ACCEPT_TRUNCATED_MSG;
 /* wait for new messages */
 gmo.WaitInterval = MQWI_UNLIMITED; /* no time limit */
 buflen = 0; /* amount of message to get */

 /*****/
 /* clear selectors to get messages in sequence */
 /*****/
 memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
 memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));

 /*****/
 /* wait for event message */
 /*****/
 printf("...>\n");
 MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 buflen, /* buffer length */
 evtmsg, /* evtmsg message buffer */
 &evtmglen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */

 /*****/
}
/*****/

```

```

/* report reason, if any */
/*****
if (Reason != MQRC_NONE && Reason != MQRC_TRUNCATED_MSG_ACCEPTED)
{
 printf("MQGET ==> %ld\n", Reason);
}
else
{
 gmo.Options = MQGMO_NO_WAIT + MQGMO_MSG_UNDER_CURSOR;
 buflen = evtmsglen; /* amount of message to get */
 evtmsg = malloc(buflen);
 if (evtmsg != NULL)
 {
 /*****
 /* clear selectors to get messages in sequence */
 /*****
 memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
 memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));

 /*****
 /* get the event message */
 /*****
 printf("...>\n");
 MQGET(Hcon, /* connection handle */
 Hobj, /* object handle */
 &md, /* message descriptor */
 &gmo, /* get message options */
 buflen, /* buffer length */
 evtmsg, /* evtmsg message buffer */
 &evtmsglen, /* message length */
 &CompCode, /* completion code */
 &Reason); /* reason code */

 /*****
 /* report reason, if any */
 /*****
 if (Reason != MQRC_NONE)
 {
 printf("MQGET ==> %ld\n", Reason);
 }
 }
 else
 {
 CompCode = MQCC_FAILED;
 }
}

/*****
/* . . . process each message received */
/*****

if (CompCode != MQCC_FAILED)
{
 /*****
 /* announce a message */
 /*****
 printf("\a\a\a\a\a\a");
 time(<time);
 printf(ctime(<time));
}

```

```

if (evtmsglen != buflen)
 printf("DataLength = %ld?\n", evtmsglen);
else
{
 /*****
 /* right let's look at the data
 /*****
 if (evtmsg->Type != MQCFT_EVENT)
 {
 printf("Something's wrong this isn't an event message,"
 " its type is %ld\n",evtmsg->Type);
 }
 else
 {
 if (evtmsg->Command == MQCMD_Q_MGR_EVENT)
 {
 printf("Queue Manager event: ");
 }
 else
 if (evtmsg->Command == MQCMD_CHANNEL_EVENT)
 {
 printf("Channel event: ");
 }
 else
 {
 printf("Unknown Event message, %ld.",
 evtmsg->Command);
 }

 if (evtmsg->CompCode == MQCC_OK)
 printf("CompCode(OK)\n");
 else if (evtmsg->CompCode == MQCC_WARNING)
 printf("CompCode(WARNING)\n");
 else if (evtmsg->CompCode == MQCC_FAILED)
 printf("CompCode(FAILED)\n");
 else
 printf("* CompCode wrong * (%ld)\n",
 evtmsg->CompCode);

 if (evtmsg->StrucLength != MQCFH_STRUC_LENGTH)
 {
 printf("it's the wrong length, %ld\n",evtmsg->StrucLength);
 }

 if (evtmsg->Version != MQCFH_VERSION_1)
 {
 printf("it's the wrong version, %ld\n",evtmsg->Version);
 }

 if (evtmsg->MsgSeqNumber != 1)
 {
 printf("it's the wrong sequence number, %ld\n",
 evtmsg->MsgSeqNumber);
 }
 }
 }
}

```

```

if (evtmsg->Control != MQCFC_LAST)
{
 printf("it's the wrong control option, %ld\n",
 evtmsg->Control);
}

printreas(evtmsg->Reason);
printf("parameter count is %ld\n", evtmsg->ParameterCount);
/*****
/* get a pointer to the start of the parameters */
*****/

paras = (MQCFST *) (evtmsg + 1);
counter = 1;
while (counter <= evtmsg->ParameterCount)
{
 switch (paras->Type)
 {
 case MQCFT_STRING:
 printfmqfst(paras);
 paras = (MQCFST *) ((char *) paras
 + paras->StrucLength);
 break;
 case MQCFT_INTEGER:
 printfmqcfin((MQCFIN*) paras);
 paras = (MQCFST *) ((char *) paras
 + paras->StrucLength);
 break;
 default:
 printf("unknown parameter type, %ld\n",
 paras->Type);
 counter = evtmsg->ParameterCount;
 break;
 }
 counter++;
}
}
} /* end evtmsg action */
free(evtmsg);
evtmsg = NULL;
} /* end process for successful GET */
} /* end message processing loop */

/*****
/* close the event queue - if it was opened */
*****/
if (OpenCode != MQCC_FAILED)
{
 C_options = 0; /* no close options */
 MQCLOSE(Hcon, /* connection handle */
 &Hobj, /* object handle */
 C_options,
 &CompCode, /* completion code */
 &Reason); /* reason code */
/*****
/* Disconnect from queue manager (unless previously connected) */
*****/
if (CReason != MQRC_ALREADY_CONNECTED)
{

```

```

 MQDISC(&Hcon, /* connection handle */
 &CompCode, /* completion code */
 &Reason); /* reason code */

/*****
/*
/* END OF EVMON
/*
/*****
}

#define PRINTPARAM(param) \
 case param: \
 { \
 char *p = #param; \
 strncpy(thestring,pmqcfst->String,min(sizeof(thestring), \
 pmqcfst->StringLength)); \
 printf("%s %s\n",p,thestring); \
 } \
 break;

#define PRINTAT(param) \
 case param: \
 printf("MQIA_APPL_TYPE = %s\n",#param); \
 break;

void printfmqcfst(MQCFST* pmqcfst)
{
 char thestring[100];

 switch (pmqcfst->Parameter)
 {
 PRINTPARAM(MQCA_BASE_Q_NAME)
 PRINTPARAM(MQCA_PROCESS_NAME)
 PRINTPARAM(MQCA_Q_MGR_NAME)
 PRINTPARAM(MQCA_Q_NAME)
 PRINTPARAM(MQCA_XMIT_Q_NAME)
 PRINTPARAM(MQCACF_APPL_NAME)

 :
 default:
 printf("Invalid parameter, %ld\n",pmqcfst->Parameter);
 break;
 }
}

void printfmqcfin(MQCFIN* pmqcfst)
{
 switch (pmqcfst->Parameter)
 {
 case MQIA_APPL_TYPE:
 switch (pmqcfst->Value)
 {
 PRINTAT(MQAT_UNKNOWN)
 PRINTAT(MQAT_OS2)
 }
 }
 }
}

```

```

 PRINTAT(MQAT_DOS)
 PRINTAT(MQAT_UNIX)
 PRINTAT(MQAT_QMGR)
 PRINTAT(MQAT_OS400)
 PRINTAT(MQAT_WINDOWS)
 PRINTAT(MQAT_CICS_VSE)
 PRINTAT(MQAT_VMS)
 PRINTAT(MQAT_GUARDIAN)
 PRINTAT(MQAT_VOS)
}
break;
case MQIA_Q_TYPE:
 if (pmqcfst->Value == MQQT_ALIAS)
 {
 printf("MQIA_Q_TYPE is MQQT_ALIAS\n");
 }
 else
:
{
 if (pmqcfst->Value == MQQT_REMOTE)
 {
 printf("MQIA_Q_TYPE is MQQT_REMOTE\n");
 if (evtmsg->Reason == MQRC_ALIAS_BASE_Q_TYPE_ERROR)
 {
 printf("but remote is not valid here\n");
 }
 }
 else
 {
 printf("MQIA_Q_TYPE is wrong, %ld\n",pmqcfst->Value);
 }
}
break;

 case MQIACF_REASON_QUALIFIER:
printf("MQIACF_REASON_QUALIFIER %ld\n",pmqcfst->Value);
break;

case MQIACF_ERROR_IDENTIFIER:
printf("MQIACF_ERROR_IDENTIFIER %ld (X'%1X')\n",
 pmqcfst->Value,pmqcfst->Value);
break;

case MQIACF_AUX_ERROR_DATA_INT_1:
printf("MQIACF_AUX_ERROR_DATA_INT_1 %ld (X'%1X')\n",
 pmqcfst->Value,pmqcfst->Value);
break;

case MQIACF_AUX_ERROR_DATA_INT_2:
printf("MQIACF_AUX_ERROR_DATA_INT_2 %ld (X'%1X')\n",
 pmqcfst->Value,pmqcfst->Value);
break;
:
default :
 printf("Invalid parameter, %ld\n",pmqcfst->Parameter);
 break;
}
}

```

```

 void printreas(MQLONG reason)
 {
 switch (reason)
 {
 PRINTREAS(MQRCCF_CFH_TYPE_ERROR)
 PRINTREAS(MQRCCF_CFH_LENGTH_ERROR)
 PRINTREAS(MQRCCF_CFH_VERSION_ERROR)
 PRINTREAS(MQRCCF_CFH_MSG_SEQ_NUMBER_ERR)

 :
 PRINTREAS(MQRC_NO_MSG_LOCKED)
 PRINTREAS(MQRC_CONNECTION_NOT_AUTHORIZED)
 PRINTREAS(MQRC_MSG_TOO_BIG_FOR_CHANNEL)
 PRINTREAS(MQRC_CALL_IN_PROGRESS)
 default:
 printf("It's an unknown reason, %ld\n",
 reason);
 break;
 }
 }

```

#### Related concepts:

“Instrumentation events” on page 830

“Event monitoring” on page 829

“Instrumentation events” on page 830

“Event monitoring” on page 829

#### Related reference:



C programming (*WebSphere MQ V7.1 Reference*)

“Sample program to monitor the logger event queue” on page 871



C programming (*WebSphere MQ V7.1 Reference*)

“Sample program to monitor the logger event queue” on page 871

---

## Message monitoring

Message monitoring is the process of identifying the route a message has taken through a queue manager network. By identifying the types of activities, and the sequence of activities performed on behalf of a message, the message route can be determined.

As a message passes through a queue manager network, various processes perform activities on behalf of the message. Use one of the following techniques to determine a message route:

- The IBM WebSphere MQ display route application (dspmqrte)
- Activity recording
- Trace-route messaging

These techniques all generate special messages that contain information about the activities performed on the message as it passed through a queue manager network. Use the information returned in these special messages to achieve the following objectives:

- Record message activity.
- Determine the last known location of a message.
- Detect routing problems in your queue manager network.

- Assist in determining the causes of routing problems in your queue manager network.
- Confirm that your queue manager network is running correctly.
- Familiarize yourself with the running of your queue manager network.
- Trace published messages.

**Related concepts:**

- “Activities and operations”
- “Message route techniques” on page 888
- “Activity recording” on page 889
- “Trace-route messaging” on page 894
- “WebSphere MQ display route application” on page 908
- “Activity report reference” on page 926
- “Trace-route message reference” on page 952
- “Trace-route reply message reference” on page 963
- “Activities and operations”
- “Message route techniques” on page 888
- “Activity recording” on page 889
- “Trace-route messaging” on page 894
- “WebSphere MQ display route application” on page 908
- “Activity report reference” on page 926
- “Trace-route message reference” on page 952
- “Trace-route reply message reference” on page 963

**Related reference:**

-  Types of message (*WebSphere MQ V7.1 Programming Guide*)
-  Types of message (*WebSphere MQ V7.1 Programming Guide*)

## Activities and operations

Activities are discrete actions that an application performs on behalf of a message. Activities consist of operations, which are single pieces of work that an application performs.

The following actions are examples of activities:

- A message channel agent (MCA) sends a message from a transmission queue down a channel
- An MCA receives a message from a channel and puts it on its target queue
- An application getting a message from a queue, and putting a reply message in response.
- The WebSphere MQ publish/subscribe engine processes a message.

Activities consist of one or more *operations*. Operations are single pieces of work that an application performs. For example, the activity of an MCA sending a message from a transmission queue down a channel consists of the following operations:

1. Getting a message from a transmission queue (a *Get* operation).
2. Sending the message down a channel (a *Send* operation).

In a publish/subscribe network, the activity of the WebSphere MQ publish/subscribe engine processing a message can consist of the following multiple operations:

1. Putting a message to a topic string (a *Put* operation).
2. Zero or more operations for each of the subscribers that are considered for receipt of the message (a *Publish* operation, a *Discarded Publish* operation or an *Excluded Publish* operation).

## Information from activities

You can identify the sequence of activities performed on a message by recording information as the message is routed through a queue manager network. You can determine the route of a message through the queue manager network from the sequence of activities performed on the message, and can obtain the following information:

### The last known location of a message

If a message does not reach its intended destination, you can determine the last known location of the message from a complete or partial message route.

### Configuration issues with a queue manager network

When studying the route of a message through a queue manager network, you might see that the message has not gone where expected. There are many reasons why this can occur, for example, if a channel is inactive, the message might take an alternative route.

For a publish/subscribe application, you can also determine the route of a message being published to a topic and any messages that flow in a queue manager network as a result of being published to subscribers. In such situations, a system administrator can determine whether there are any problems in the queue manager network, and if appropriate, correct them.

## Message routes

Depending on your reason for determining a message route, you can use the following general approaches:

### Using activity information recorded for a trace-route message

Trace-route messages record activity information for a specific purpose. You can use them to determine configuration issues with a queue manager network, or to determine the last known location of a message. If a trace-route message is generated to determine the last known location of a message that did not reach its intended destination, it can mimic the original message. This gives the trace-route message the greatest chance of following the route taken by the original message.

The WebSphere MQ display route application can generate trace-route messages.

### Using activity information recorded for the original message

You can enable any message for activity recording and have activity information recorded on its behalf. If a message does not reach its intended destination, you can use the recorded activity information to determine the last known location of the message. By using activity information from the original message, the most accurate possible message route can be determined, leading to the last known location. To use this approach, the original message must be enabled for activity recording.

**Warning:** Avoid enabling all messages in a queue manager network for activity recording. Messages enabled for activity recording can have many activity reports generated on their behalf. If every message in a queue manager network is enabled for activity recording, the queue manager network traffic can increase to an unacceptable level.

## Related concepts:

“Message monitoring” on page 885

“Message route techniques”

 Writing your own message channel agents (*WebSphere MQ V7.1 Installing Guide*)

“Trace-route messaging” on page 894

## Message route techniques

Activity recording and trace-route messaging are techniques that allow you to record activity information for a message as it is routed through a queue manager network.

### Activity recording

If a message has the appropriate report option specified, it requests that applications generate *activity reports* as it is routed through a queue manager network. When an application performs an activity on behalf of a message, an activity report can be generated, and delivered to an appropriate location. An activity report contains information about the activity that was performed on the message.

The activity information collected using activity reports must be arranged in order before a message route can be determined.

### Trace-route messaging

*Trace-route messaging* is a technique that involves sending a *trace-route message* into a queue manager network. When an application performs an activity on behalf of the trace-route message, activity information can be accumulated in the message data of the trace-route message, or activity reports can be generated. If activity information is accumulated in the message data of the trace-route message, when it reaches its target queue a trace-route reply message containing all the information from the trace-route message can be generated and delivered to an appropriate location.

Because a trace-route message is dedicated to recording the sequence of activities performed on its behalf, there are more processing options available compared with normal messages that request activity reports.

## Comparison of activity recording and trace-route messaging

Both activity recording and trace-route messaging can provide activity information to determine the route a message has taken through a queue manager network. Both methods have their own advantages.

| Benefit                                                                                 | Activity recording | Trace-route messaging |
|-----------------------------------------------------------------------------------------|--------------------|-----------------------|
| Can determine the last known location of a message                                      | Yes                | Yes                   |
| Can determine configuration issues with a queue manager network                         | Yes                | Yes                   |
| Can be requested by any message<br>(is not restricted to use with trace-route messages) | Yes                | No                    |
| Message data is left unmodified                                                         | Yes                | No                    |
| Message processed normally                                                              | Yes                | No                    |
| Activity information can be accumulated in the message data                             | No                 | Yes                   |
| Optional message delivery to target queue                                               | No                 | Yes                   |
| If a message is caught in an infinite loop, it can be detected and dealt with           | No                 | Yes                   |
| Activity information can be put in order reliably                                       | No                 | Yes                   |

| Benefit                                                  | Activity recording | Trace-route messaging |
|----------------------------------------------------------|--------------------|-----------------------|
| Application provided to display the activity information | No                 | Yes                   |

## Message route completeness

In some cases it is not possible to identify the full sequence of activities performed on behalf of a message, so only a partial message route can be determined. The completeness of a message route is directly influenced by the queue manager network that the messages are routed through. The completeness of a message route depends on the level of the queue managers in the queue manager network, as follows:

### Queue managers at WebSphere MQ Version 6.0 and subsequent releases

MCAs and user-written applications connected to queue managers at WebSphere MQ Version 6.0 or subsequent releases can record information related to the activities performed on behalf of a message. The recording of activity information is controlled by the queue manager attributes `ACTIVREC` and `ROUTEREC`. If a queue manager network consists of queue managers at WebSphere MQ Version 6.0 or subsequent releases only, complete message routes can be determined.

### WebSphere MQ queue managers before Version 6.0

Applications connected to WebSphere MQ queue managers before Version 6.0 **do not** record the activities that they have performed on behalf of a message. If a queue manager network contains any WebSphere MQ queue manager prior to Version 6.0, only a partial message route can be determined.

## How activity information is stored

WebSphere MQ stores activity information in activity reports, trace-route messages, or trace-route reply messages. In each case the information is stored in a structure called the *Activity* PCF group. A trace-route message or trace-route reply message can contain many Activity PCF groups, depending on the number of activities performed on the message. Activity reports contain one Activity PCF group because a separate activity report is generated for every recorded activity.

With trace-route messaging, additional information can be recorded. This additional information is stored in a structure called the *TraceRoute* PCF group. The TraceRoute PCF group contains a number of PCF structures that are used to store additional activity information, and to specify options that determine how the trace-route message is handled as it is routed through a queue manager network.

### Related concepts:

“Activity recording”

“Trace-route messaging” on page 894

### Related reference:

“The TraceRoute PCF group” on page 900

“Activity report message data” on page 935

## Activity recording

Activity recording is a technique for determining the routes that messages take through a queue manager network. To determine the route that a message has taken, the activities performed on behalf of the message are recorded.

When using activity recording, each activity performed on behalf of a message can be recorded in an activity report. An activity report is a type of report message. Each activity report contains information about the application that performed the activity on behalf of the message, when the activity took place, and information about the operations that were performed as part of the activity. Activity reports are

typically delivered to a reply-to queue where they are collected together. By studying the activity reports related to a message, you can determine the route that the message took through the queue manager network.

## Activity report usage

When messages are routed through a queue manager network, activity reports can be generated. You can use activity report information in the following ways:

### Determine the last known location of a message

If a message that is enabled for activity recording does not reach its intended destination, activity reports generated for the message as it was routed through a queue manager network can be studied to determine the last known location of the message.

### Determine configuration issues with a queue manager network

A number of messages enabled for activity recording can be sent into a queue manager network. By studying the activity reports related to each message it can become apparent that they have not taken the expected route. There are many reasons why this can occur, for example, a channel could have stopped, forcing the message to take an alternative route. In these situations, a system administrator can determine whether there are any problems in the queue manager network, and if there are, correct them.

**Note:** You can use activity recording in conjunction with trace-route messages by using the WebSphere MQ display route application.

## Activity report format

Activity reports are PCF messages generated by applications that have performed an activity on behalf of a message. Activity reports are standard WebSphere MQ report messages containing a message descriptor and message data, as follows:

### The message descriptor

- An MQMD structure

### Message data

- An embedded PCF header (MQEPH)
- Activity report message data

Activity report message data consists of the *Activity* PCF group, and if generated for a trace-route message, the *TraceRoute* PCF group.

### Related concepts:

“Circumstances where activity information is not acquired” on page 894

### Related tasks:

“Controlling activity recording” on page 891

“Setting up a common queue for activity reports” on page 892

“Determining message route information” on page 892

“Retrieving further activity reports” on page 893

### Related reference:

 MQMD – Message descriptor (*WebSphere MQ V7.1 Reference*)

 MQEPH – Embedded PCF header (*WebSphere MQ V7.1 Reference*)

## Controlling activity recording

Enable activity recording at the queue manager level. To enable an entire queue manager network, individually enable every queue manager in the network for activity recording. If you enable more queue managers, more activity reports are generated.

### About this task

To generate activity reports for a message as it is routed through a queue manager: define the message to request activity reports; enable the queue manager for activity recording; and ensure that applications performing activities on the message are capable of generating activity reports.

If you do *not* want activity reports to be generated for a message as it is routed through a queue manager, *disable* the queue manager for activity recording.

### Procedure

1. Request activity reports for a message
  - a. In the message descriptor of the message, specify MQRO\_ACTIVITY in the *Report* field.
  - b. In the message descriptor of the message, specify the name of a reply-to queue in the *ReplyToQ* field.

**Warning:** Avoid enabling all messages in a queue manager network for activity recording. Messages enabled for activity recording can have many activity reports generated on their behalf. If every message in a queue manager network is enabled for activity recording, the queue manager network traffic can increase to an unacceptable level.

2. Enable or disable the queue manager for activity recording. Use the MQSC command ALTER QMGR, specifying the parameter ACTIVREC, to change the value of the queue manager attribute. The value can be:

**MSG** The queue manager is enabled for activity recording. Any activity reports generated are delivered to the reply-to queue specified in the message descriptor of the message. This is the default value.

#### QUEUE

The queue manager is enabled for activity recording. Any activity reports generated are delivered to the local system queue SYSTEM.ADMIN.ACTIVITY.QUEUE. The system queue can also be used to forward activity reports to a common queue.

#### DISABLED

The queue manager is disabled for activity recording. No activity reports are generated while in the scope of this queue manager.

For example, to enable a queue manager for activity recording and specify that any activity reports generated are delivered to the local system queue SYSTEM.ADMIN.ACTIVITY.QUEUE, use the following MQSC command:

```
ALTER QMGR ACTIVREC(Queue)
```

**Remember:** When you modify the *ACTIVREC* queue manager attribute, a running MCA does not detect the change until the channel is restarted.

3. Ensure that your application uses the same algorithm as MCAs use to determine whether to generate an activity report for a message:
  - a. Verify that the message has requested activity reports to be generated
  - b. Verify that the queue manager where the message currently resides is enabled for activity recording
  - c. Put the activity report on the queue determined by the *ACTIVREC* queue manager attribute

## Setting up a common queue for activity reports

To determine the locations of the activity reports related to a specific message when the reports are delivered to the local system queue, it is more efficient to use a common queue on a single node

### Before you begin

Set the `ACTIVREC` parameter to enable the queue manager for activity recording and to specify that any activity reports generated are delivered to the local system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE`.

### About this task

If a number of queue managers in a queue manager network are set to deliver activity reports to the local system queue, it can be time consuming to determine the locations of the activity reports related to a specific message. Alternatively, use a single node, which is a queue manager that hosts a common queue. All the queue managers in a queue manager network can deliver activity reports to this common queue. The benefit of using a common queue is that queue managers do not have to deliver activity reports to the reply-to queue specified in a message and, when determining the locations of the activity reports related to a message, you query one queue only.

To set up a common queue, perform the following steps:

### Procedure

1. Select or define a queue manager as the single node
2. On the single node, select or define a queue for use as the common queue
3. On all queue managers where activity reports are to be delivered to the common queue, redefine the local system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE` as a remote queue definition:
  - a. Specify the name of the single node as the remote queue manager name
  - b. Specify the name of the common queue as the remote queue name

## Determining message route information

To determine a message route, obtain the information from the activity reports collected. Determine whether enough activity reports are on the reply-to queue to enable you to determine the required information and arrange the activity reports in order.

### About this task

The order that activity reports are put on the reply-to queue does not necessarily correlate to the order in which the activities were performed. You must order activity reports manually, unless they are generated for a trace-route message, in which case you can use the WebSphere MQ display route application to order the activity reports.

Determine whether enough activity reports are on the reply-to queue for you to obtain the necessary information:

### Procedure

1. Identify all related activity reports on the reply-to queue by comparing identifiers of the activity reports and the original message. Ensure you set the report option of the original message such that the activity reports can be correlated with the original message.
2. Order the identified activity reports from the reply-to queue. You can use the following parameters from the activity report:

*OperationType*

The types of operations performed might enable you to determine the activity report that was generated directly before, or after, the current activity report.

For example, an activity report details that an MCA sent a message from a transmission queue down a channel. The last operation detailed in the activity report has an *OperationType* of **send** and details that the message was sent using the channel, CH1, to the destination queue manager, QM1. This means that the next activity performed on the message will have occurred on queue manager, QM1, and that it will have begun with a **receive** operation from channel, CH1. By using this information you can identify the next activity report, providing it exists and has been acquired.

#### *OperationDate* and *OperationTime*

You can determine the general order of the activities from the dates and times of the operations in each activity report.

**Warning:** Unless every queue manager in the queue manager network has their system clocks synchronized, ordering by date and time does not guarantee that the activity reports are in the correct sequence. You must establish the order manually.

The order of the activity reports represents the route, or partial route, that the message took through the queue manager network.

3. Obtain the information you need from the activity information in the ordered activity reports. If you have insufficient information about the message, you might be able to acquire further activity reports.

### Retrieving further activity reports

To determine a message route, sufficient information must be available from the activity reports collected. If you retrieve the activity reports related to a message from the reply-to queue that the message specified, but you not have the necessary information, look for further activity reports.

#### About this task

To determine the locations of any further activity reports, perform the following steps:

#### Procedure

1. For any queue managers in the queue manager network that deliver activity reports to a common queue, retrieve activity reports from the common queue that have a *CorrelId* that matches the *MsgId* of the original message.
2. For any queue managers in the queue manager network that do not deliver activity reports to a common queue, retrieve activity reports as follows:
  - a. Examine the existing activity reports to identify queue managers through which the message was routed.
  - b. For these queue managers, identify the queue managers that are enabled for activity recording.
  - c. For these queue managers, identify any that did not return activity reports to the specified reply-to queue.
  - d. For each of the queue managers that you identify, check the system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE` and retrieve any activity reports that have a *CorrelId* that matches the *MsgId* of the original message.
  - e. If you find no activity reports on the system queue, check the queue manager dead letter queue, if one exists. An activity report can only be delivered to a dead letter queue if the report option, `MQRO_DEAD_LETTER_Q`, is set.
3. Arrange all the acquired activity reports in order. The order of the activity reports then represents the route, or partial route, that the message took.
4. Obtain the information you need from the activity information in the ordered activity reports. In some circumstances, recorded activity information cannot reach the specified reply-to queue, a common queue, or a system queue.

## Circumstances where activity information is not acquired

To determine the complete sequence of activities performed on behalf of a message, information related to every activity must be acquired. If the information relating to any activity has not been recorded, or has not been acquired, you can determine only a partial sequence of activities.

Activity information is not recorded in the following circumstances:

- The message is processed by a WebSphere MQ queue manager earlier than Version 6.0.
- The message is processed by a queue manager that is not enabled for activity recording.
- The application that expected to process the message is not running.

Recorded activity information is unable to reach the specified reply-to queue in the following circumstances:

- There is no channel defined to route activity reports to the reply-to queue.
- The channel to route activity reports to the reply-to queue is not running.
- The remote queue definition to route activity reports back to the queue manager where the reply-to queue resides (the queue manager alias), is not defined.
- The user that generated the original message does not have open, or put, authority to the queue manager alias.
- The user that generated the original message does not have open, or put, authority to the reply-to queue.
- The reply-to queue is put inhibited.

Recorded activity information is unable to reach the system queue, or a common queue, in the following circumstances:

- If a common queue is to be used and there is no channel defined to route activity reports to the common queue.
- If a common queue is to be used and the channel to route activity reports to the common queue is not running.
- If a common queue is to be used and the system queue is incorrectly defined.
- The user that generated the original message does not have open, or put, authority to the system queue.
- The system queue is put inhibited.
- If a common queue is to be used and the user that generated the original message does not have open, or put, authority to the common queue.
- If a common queue is to be used and the common queue is put inhibited.

In these circumstances, providing the activity report does not have the report option `MQRO_DISCARD_MSG` specified, the activity report can be retrieved from a dead letter queue if one was defined on the queue manager where the activity report was rejected. An activity report will only have this report option specified if the original message, from which the activity report was generated, had both `MQRO_PASS_DISCARD_AND_EXPIRY` and `MQRO_DISCARD_MSG` specified in the Report field of the message descriptor.

## Trace-route messaging

Trace-route messaging is a technique that uses *trace-route messages* to record activity information for a message. Trace-route messaging involves sending a trace-route message into a queue manager network.

As the trace-route message is routed through the queue manager network, activity information is recorded. This activity information includes information about the applications that performed the activities, when they were performed, and the operations that were performed as part of the activities. You can use the information recorded using trace-route messaging for the following purposes:

### **To determine the last known location of a message**

If a message does not reach its intended destination, you can use the activity information recorded for a trace-route message to determine the last known location of the message. A trace-route message is sent into a queue manager network with the same target destination as the original message, intending that it follows the same route. Activity information can be accumulated in the message data of the trace-route message, or recorded using activity reports. To increase the probability that the trace-route message follows the same route as the original message, you can modify the trace-route message to mimic the original message.

### **To determine configuration issues with a queue manager network**

Trace-route messages are sent into a queue manager network and activity information is recorded. By studying the activity information recorded for a trace-route message, it can become apparent that the trace-route message did not follow the expected route. There are many reasons why this can occur, for example, a channel might be inactive, forcing the message to take an alternative route. In these situations, a system administrator can determine whether there are any problems in the queue manager network, and if there are, correct them.

You can use the WebSphere MQ display route application to configure, generate, and put trace-route messages into a queue manager network.

**Warning:** If you put a trace-route message to a distribution list, the results are undefined.

#### **Related concepts:**

“How activity information is recorded”

“Generating and configuring a trace-route message” on page 899

“Trace-route message reference” on page 952

#### **Related tasks:**

“Acquiring recorded activity information” on page 896

“Controlling trace-route messaging” on page 896

“Setting up a common queue for trace-route reply messages” on page 903

“Acquiring and using recorded information” on page 904

#### **Related reference:**

“Additional activity information” on page 905

### **How activity information is recorded**

With trace-route messaging, you can record activity information in the message data of the trace-route message, or use activity reports. Alternatively, you can use both techniques.

### **Accumulating activity information in the message data of the trace-route message**

As a trace-route message is routed through a queue manager network, information about the activities performed on behalf of the trace-route message can be accumulated in the message data of the trace-route message. The activity information is stored in *Activity* PCF groups. For every activity performed on behalf of the trace-route message, an *Activity* PCF group is written to the end of the PCF block in the message data of the trace-route message.

Additional activity information is recorded in trace-route messaging, in a PCF group called the *TraceRoute* PCF group. The additional activity information is stored in this PCF group, and can be used to help determine the sequence of recorded activities. This technique is controlled by the *Accumulate* parameter in the *TraceRoute* PCF group.

### **Recording activity information using activity reports**

As a trace-route message is routed through a queue manager network, an activity report can be generated for every activity that was performed on behalf of the trace-route message. The activity information is

stored in the *Activity* PCF group. For every activity performed on behalf of a trace-route message, an activity report is generated containing an *Activity* PCF group. Activity recording for trace-route messages works in the same way as for any other message.

Activity reports generated for trace-route messages contain additional activity information compared to the those generated for any other message. The additional information is returned in a *TraceRoute* PCF group. The information contained in the *TraceRoute* PCF group is accurate only from the time the activity report was generated. You can use the additional information to help determine the sequence of activities performed on behalf of the trace-route message.

## Acquiring recorded activity information

When a trace-route message has reached its intended destination, or is discarded, the method that you use to acquire the activity information depends on how that information was recorded.

### Before you begin

If you are unfamiliar with activity information, refer to “How activity information is recorded” on page 895.

### About this task

Use the following methods to acquire the activity information after the trace-route message has reached its intended destination, or is discarded:

### Procedure

- Retrieve the trace-route message. The *Deliver* parameter, in the *TraceRoute* PCF group, controls whether a trace-route message is placed on the target queue on arrival, or whether it is discarded. If the trace-route message is delivered to the target queue, you can retrieve the trace-route message from this queue. Then, you can use the WebSphere MQ display route application to display the activity information.

To request that activity information is accumulated in the message data of a trace-route message, set the *Accumulate* parameter in the *TraceRoute* PCF group to `MQRROUTE_ACCUMULATE_IN_MSG`.

- Use a trace-route reply message. When a trace-route message reaches its intended destination, or the trace-route message cannot be routed any further in a queue manager network, a trace-route reply message can be generated. A trace-route reply message contains a duplicate of all the activity information from the trace-route message, and is either delivered to a specified reply-to queue, or the system queue `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE`. You can use the WebSphere MQ display route application to display the activity information.

To request a trace-route reply message, set the *Accumulate* parameter in the *TraceRoute* PCF group to `MQRROUTE_ACCUMULATE_AND_REPLY`.

- Use activity reports. If activity reports are generated for a trace-route message, you must locate the activity reports before you can acquire the activity information. Then, to determine the sequence of activities, you must order the activity reports.

## Controlling trace-route messaging

Enable trace-route messaging at the queue manager level, so that applications in the scope of that queue manager can write activity information to a trace-route message. To enable an entire queue manager network, individually enable every queue manager in the network for trace-route messaging. If you enable more queue managers, more activity reports are generated.

### Before you begin

If you are using activity reports to record activity information for a trace-route message, refer to “Controlling activity recording” on page 891.

## About this task

To record activity information for a trace-route message as it is routed through a queue manager, perform the following steps:

### Procedure

- Define how activity information is to be recorded for the trace-route message. Refer to “Generating and configuring a trace-route message” on page 899
- If you want to accumulate activity information in the trace-route message, ensure that the queue manager is enabled for trace-route messaging
- If you want to accumulate activity information in the trace-route message, ensure that applications performing activities on the trace-route message are capable of writing activity information to the message data of the trace-route message

### Related concepts:

“Generating and configuring a trace-route message” on page 899

### Related tasks:

“Enabling applications for trace-route messaging” on page 898

“Controlling activity recording” on page 891

### Related reference:

“Enabling queue managers for trace-route messaging”

### Enabling queue managers for trace-route messaging:

To control whether queue managers are enabled or disabled for trace-route messaging use the queue manager attribute ROUTEREC.

Use the MQSC command ALTER QMGR, specifying the parameter ROUTEREC to change the value of the queue manager attribute. The value can be:

**MSG** The queue manager is enabled for trace-route messaging. Applications within the scope of the queue manager can write activity information to the trace-route message.

If the *Accumulate* parameter in the *TraceRoute* PCF group is set as MQRROUTE\_ACCUMULATE\_AND\_REPLY, and the next activity to be performed on the trace-route message:

- is a discard
- is a put to a local queue (target queue or dead-letter queue)
- will cause the total number of activities performed on the trace-route message to exceed the value of parameter the *MaxActivities*, in the *TraceRoute* PCF group .

a trace-route reply message is generated, and delivered to the reply-to queue specified in the message descriptor of the trace-route message.

### QUEUE

The queue manager is enabled for trace-route messaging. Applications within the scope of the queue manager can write activity information to the trace-route message.

If the *Accumulate* parameter in the *TraceRoute* PCF group is set as MQRROUTE\_ACCUMULATE\_AND\_REPLY, and the next activity to be performed on the trace-route message:

- is a discard
- is a put to a local queue (target queue or dead-letter queue)
- will cause the total number of activities performed on the trace-route message to exceed the value of parameter the *MaxActivities*, in the *TraceRoute* PCF group .

a trace-route reply message is generated, and delivered to the local system queue SYSTEM.ADMIN.TRACE.ROUTE.QUEUE.

## DISABLED

The queue manager is disabled for trace-route messaging. Activity information is not accumulated in the the trace-route message, however the *TraceRoute* PCF group can be updated while in the scope of this queue manager.

For example, to disable a queue manager for trace-route messaging, use the following MQSC command:  
ALTER QMGR ROUTEREC(DISABLED)

**Remember:** When you modify the *ROUTEREC* queue manager attribute, a running MCA does not detect the change until the channel is restarted.

### Enabling applications for trace-route messaging:

To enable trace-route messaging for a user application, base your algorithm on the algorithm used by message channel agents (MCAs)

### Before you begin

If you are not familiar with the format of a trace-route message, see “Trace-route message reference” on page 952.

### About this task

Message channel agents (MCAs) are enabled for trace-route messaging. To enable a user application for trace-route messaging, use the following steps from the algorithm that MCAs use:

### Procedure

1. Determine whether the message being processed is a trace-route message. If the message does not conform to the format of a trace-route message, the message is not processed as a trace-route message.
2. Determine whether activity information is to be recorded. If the detail level of the performed activity is not less than the level of detail specified by the *Detail* parameter, activity information is recorded under specific circumstances. This information is only recorded if the trace-route message requests accumulation, and the queue manager is enabled for trace-route messaging, or if the trace-route message requests an activity report and the queue manager is enabled for activity recording.
  - If activity information is to be recorded, increment the *RecordedActivities* parameter.
  - If activity information is not to be recorded, increment the *UnrecordedActivities* parameter.
3. Determine whether the total number of activities performed on the trace-route message exceeds the value of the *MaxActivities* parameter.

The total number of activities is the sum of *RecordedActivities*, *UnrecordedActivities*, and *DiscontinuityCount*.

If the total number of activities exceeds *MaxActivities*, reject the message with feedback *MQFB\_MAX\_ACTIVITIES*.
4. If value of *Accumulate* is set as *MQROUTE\_ACCUMULATE\_IN\_MSG* or *MQROUTE\_ACCUMULATE\_AND\_REPLY*, and the queue manager is enabled for trace-route messaging, write an Activity PCF group to the end of the PCF block in the message data of a trace-route message.
5. Deliver the trace-route message to a local queue.
  - If the parameter, *Deliver*, is specified as *MQROUTE\_DELIVER\_NO*, reject the trace-route message with feedback *MQFB\_NOT\_DELIVERED*.
  - If the parameter, *Deliver*, is specified as *MQROUTE\_DELIVER\_YES*, deliver the trace-route message to the local queue.
6. Generate a trace-route reply message if all the following conditions are true:
  - The trace-route message was delivered to a local queue or rejected

- The value of the parameter, *Accumulate*, is MQROUTE\_ACCUMULATE\_AND\_REPLY
- The queue manager is enabled for trace-route messaging

The trace-route reply message is put on the queue determined by the ROUTEREC queue manager attribute.

7. If the trace-route message requested an activity report and the queue manager is enabled for activity recording, generate an activity report. The activity report is put on the queue determined by the ACTIVREC queue manager attribute.

## Generating and configuring a trace-route message

A trace-route message comprises specific message descriptor and message data parts. To generate a trace-route message, either create the message manually or use the WebSphere MQ display route application.

A trace-route message consists of the following parts:

### Message descriptor

An MQMD structure, with the *Format* field set to MQFMT\_ADMIN or MQFMT\_EMBEDDED\_PCF.

### Message data

One of the following combinations:

- A PCF header (MQCFH) and trace-route message data, if *Format* is set to MQFMT\_ADMIN
- An embedded PCF header (MQEPH), trace-route message data, and additional user-specified message data, if *Format* is set to MQFMT\_EMBEDDED\_PCF

The trace-route message data consists of the *TraceRoute* PCF group and one or more *Activity* PCF groups.

## Manual generation

When generating a trace-route message manually, an *Activity* PCF group is not required. *Activity* PCF groups are written to the message data of the trace-route message when an MCA or user-written application performs an activity on its behalf.

## The WebSphere MQ display route application

Use the WebSphere MQ display route application, **dspmqrte**, to configure, generate and put a trace-route message into a queue manager network. Set the *Format* parameter in the message descriptor to MQFMT\_ADMIN. You cannot add user data to the trace-route message generated by the WebSphere MQ display route application.

**Restriction:** **dspmqrte** cannot be issued on queue managers before WebSphere MQ Version 6.0 or on WebSphere MQ for z/OS queue managers. If you want the first queue manager the trace-route message is routed through to be a queue manager of this type, connect to the queue manager as a WebSphere MQ Version 6.0 or later client using the optional parameter *-c*.

**Related reference:**

“Mimicking the original message”

“The TraceRoute PCF group”

**Mimicking the original message:**

When using a trace-route message to determine the route another message has taken through a queue manager network, the more closely a trace-route message mimics the original message, the greater the chance that the trace-route message will follow the same route as the original message.

The following message characteristics can affect where a message is forwarded to within a queue manager network:

**Priority**

The priority can be specified in the message descriptor of the message.

**Persistence**

The persistence can be specified in the message descriptor of the message.

**Expiration**

The expiration can be specified in the message descriptor of the message.

**Report options**

Report options can be specified in the message descriptor of the message.

**Message size**

To mimic the size of a message, additional data can be written to the message data of the message. For this purpose, additional message data can be meaningless.

**Tip:** The WebSphere MQ display route application cannot specify message size.

**Message data**

Some queue manager networks use content based routing to determine where messages are forwarded. In these cases the message data of the trace-route message needs to be written to mimic the message data of the original message.

**Tip:** The WebSphere MQ display route application cannot specify message data.

**The TraceRoute PCF group:**

Attributes in the *TraceRoute* PCF group control the behavior of a trace-route message. The *TraceRoute* PCF group is in the message data of every trace-route message.

The following table lists the parameters in the *TraceRoute* group that an MCA recognizes. Further parameters can be added if user-written applications are written to recognize them, as described in “Additional activity information” on page 905.

Table 105. *TraceRoute* PCF group

| Parameter            | Type   |
|----------------------|--------|
| TraceRoute           | MQCFGR |
| Detail               | MQCFIN |
| RecordedActivities   | MQCFIN |
| UnrecordedActivities | MQCFIN |
| DiscontinuityCount   | MQCFIN |
| MaxActivities        | MQCFIN |
| Accumulate           | MQCFIN |
| Forward              | MQCFIN |
| Deliver              | MQCFIN |

Descriptions of each parameter in the *TraceRoute* PCF group follows:

*Detail* Specifies the detail level of activity information that is to be recorded. The value can be:

**MQROUTE\_DETAIL\_LOW**

Only activities performed by user application are recorded.

**MQROUTE\_DETAIL\_MEDIUM**

Activities specified in *MQROUTE\_DETAIL\_LOW* should be recorded. Additionally, activities performed by MCAs are recorded.

**MQROUTE\_DETAIL\_HIGH**

Activities specified in *MQROUTE\_DETAIL\_LOW*, and *MQROUTE\_DETAIL\_MEDIUM* should be recorded. MCAs do not record any further activity information at this level of detail. This option is only available to user applications that are to record further activity information. For example, if a user application determines the route a message takes by considering certain message characteristics, the information about the routing logic could be included with this level of detail.

*RecordedActivities*

Specifies the number of recorded activities performed on behalf of the trace-route message. An activity is considered to be recorded if information about it has been written to the trace-route message, or if an activity report has been generated. For every recorded activity, *RecordedActivities* increments by one.

*UnrecordedActivities*

Specifies the number of unrecorded activities performed on behalf of the trace-route message. An activity is considered to be unrecorded if an application that is enabled for trace-route messaging neither accumulates, nor writes the related activity information to an activity report.

An activity performed on behalf of a trace-route message is unrecorded in the following circumstances:

- The detail level of the performed activity is less than the level of detail specified by the parameter *Detail*.
- The trace-route message requests an activity report but not accumulation, and the queue manager is not enabled for activity recording.
- The trace-route message requests accumulation but not an activity report, and the queue manager is not enabled for trace-route messaging.
- The trace-route message requests both accumulation and an activity report, and the queue manager is not enabled for activity recording and trace route messaging.
- The trace-route message requests neither accumulation nor an activity report.

For every unrecorded activity the parameter, *UnrecordedActivities*, increments by one.

*DiscontinuityCount*

Specifies the number of times the trace-route message has been routed through a queue manager with applications that were not enabled for trace-route messaging. This value is incremented by the queue manager. If this value is greater than 0, only a partial message route can be determined.

*MaxActivities*

Specifies the maximum number of activities that can be performed on behalf of the trace-route message.

The total number of activities is the sum of *RecordedActivities*, *UnrecordedActivities*, and *DiscontinuityCount*. The total number of activities must not exceed the value of *MaxActivities*.

The value of *MaxActivities* can be:

**A positive integer**

The maximum number of activities.

If the maximum number of activities is exceeded, the trace-route message is rejected with feedback MQFB\_MAX\_ACTIVITIES. This can prevent the trace-route message from being forwarded indefinitely if caught in an infinite loop.

**MQROUTE\_UNLIMITED\_ACTIVITIES**

An unlimited number of activities can be performed on behalf of the trace-route message.

*Accumulate*

Specifies the method used to accumulate activity information. The value can be:

**MQROUTE\_ACCUMULATE\_IN\_MSG**

If the queue manager is enabled for trace-route messaging, activity information is accumulated in the message data of the trace-route message.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.
- Zero or more *Activity* PCF groups.

**MQROUTE\_ACCUMULATE\_AND\_REPLY**

If the queue manager is enabled for trace-route messaging, activity information is accumulated in the message data of the trace-route message, and a trace-route reply message is generated if any of the following occur:

- The trace-route message is discarded by a WebSphere MQ Version 6 (or later) queue manager.
- The trace-route message is put to a local queue (target queue or dead-letter queue) by a WebSphere MQ Version 6 (or later) queue manager.
- The number of activities performed on the trace-route message exceeds the value of *MaxActivities*.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.
- Zero or more *Activity* PCF groups.

**MQROUTE\_ACCUMULATE\_NONE**

Activity information is not accumulated in the message data of the trace-route message.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.

*Forward*

Specifies where a trace-route message can be forwarded to. The value can be:

**MQROUTE\_FORWARD\_IF\_SUPPORTED**

The trace-route message is only forwarded to queue managers that will honor the value of the *Deliver* parameter from the *TraceRoute* group.

**MQROUTE\_FORWARD\_ALL**

The trace-route message is forwarded to any queue manager, regardless of whether the value of the *Deliver* parameter will be honored.

Queue managers use the following algorithm when determining whether to forward a trace-route message to a remote queue manager:

1. Determine whether the remote queue manager is capable of supporting trace-route messaging.
  - If the remote queue manager is capable of supporting trace-route messaging, the algorithm continues to step 4 on page 903.

- If the remote queue manager is not capable of supporting trace-route messaging, the algorithm continues to step 2
2. Determine whether the *Deliver* parameter from the *TraceRoute* group contains any unrecognized delivery options in the MQROUTE\_DELIVER\_REJ\_UNSUP\_MASK bit mask.
    - If any unrecognized delivery options are found, the trace-route message is rejected with feedback MQFB\_UNSUPPORTED\_DELIVERY.
    - If no unrecognized delivery options are found, the algorithm continues to step 3.
  3. Determine the value of the parameter *Deliver* from the *TraceRoute* PCF group in the trace-route message.
    - If *Deliver* is specified as MQROUTE\_DELIVER\_YES, the trace-route message is forwarded to the remote queue manager.
    - If *Deliver* is specified as MQROUTE\_DELIVER\_NO, the algorithm continues to step 4.
  4. Determine whether the *Forward* parameter from the *TraceRoute* group contains any unrecognized forwarding options in the MQROUTE\_FORWARDING\_REJ\_UNSUP\_MASK bit mask.
    - If any unrecognized forwarding options are found, the trace-route message is rejected with feedback MQFB\_UNSUPPORTED\_FORWARDING.
    - If no unrecognized forwarding options are found, the algorithm continues to step 5.
  5. Determine the value of the parameter *Forward* from the *TraceRoute* PCF group in the trace-route message.
    - If *Forward* is specified as MQROUTE\_FORWARD\_IF\_SUPPORTED, the trace-route message is rejected with feedback MQFB\_NOT\_FORWARDED.
    - If *Forward* is specified as MQROUTE\_FORWARD\_ALL, trace-route message can be forwarded to the remote queue manager.

*Deliver* Specifies the action to be taken if the trace-route message reaches its intended destination. User-written applications must check this attribute before placing a trace-route message on its target queue. The value can be:

**MQROUTE\_DELIVER\_YES**

On arrival, the trace-route message is put on the target queue. Any application performing a get operation on the target queue can retrieve the trace-route message.

**MQROUTE\_DELIVER\_NO**

On arrival, the trace-route message is not delivered to the target queue. The message is processed according to its report options.

**Setting up a common queue for trace-route reply messages**

To determine the locations of the trace-route reply messages related to a specific message when the reports are delivered to the local system queue, it is more efficient to use a common queue on a single node

**Before you begin**

Set the ROUTEREC parameter to enable the queue manager for trace-route messaging and to specify that any trace-route reply messages generated are delivered to the local system queue SYSTEM.ADMIN.TRACE.ROUTE.QUEUE.

**About this task**

If a number of queue managers in a queue manager network are set to deliver trace-route reply messages to the local system queue, it can be time consuming to determine the locations of the trace-route reply

messages related to a specific message. Alternatively, use a single node, which is a queue manager that hosts a common queue. All the queue managers in a queue manager network can deliver trace-route reply messages to this common queue. The benefit of using a common queue is that queue managers do not have to deliver trace-route reply messages to the reply-to queue specified in a message and, when determining the locations of the trace-route reply messages related to a message, you query one queue only.

To set up a common queue, perform the following steps:

### **Procedure**

1. Select or define a queue manager as the single node
2. On the single node, select or define a queue for use as the common queue
3. On all queue managers that forward trace-route reply messages to the common queue, redefine the local system queue `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE` as a remote queue definition
  - a. Specify the name of the single node as the remote queue manager name
  - b. Specify the name of the common queue as the remote queue name

### **Acquiring and using recorded information**

Use any of the following techniques to acquire recorded activity information for a trace-route message

#### **About this task**

Note that the circumstances in which activity information is not acquired apply also to trace-route reply messages.

Activity information is not recorded when a trace-route message is processed by a queue manager that is disabled for both activity recording and trace-route messaging.

#### **Related tasks:**

“Acquiring information from trace-route reply messages”

“Acquiring information from trace-route messages” on page 905

“Acquiring information from activity reports” on page 905

#### **Acquiring information from trace-route reply messages:**

To acquire activity information you locate the trace-route reply message. Then you retrieve the message and analyze the activity information.

#### **About this task**

You can acquire activity information from a trace-route reply message only if you know the location of the trace-route reply message. Locate the message and process the activity information as follows:

### **Procedure**

1. Check the reply-to queue that was specified in the message descriptor of the trace-route message. If the trace-route reply message is not on the reply-to queue, check the following locations:
  - The local system queue, `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE`, on the target queue manager of the trace-route message
  - The common queue, if you have set up a common queue for trace-route reply messages
  - The local system queue, `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE`, on any other queue manager in the queue manager network, which can occur if the trace-route message has been put to a dead-letter queue, or the maximum number of activities was exceeded
2. Retrieve the trace-route reply message

3. Use the WebSphere MQ display route application to display the recorded activity information
4. Study the activity information and obtain the information that you need

#### **Acquiring information from trace-route messages:**

To acquire activity information you locate the trace-route message, which must have the appropriate parameters in the *TraceRoute* PCF group. Then you retrieve the message and analyze the activity information.

#### **About this task**

You can acquire activity information from a trace-route message only if you know the location of the trace-route message and it has the parameter *Accumulate* in the *TraceRoute* PCF group specified as either `MQROUTE_ACCUMULATE_IN_MSG` or `MQROUTE_ACCUMULATE_AND_REPLY`.

For the trace-route message to be delivered to the target queue the *Deliver* parameter in the *TraceRoute* PCF group must be specified as `MQROUTE_DELIVER_YES`.

#### **Procedure**

1. Check the target queue. If the trace-route message is not on the target queue, you can try to locate the trace-route message using a trace-route message enabled for activity recording. With the generated activity reports try to determine the last known location of the trace-route message.
2. Retrieve the trace-route message
3. Use the WebSphere MQ display route application to display the recorded activity information
4. Study the activity information and obtain the information that you need

#### **Acquiring information from activity reports:**

To acquire activity information you locate the activity report, which must have the report option specified in the message descriptor. Then you retrieve the activity report and analyze the activity information.

#### **About this task**

You can acquire activity information from an activity report only if you know the location of the activity report and the report option `MQRO_ACTIVITY` was specified in the message descriptor of the trace-route message.

#### **Procedure**

1. Locate and order the activity reports generated for a trace-route message. When you have located the activity reports, you can order them manually or use the WebSphere MQ display route application to order and display the activity information automatically.
2. Study the activity information and obtain the information that you need

#### **Additional activity information**

As a trace-route message is routed through a queue manager network, user applications can record additional information by including one or more additional PCF parameters when writing the *Activity* group to the message data of the trace-route message or activity report.

Additional activity information can help system administrators to identify the route taken by a trace-route message took, or why that route was taken.

If you use the WebSphere MQ display route application to display the recorded information for a trace-route message, any additional PCF parameters can only be displayed with a numeric identifier, unless the parameter identifier of each parameter is recognized by the WebSphere MQ display route

application. To recognize a parameter identifier, additional information must be recorded using the following PCF parameters. Include these PCF parameters in an appropriate place in the *Activity* PCF group.

#### *GroupName*

Description: Grouped parameters specifying the additional information.  
Identifier: MQGACF\_VALUE\_NAMING.  
Data type: MQCFGR.  
Parameters in group: *ParameterName*  
*ParameterValue*

#### *ParameterName*

Description: Contains the **name** to be displayed by the WebSphere MQ display route application, which puts the value of *ParameterValue* into context.  
Identifier: MQCA\_VALUE\_NAME.  
Data type: MQCFST.  
Included in PCF group: *GroupName*.  
Value: The name to be displayed.

#### *ParameterValue*

Description: Contains the **value** to be displayed by the WebSphere MQ display route application.  
Identifier: The PCF structure identifier for the additional information.  
Data type: The PCF structure data type for the additional information.  
Included in PCF group: *GroupName*.  
Value: The value to be displayed.

## Examples of recording additional activity information

The following examples illustrate how a user application can record additional information when performing an activity on behalf of a trace-route message. In both examples, the WebSphere MQ display route application is used to generate a trace-route message, and display the activity information returned to it.

### Related concepts:

“Example 1”

“Example 2” on page 907

### Example 1:

Additional activity information is recorded by a user application in a format where the parameter identifier *is not* recognized by the WebSphere MQ display route application.

1. The WebSphere MQ display route application is used to generate and put a trace-route message into a queue manager network. The necessary options are set to request the following:
  - Activity information is accumulated in the message data of the trace-route message.
  - On arrival at the target queue the trace-route message is discarded, and a trace-route reply message is generated and delivered to a specified reply-to queue.
  - On receipt of the trace-route reply message, the WebSphere MQ display route application displays the accumulated activity information.

The trace-route message is put into the queue manager network.

- As the trace-route message is routed through the queue manager network a user application, that is enabled for trace-route messaging, performs a low detail activity on behalf of the message. In addition to writing the standard activity information to the trace-route message, the user application writes the following PCF parameter to the end of the Activity group:

*ColorValue*

Identifier: 65536.  
 Data type: MQCFST.  
 Value: 'Red'

This additional PCF parameter gives further information about the activity that was performed, however it is written in a format where the parameter identifier *is not* recognized by the WebSphere MQ display route application.

- The trace-route messages reaches the target queue and a trace-route reply message is returned to the WebSphere MQ display route application. The additional activity information is displayed as follows:  
 65536: 'Red'

The WebSphere MQ display route application does not recognize the parameter identifier of the PCF parameter and displays it as a numeric value. The context of the additional information is not clear. For an example of when the WebSphere MQ display route application does recognize the parameter identifier of the PCF parameter, see “Example 2.”

### Example 2:

Additional activity information is recorded by a user application in a format where the parameter identifier *is* recognized by the WebSphere MQ display route application.

- The WebSphere MQ display route application is used to generate and put a trace-route message into a queue manager network in the same fashion as in “Example 1” on page 906.
- As the trace-route message is routed through the queue manager network a user application, that is enabled for trace-route messaging, performs a low detail activity on behalf of the message. In addition to writing the standard activity information to the trace-route message, the user application writes the following PCF parameters to the end of the Activity group:

*ColorInfo*

Description: Grouped parameters specifying information about a color.  
 Identifier: MQGACF\_VALUE\_NAMING.  
 Data type: MQCFGR.  
 Parameters in group: *ColorName*  
                           *ColorValue*

*ColorName*

|                        |                                                                                                                                                |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:           | Contains the <b>name</b> to be displayed by the WebSphere MQ display route application which puts the value of <i>ColorValue</i> into context. |
| Identifier:            | MQCA_VALUE_NAME.                                                                                                                               |
| Data type:             | MQCFST.                                                                                                                                        |
| Included in PCF group: | <i>ColorInfo</i> .                                                                                                                             |
| Value:                 | 'Color'                                                                                                                                        |

### *ColorValue*

|                        |                                                                                          |
|------------------------|------------------------------------------------------------------------------------------|
| Description:           | Contains the <b>value</b> to be displayed by the WebSphere MQ display route application. |
| Identifier:            | 65536.                                                                                   |
| Data type:             | MQCFST.                                                                                  |
| Included in PCF group: | <i>ColorInfo</i> .                                                                       |
| Value:                 | 'Red'                                                                                    |

These additional PCF parameters gives further information about the activity that was performed. These PCF parameters are written in a format where the parameter identifier *is* recognized by the WebSphere MQ display route application.

- The trace-route messages reaches the target queue and a trace-route reply message is returned to the WebSphere MQ display route application. The additional activity information is displayed as follows:  
Color: 'Red'

The WebSphere MQ display route application recognizes that the parameter identifier of the PCF structure containing the value of the additional activity information has a corresponding name. The corresponding name is displayed instead of the numeric value.

## WebSphere MQ display route application

Use the WebSphere MQ display route application (**dspmqrte**) to work with trace-route messages and activity information related to a trace-route message, using a command-line interface. The WebSphere MQ display route application is not available for WebSphere MQ for z/OS queue managers.

You can run the WebSphere MQ display route application (**dspmqrte**) on all WebSphere MQ Version 7.0 queue managers, with the exception of WebSphere MQ for z/OS queue managers. You can run the WebSphere MQ display route application as a client to a WebSphere MQ for z/OS Version 7.0 queue manager by specifying the **-c** parameter when issuing the **dspmqrte** command.

**Note:** To run a Client Application against a queue manager, the Client Attachment feature must be installed.

You can use the WebSphere MQ display route application for the following purposes:

- To configure, generate, and put a trace-route message into a queue manager network.  
By putting a trace-route message into a queue manager network, activity information can be collected and used to determine the route that the trace-route message took. You can specify the characteristics of the trace-route messages as follows:
  - The destination of the trace-route message.
  - How the trace-route message mimics another message.
  - How the trace-route message should be handled as it is routed through a queue manager network.
  - Whether activity recording or trace-route messaging are used to record activity information.
- To order and display activity information related to a trace-route message.

If the WebSphere MQ display route application has put a trace-route message into a queue manager network, after the related activity information has been returned, the information can be ordered and displayed immediately. Alternatively, the WebSphere MQ display route application can be used to order, and display, activity information related to a trace-route message that was previously generated.

**Related concepts:**

“WebSphere MQ display route application examples” on page 916

**Related reference:**

“Parameters for trace-route messages”

“Display of activity information” on page 914



dspmqrte (*WebSphere MQ V7.1 Reference*)

## Parameters for trace-route messages

Use this page to obtain an overview of the parameters provided by the WebSphere MQ display route application, **dspmqrte**, to determine the characteristics of a trace-route message, including how it is treated as it is routed through a queue manager network.

**Related reference:**

“Queue manager connection”

“The target destination”

“The publication topic” on page 910

“Message mimicking” on page 910

“Recorded activity information” on page 911

“How the trace-route message is handled” on page 913



dspmqrte (*WebSphere MQ V7.1 Reference*)

### Queue manager connection:

Use this page to specify the queue manager that the WebSphere MQ display route application connects to

**-c**

Specifies that the WebSphere MQ display route application connects as a client application.

If you do not specify this parameter, the WebSphere MQ display route application does not connect as a client application.

**-m** *QMgrName*

The name of the queue manager to which the WebSphere MQ display route application connects. The name can contain up to 48 characters.

If you do not specify this parameter, the default queue manager is used.

### The target destination:

Use this page to specify the target destination of a trace-route message

**-q** *TargetQName*

If the WebSphere MQ display route application is being used to send a trace-route message into a queue manager network, *TargetQName* specifies the name of the target queue.

**-ts** *TargetTopicString*

Specifies the topic string.

**-qmqm** *TargetQMGr*

Qualifies the target destination; normal queue manager name resolution will then apply. The target destination is specified with **-q** *TargetQName* or **-ts** *TargetTopicString*.

If you do not specify this parameter, the queue manager to which the WebSphere MQ display route application is connected is used as the target queue manager.

- o Specifies that the target destination is not bound to a specific destination. Typically this parameter is used when the trace-route message is to be put across a cluster. The target destination is opened with option MQOO\_BIND\_NOT\_FIXED.

If you do not specify this parameter, the target destination is bound to a specific destination.

### The publication topic:

For publish/subscribe applications, use this page to specify the topic string of a trace-route message for the WebSphere MQ display route application to publish

#### -ts *TopicName*

Specifies a topic string to which the WebSphere MQ display route application is to publish a trace-route message, and puts this application into topic mode. In this mode, the application traces all of the messages that result from the publish request.

You can also use the WebSphere MQ display route application to display the results from an activity report that was generated for publish messages.

### Message mimicking:

Use this page to configure a trace-route message to mimic a message, for example when the original message did not reach its intended destination

One use of trace-route messaging is to help determine the last known location of a message that did not reach its intended destination. The WebSphere MQ display route application provides parameters that can help configure a trace-route message to mimic the original message. When mimicking a message, you can use the following parameters:

#### -l *Persistence*

Specifies the persistence of the generated trace-route message. Possible values for *Persistence* are:

|            |                                                                                                                                                                                          |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>yes</b> | The generated trace-route message is persistent. (MQPER_PERSISTENT).                                                                                                                     |
| <b>no</b>  | The generated trace-route message is <b>not</b> persistent. (MQPER_NOT_PERSISTENT).                                                                                                      |
| <b>q</b>   | The generated trace-route message inherits its persistence value from the destination specified by <i>-q TargetQName</i> or <i>-ts TargetTopicString</i> . (MQPER_PERSISTENCE_AS_Q_DEF). |

A trace-route reply message, or any report messages, returned will share the same persistence value as the original trace-route message.

If *Persistence* is specified as **yes**, you must specify the parameter *-rq ReplyToQ*. The reply-to queue must not resolve to a temporary dynamic queue.

If you do not specify this parameter, the generated trace-route message is **not** persistent.

#### -p *Priority*

Specifies the priority of the trace-route message. The value of *Priority* is either greater than or equal to 0, or MQPRI\_PRIORITY\_AS\_Q\_DEF. MQPRI\_PRIORITY\_AS\_Q\_DEF specifies that the priority value is taken from the destination specified by *-q TargetQName* or *-ts TargetTopicString*.

If you do not specify this parameter, the priority value is taken from the destination specified by *-q TargetQName* or *-ts TargetTopicString*.

#### -xs *Expiry*

Specifies the expiry time for the trace-route message, in seconds.

If you do not specify this parameter, the expiry time is specified as 60 seconds.

**-ro none** | *ReportOption*

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>none</b>                | Specifies no report options are set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b><i>ReportOption</i></b> | Specifies report options for the trace-route message. Multiple report options can be specified using a comma as a separator. Possible values for <i>ReportOption</i> are:<br><b>activity</b> The report option MQRO_ACTIVITY is set.<br><b>coa</b> The report option MQRO_COA_WITH_FULL_DATA is set.<br><b>cod</b> The report option MQRO_COD_WITH_FULL_DATA is set.<br><b>exception</b> The report option MQRO_EXCEPTION_WITH_FULL_DATA is set.<br><b>expiration</b> The report option MQRO_EXPIRATION_WITH_FULL_DATA is set.<br><b>discard</b> The report option MQRO_DISCARD_MSG is set. |

If neither *-ro ReportOption* nor *-ro none* are specified, then the MQRO\_ACTIVITY and MQRO\_DISCARD\_MSG report options are specified.

The WebSphere MQ display route application does not allow you to add user data to the trace-route message. If you require user data to be added to the trace-route message you must generate the trace-route message manually.

**Recorded activity information:**

Use this page to specify the method used to return recorded activity information, which you can then use to determine the route that a trace-route message has taken

Recorded activity information can be returned as follows:

- In activity reports
- In a trace-route reply message
- In the trace-route message itself (having been put on the target queue)

When using **dspmqrte**, the method used to return recorded activity information is determined using the following parameters:

**The activity report option, specified using -ro**

Specifies that activity information is returned using activity reports. By default activity recording is enabled.

**-ac -ar**

Specifies that activity information is accumulated in the trace-route message, and that a trace-route reply message is to be generated.

**-ac**

Specifies that activity information is to be accumulated within the trace-route message.

If you do not specify this parameter, activity information is **not** accumulated within the trace-route message.

**-ar**

Requests that a trace-route reply message containing all accumulated activity information is generated in the following circumstances:

- The trace-route message is discarded by a WebSphere MQ queue manager.

- The trace-route message is put to a local queue (target queue or dead-letter queue) by a WebSphere MQ queue manager.
- The number of activities performed on the trace-route message exceeds the value of specified in *-s Activities*.

**-ac -d yes**

Specifies that activity information is accumulated in the trace-route message, and that on arrival, the trace-route message will be put on the target queue.

**-ac**

Specifies that activity information is to be accumulated within the trace-route message.

If you do not specify this parameter, activity information is **not** accumulated within the trace-route message.

**-d yes**

On arrival, the trace-route message is put to the target queue, even if the queue manager does not support trace-route messaging.

If you do not specify this parameter, the trace-route message is **not** put to the target queue.

The trace-route message can then be retrieved from the target queue, and the recorded activity information acquired.

You can combine these methods as required.

Additionally, the detail level of the recorded activity information can be specified using the following parameter:

**-t *Detail***

Specifies the activities that are recorded. The possible values for *Detail* are:

|               |                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>low</b>    | Activities performed by user-defined application are recorded only.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>medium</b> | Activities specified in <b>low</b> are recorded. Additionally, publish activities and activities performed by MCAs are recorded.                                                                                                                                                                                                                                                                                                                             |
| <b>high</b>   | Activities specified in <b>low</b> , and <b>medium</b> are recorded. MCAs do not expose any further activity information at this level of detail. This option is available to user-defined applications that are to expose further activity information only. For example, if a user-defined application determines the route a message takes by considering certain message characteristics, the routing logic could be included with this level of detail. |

If you do not specify this parameter, medium level activities are recorded.

By default the WebSphere MQ display route application uses a temporary dynamic queue to store the returned messages. When the WebSphere MQ display route application ends, the temporary dynamic queue is closed, and any messages are purged. If the returned messages are required beyond the current execution of the WebSphere MQ display route application ends, then a permanent queue must be specified using the following parameters:

**-rq *ReplyToQ***

Specifies the name of the reply-to queue that all responses to the trace-route message are sent to. If the trace-route message is persistent, or if the *-n* parameter is specified, a reply-to queue must be specified that is **not** a temporary dynamic queue.

If you do not specify this parameter then a dynamic reply-to queue is created using the system default model queue, SYSTEM.DEFAULT.MODEL.QUEUE.

**-rqm** *ReplyToQMgr*

Specifies the name of the queue manager where the reply-to queue resides. The name can contain up to 48 characters.

If you do not specify this parameter, the queue manager to which the WebSphere MQ display route application is connected is used as the reply-to queue manager.

**How the trace-route message is handled:**

Use this page to control how a trace-route message is handled as it is routed through a queue manager network.

The following parameters can restrict where the trace-route message can be routed in the queue manager network:

**-d** *Deliver*

Specifies whether the trace-route message is to be delivered to the target queue on arrival. Possible values for *Deliver* are:

|            |                                                                                                                                   |
|------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>yes</b> | On arrival, the trace-route message is put to the target queue, even if the queue manager does not support trace-route messaging. |
| <b>no</b>  | On arrival, the trace-route message is <b>not</b> put to the target queue.                                                        |

If you do not specify this parameter, the trace-route message is **not** put to the target queue.

**-f** *Forward*

Specifies the type of queue manager that the trace-route message can be forwarded to. For details of the algorithm that queue managers use to determine whether to forward a message to a remote queue manager, refer to “The TraceRoute PCF group” on page 900. The possible values for *Forward* are:

|                  |                                                                                                                                                                                                                                                                                                 |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>all</b>       | The trace-route message is forwarded to any queue manager.<br><b>Warning:</b> If forwarded to a WebSphere MQ queue manager earlier than Version 6.0, the trace-route message will not be recognized and can be delivered to a local queue despite the value of the <i>-d Deliver</i> parameter. |
| <b>supported</b> | The trace-route message is only forwarded to a queue manager that will honor the <i>Deliver</i> parameter from the <i>TraceRoute</i> PCF group.                                                                                                                                                 |

If you do not specify this parameter, the trace-route message will only be forwarded to a queue manager that will honor the *Deliver* parameter.

The following parameters can prevent a trace-route message from remaining in a queue manager network indefinitely:

**-s** *Activities*

Specifies the maximum number of recorded activities that can be performed on behalf of the trace-route message before it is discarded. This prevents the trace-route message from being forwarded indefinitely if caught in an infinite loop. The value of *Activities* is either greater than or equal to 1, or MQROUTE\_UNLIMITED\_ACTIVITIES. MQROUTE\_UNLIMITED\_ACTIVITIES specifies that an unlimited number of activities can be performed on behalf of the trace-route message.

If you do not specify this parameter, an unlimited number of activities can be performed on behalf of the trace-route message.

**-xs** *Expiry*

Specifies the expiry time for the trace-route message, in seconds.

If you do not specify this parameter, the expiry time is specified as 60 seconds.

**-xp** *PassExpiry*

Specifies whether the expiry time from the trace-route message is passed on to a trace-route reply message. Possible values for *PassExpiry* are:

|            |                                                                                                                                                                                                                                                                                                                                                        |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>yes</b> | The report option MQRO_PASS_DISCARD_AND_EXPIRY is specified in the message descriptor of the trace-route message.<br><br>If a trace-route reply message, or activity reports, are generated for the trace-route message, the MQRO_DISCARD report option (if specified), and the remaining expiry time are passed on.<br><br>This is the default value. |
| <b>no</b>  | The report option MQRO_PASS_DISCARD_AND_EXPIRY is not specified.<br><br>If a trace-route reply message is generated for the trace-route message, the discard option and expiry time from the trace-route message are <b>not</b> passed on.                                                                                                             |

If you do not specify this parameter, MQRO\_PASS\_DISCARD\_AND\_EXPIRY is not specified.

**The discard report option, specified using -ro**

Specifies the MQRO\_DISCARD\_MSG report option. This can prevent the trace-route message remaining in the queue manager network indefinitely.

## Display of activity information

The WebSphere MQ display route application can display activity information for a trace-route message that it has just put into a queue manager network, or it can display activity information for a previously generated trace-route message. It can also display additional information recorded by user-written applications.

To specify whether activity information returned for a trace-route message is displayed, specify the following parameter:

**-n** Specifies that activity information returned for the trace-route message is not to be displayed.

If this parameter is accompanied by a request for a trace-route reply message, (*-ar*), or any of the report generating options from (*-ro ReportOption*), then a specific (non-model) reply-to queue must be specified using *-rq ReplyToQ*. By default, only activity report messages are requested.

After the trace-route message is put to the specified target queue, a 48 character hexadecimal string is displayed containing the message identifier of the trace-route message. The message identifier can be used by the WebSphere MQ display route application to display the activity information for the trace-route message at a later time, using the *-i CorrelId* parameter.

If you do not specify this parameter, activity information returned for the trace-route message is displayed in the form specified by the *-v* parameter.

When displaying activity information for a trace-route message that has just been put into a queue manager network, the following parameter can be specified:

**-w** *WaitTime*

Specifies the time, in seconds, that the WebSphere MQ display route application will wait for activity reports, or a trace-route reply message, to return to the specified reply-to queue.

If you do not specify this parameter, the wait time is specified as the expiry time of the trace-route message, plus 60 seconds.

When displaying previously accumulated activity information the following parameters must be set:

**-q** *TargetQName*

If the WebSphere MQ display route application is being used to view previously gathered activity information, *TargetQName* specifies the name of the queue where the activity information is stored.

**-i CorrelId**

This parameter is used when the WebSphere MQ display route application is used to display previously accumulated activity information only. There can be many activity reports and trace-route reply messages on the queue specified by *-q TargetQName*. *CorrelId* is used to identify the activity reports, or a trace-route reply message, related to a trace-route message. Specify the message identifier of the original trace-route message in *CorrelId*.

The format of *CorrelId* is a 48 character hexadecimal string.

The following parameters can be used when displaying previously accumulated activity information, or when displaying current activity information for a trace-route message:

- b** Specifies that the WebSphere MQ display route application will only browse activity reports or a trace-route reply message related to a message. This allows activity information to be displayed again at a later time.

If you do not specify this parameter, the WebSphere MQ display route application will destructively get activity reports or a trace-route reply message related to a message.

**-v summary | all | none | outline DisplayOption**

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>summary</b>               | The queues that the trace-route message was routed through are displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>all</b>                   | All available information is displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>none</b>                  | No information is displayed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>outline DisplayOption</b> | Specifies display options for the trace-route message. Multiple display options can be specified using a comma as a separator.<br><br>If no values are supplied the following is displayed: <ul style="list-style-type: none"><li>• The application name</li><li>• The type of each operation</li><li>• Any operation specific parameters</li></ul> Possible values for <i>DisplayOption</i> are:<br><b>activity</b> All non-PCF group parameters in <i>Activity</i> PCF groups are displayed.<br><b>identifiers</b> Values with parameter identifiers MQBACF_MSG_ID or MQBACF_CORREL_ID are displayed. This overrides <i>msgdelta</i> .<br><b>message</b> All non-PCF group parameters in <i>Message</i> PCF groups are displayed. When this value is specified, you cannot specify <i>msgdelta</i> .<br><b>msgdelta</b> All non-PCF group parameters in <i>Message</i> PCF groups, that have changed since the last operation, are displayed. When this value is specified, you cannot specify <i>message</i> .<br><b>operation</b> All non-PCF group parameters in <i>Operation</i> PCF groups are displayed.<br><b>traceroute</b> All non-PCF group parameters in <i>TraceRoute</i> PCF groups are displayed. |

If you do not specify this parameter, a summary of the message route is displayed.

## Display of additional information

As a trace-route message is routed through a queue manager network, user-written applications can record additional information by writing one or more additional PCF parameters to the message data of the trace-route message or to the message data of an activity report. For the WebSphere MQ display route application to display additional information in a readable form it must be recorded in a specific format, as described in “Additional activity information” on page 905.

### WebSphere MQ display route application examples

The following examples show how you can use the WebSphere MQ display route application. In each example, two queue managers (QM1 and QM2) are inter-connected by two channels (QM2.TO.QM1 and QM1.TO.QM2).

#### Related concepts:

“Example 1 - Requesting activity reports”

“Example 2 - Requesting a trace-route reply message” on page 919

“Example 3 - Delivering activity reports to the system queue” on page 922

“Example 4 - Diagnosing a channel problem” on page 925

#### Example 1 - Requesting activity reports:

Display activity information from a trace-route message delivered to the target queue

In this example the WebSphere MQ display route application connects to queue manager, QM1, and is used to generate and deliver a trace-route message to the target queue, TARGET.Q, on remote queue manager, QM2. The necessary report option is specified so that activity reports are requested as the trace-route reply message is routed. On arrival at the target queue the trace-route message is discarded. Activity information returned to the WebSphere MQ display route application using activity reports is put in order and displayed.

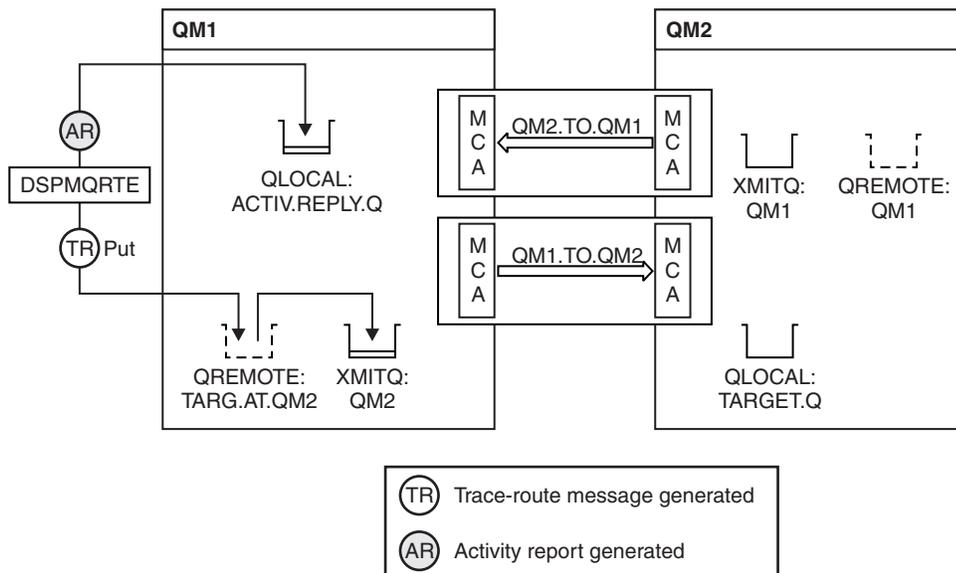


Figure 117. Requesting activity reports, Diagram 1

- The ACTIVREC attribute of each queue manager (QM1 and QM2) is set to MSG.
- The following command is issued:  
`dspmqrte -m QM1 -q TARG.AT.QM2 -rq ACTIV.REPLY.Q`

QM1 is the name of the queue manager to which the WebSphere MQ display route application connects, TARG.AT.QM2 is the name of the target queue, and ACTIV.REPLY.Q is the name of the queue to which it is requested that all responses to the trace-route message are sent.

Default values are assumed for all options that are not specified, but note in particular the -f option (the trace-route message is forwarded only to a queue manager that honors the Deliver parameter of the TraceRoute PCF group), the -d option (on arrival, the trace-route message is not put on the target queue), the -ro option (MQRO\_ACTIVITY and MQRO\_DISCARD\_MSG report options are specified), and the -t option (medium detail level activity is recorded).

- DSPMQRTE generates the trace-route message and puts it on the remote queue TARG.AT.QM2.
- DSPMQRTE then looks at the value of the ACTIVREC attribute of queue manager QM1. The value is MSG, therefore DSPMQRTE generates an activity report and puts it on the reply queue ACTIV.REPLY.Q.

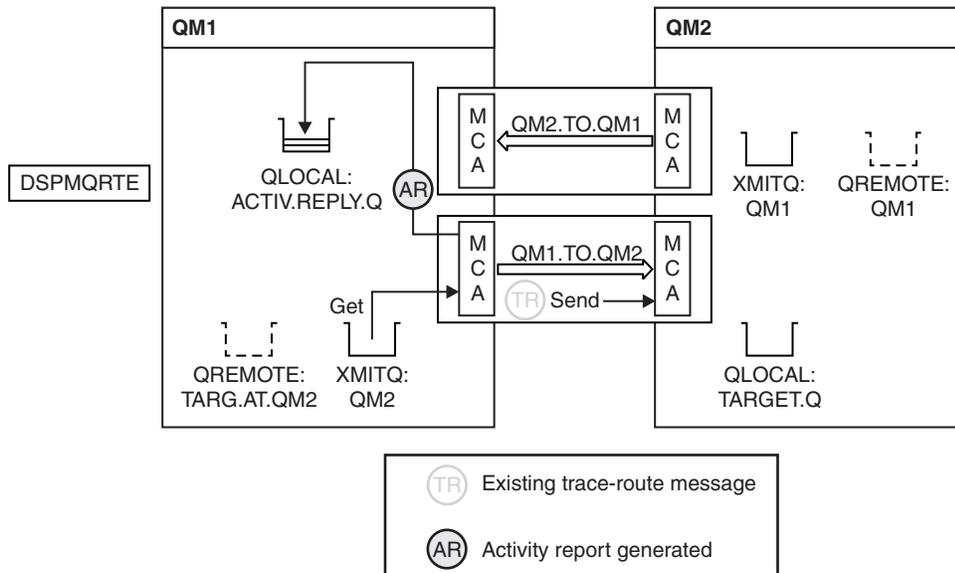


Figure 118. Requesting activity reports, Diagram 2

- The sending message channel agent (MCA) gets the trace-route message from the transmission queue. The message is a trace-route message, therefore the MCA begins to record the activity information.
- The ACTIVREC attribute of the queue manager (QM1) is MSG, and the MQRO\_ACTIVITY option is specified in the Report field of the message descriptor, therefore the MCA will later generate an activity report. The RecordedActivities parameter value in the TraceRoute PCF group is incremented by 1.
- The MCA checks that the MaxActivities value in the TraceRoute PCF group has not been exceeded.
- Before the message is forwarded to QM2 the MCA follows the algorithm that is described in Forwarding (steps 1 on page 902, 4 on page 903, and 5 on page 903) and the MCA chooses to send the message.
- The MCA then generates an activity report and puts it on the reply queue (ACTIV.REPLY.Q).

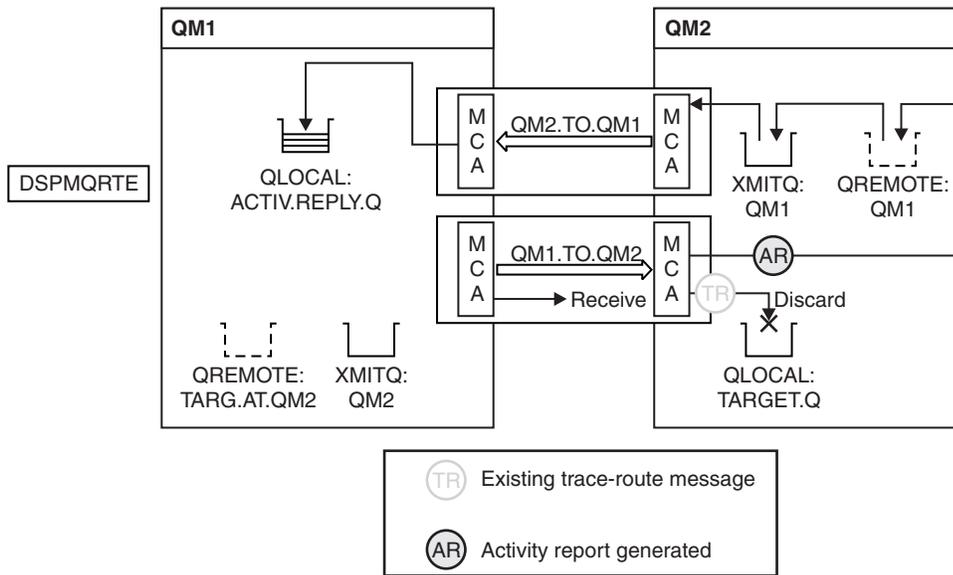


Figure 119. Requesting activity reports, Diagram 3

- The receiving MCA receives the trace-route message from the channel. The message is a trace-route message, therefore the MCA begins to record the information about the activity.
- If the queue manager that the trace-route message has come from is Version 5.3.1 or earlier, the MCA increments the DiscontinuityCount parameter of the TraceRoute PCF by 1. This is not the case here.
- The ACTIVREC attribute of the queue manager (QM2) is MSG, and the MQRO\_ACTIVITY option is specified, therefore the MCA will generate an activity report. The RecordedActivities parameter value is incremented by 1.
- The target queue is a local queue, therefore the message is discarded with feedback MQFB\_NOT\_DELIVERED, in accordance with the Deliver parameter value in the TraceRoute PCF group.
- The MCA then generates the final activity report and puts it on the reply queue. This resolves to the transmission queue that is associated with queue manager QM1 and the activity report is returned to queue manager QM1 (ACTIV.REPLY.Q).

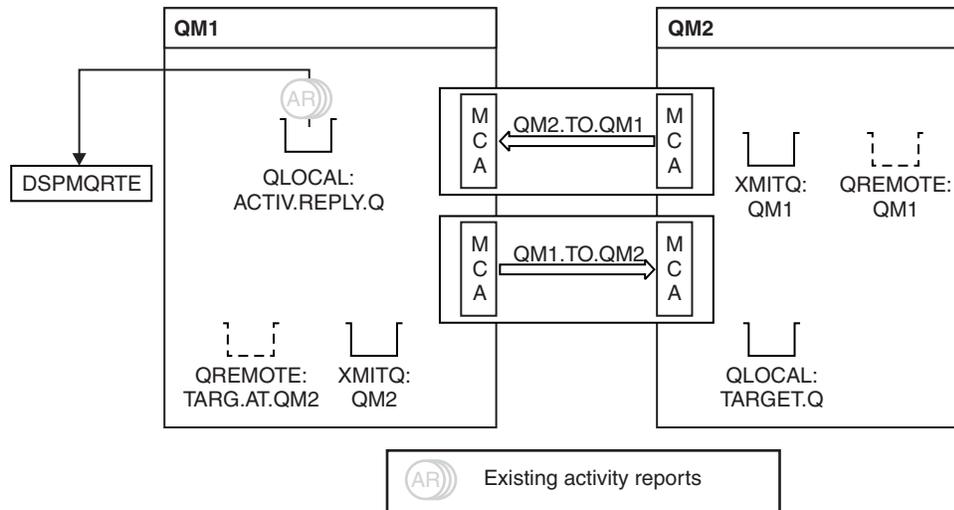


Figure 120. Requesting activity reports, Diagram 4

- Meanwhile, DSPMQRTE has been continually performing MQGETs on the reply queue (ACTIV.REPLY.Q), waiting for activity reports. It will wait for up to 120 seconds (60 seconds longer than the expiry time of the trace-route message) since -w was not specified when DSPMQRTE was started.
  - DSPMQRTE gets the 3 activity reports off the reply queue.
  - The activity reports are ordered using the RecordedActivities, UnrecordedActivities, and DiscontinuityCount parameters in the TraceRoute PCF group for each of the activities. The only value that is non-zero in this example is RecordedActivities, therefore this is the only parameter that is actually used.
  - The program ends as soon as the discard operation is displayed. Even though the final operation was a discard, it is treated as though a put took place because the feedback is MQFB\_NOT\_DELIVERED.
- The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2
-rq ACTIV.REPLY.Q'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2',
queue manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
AMQ8666: Queue 'QM2' on queue manager 'QM1'.
AMQ8666: Queue 'TARGET.Q' on queue manager 'QM2'.
AMQ8652: DSPMQRTE command has finished.
```

### Example 2 - Requesting a trace-route reply message:

Generate and deliver a trace-route message to the target queue

In this example the WebSphere MQ display route application connects to queue manager, QM1, and is used to generate and deliver a trace-route message to the target queue, TARGET.Q, on remote queue manager, QM2. The necessary option is specified so that activity information is accumulated in the trace-route message. On arrival at the target queue a trace-route reply message is requested, and the trace-route message is discarded.

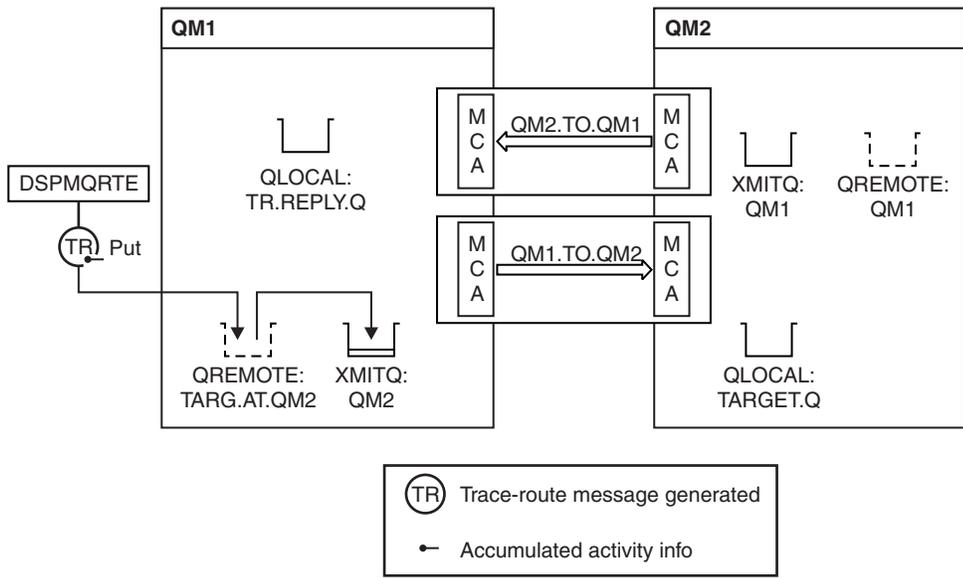


Figure 121. Requesting a trace-route reply message, Diagram 1

- The ROUTEREC attribute of each queue manager (QM1 and QM2) is set to MSG.
- The following command is issued:  
`dspmqrte -m QM1 -q TARG.AT.QM2 -rq TR.REPLY.Q -ac -ar -ro discard`

QM1 is the name of the queue manager to which the WebSphere MQ display route application connects, TARG.AT.QM2 is the name of the target queue, and ACTIV.REPLY.Q is the name of the queue to which it is requested that all responses to the trace-route message are sent. The -ac option specifies that activity information is accumulated in the trace-route message, the -ar option specifies that all accumulated activity is sent to the reply-to queue that is specified by the -rq option (that is, TR.REPLY.Q). The -ro option specifies that report option MQRO\_DISCARD\_MSG is set which means that activity reports are not generated in this example.

- DSPMQRTE accumulates activity information in the trace-route message before the message is put on the target route. The queue manager attribute ROUTEREC must not be DISABLED for this to happen.

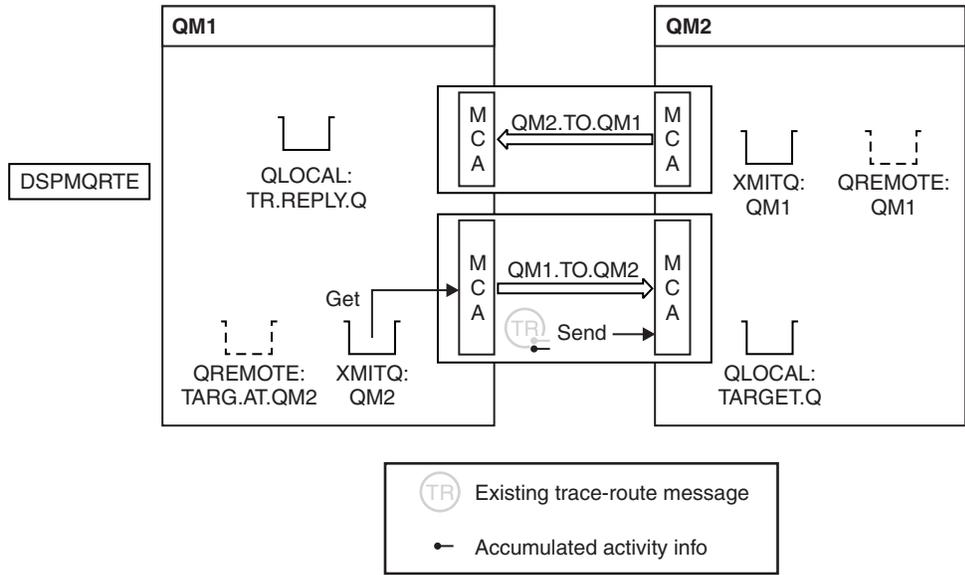


Figure 122. Requesting a trace-route reply message, Diagram 2

- The message is a trace-route message, therefore the sending MCA begins to record information about the activity.
- The queue manager attribute ROUTEREC on QM1 is not DISABLED, therefore the MCA accumulates the activity information within the message, before the message is forwarded to queue manager QM2.

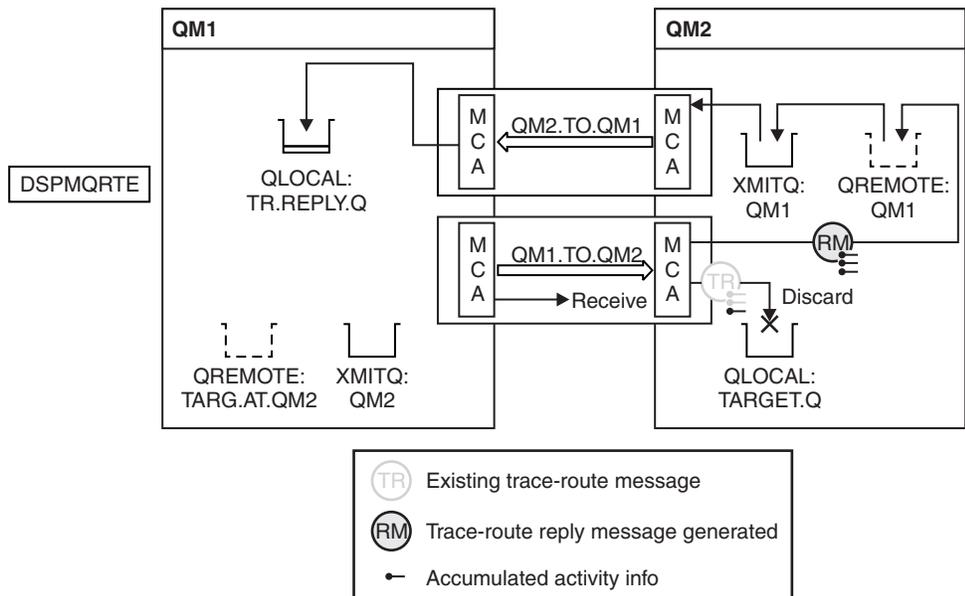


Figure 123. Requesting a trace-route reply message, Diagram 3

- The message is a trace-route message, therefore the receiving MCA begins to record information about the activity.
- The queue manager attribute ROUTEREC on QM2 is not DISABLED, therefore the MCA accumulates the information within the message.

- The target queue is a local queue, therefore the message is discarded with feedback MQFB\_NOT\_DELIVERED, in accordance with the Deliver parameter value in the TraceRoute PCF group.
- This is the last activity that will take place on the message, and because the queue manager attribute ROUTEREC on QM1 is not DISABLED, the MCA generates a trace-route reply message in accordance with the Accumulate value. The value of ROUTEREC is MSG, therefore the reply message is put on the reply queue. The reply message contains all the accumulated activity information from the trace-route message.

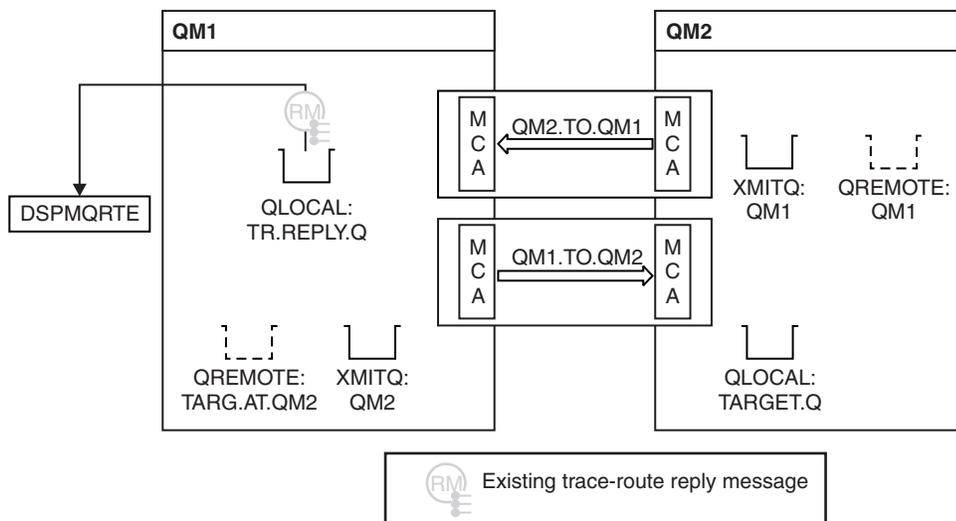


Figure 124. Requesting a trace-route reply message, Diagram 4

- Meanwhile DSPMQRTE is waiting for the trace-route reply message to return to the reply queue. When it returns, DSPMQRTE parses each activity that it contains and prints it out. The final operation is a discard operation. DSPMQRTE ends after it has been printed.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2 -rq
TR.REPLY.Q'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2', queue
manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
AMQ8666: Queue 'QM2' on queue manager 'QM1'.
AMQ8666: Queue 'TARGET.Q' on queue manager 'QM2'.
AMQ8652: DSPMQRTE command has finished.
```

### Example 3 - Delivering activity reports to the system queue:

Detect when activity reports are delivered to queues other than the reply-to queue and use the WebSphere MQ display route application to read activity reports from the other queue.

This example is the same as “Example 1 - Requesting activity reports” on page 916, except that QM2 now has the value of the ACTIVREC queue manage attribute set to QUEUE. Channel QM1.TO.QM2 must have been restarted for this to take effect.

This example demonstrates how to detect when activity reports are delivered to queues other than the reply-to queue. Once detected, the WebSphere MQ display route application is used to read activity reports from another queue.

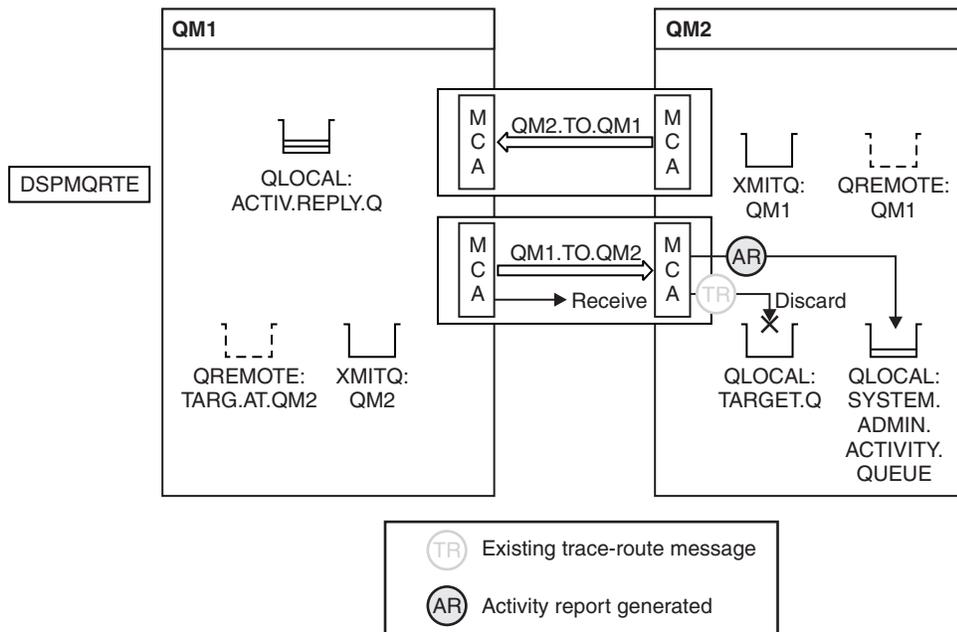


Figure 125. Delivering activity reports to the system queue, Diagram 1

- The message is a trace-route message, therefore the receiving MCA begins to record information about the activity.
- The value of the ACTIVREC queue manager attribute on QM2 is now QUEUE, therefore the MCA generates an activity report, but puts it on the system queue (SYSTEM.ADMIN.ACTIVITY.QUEUE) and not on the reply queue (ACTIV.REPLY.Q).

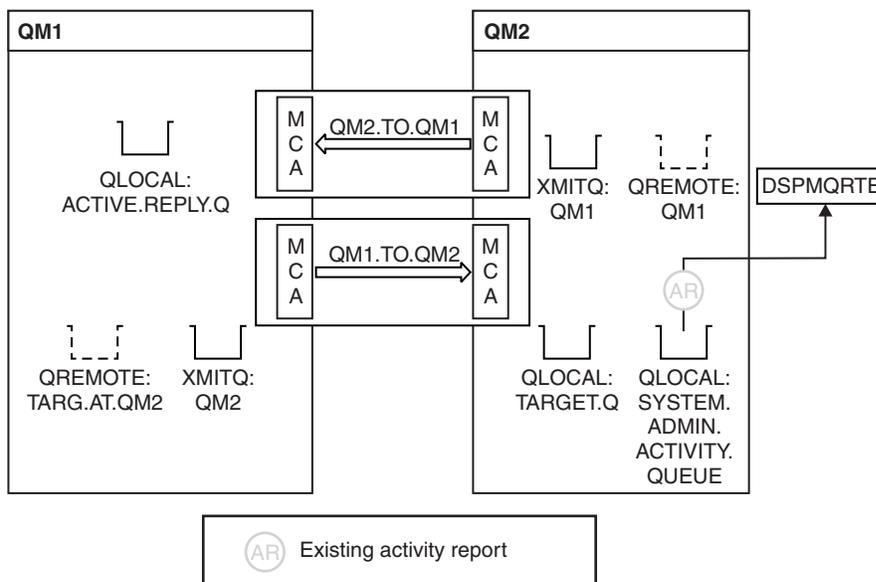


Figure 126. Delivering activity reports to the system queue, Diagram 2

- Meanwhile DSPMQRTE has been waiting for activity reports to arrive on ACTIV.REPLY.Q. Only two arrive. DSPMQRTE continues waiting for 120 seconds because it seems that the route is not yet complete.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2 -rq
 ACTIV.REPLY.Q -v outline identifiers'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2', queue
 manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
```

```

Activity:
 ApplName: 'cann\output\bin\dspmqrte.exe'
```

```
Operation:
 OperationType: Put
```

```
Message:
```

```
MQMD:
 MsgId: X'414D51204C41524745512020202020A3C9154220001502'
 CorrelId: X'414D51204C41524745512020202020A3C9154220001503'
 QMgrName: 'QM1'
 QName: 'TARG.AT.QM2'
 ResolvedQName: 'QM2'
 RemoteQName: 'TARGET.Q'
 RemoteQMgrName: 'QM2'
```

```

Activity:
 ApplName: 'cann\output\bin\runmqchl.EXE'
```

```
Operation:
 OperationType: Get
```

```
Message:
```

```
MQMD:
 MsgId: X'414D51204C41524745512020202020A3C9154220001505'
 CorrelId: X'414D51204C41524745512020202020A3C9154220001502'
```

```
EmbeddedMQMD:
 MsgId: X'414D51204C41524745512020202020A3C9154220001502'
 CorrelId: X'414D51204C41524745512020202020A3C9154220001503'
```

```
QMgrName: 'QM1'
QName: 'QM2'
ResolvedQName: 'QM2'
```

```
Operation:
 OperationType: Send
```

```
Message:
```

```
MQMD:
 MsgId: X'414D51204C41524745512020202020A3C9154220001502'
 CorrelId: X'414D51204C41524745512020202020A3C9154220001503'
 QMgrName: 'QM1'
 RemoteQMgrName: 'QM2'
 ChannelName: 'QM1.TO.QM2'
 ChannelType: Sender
 XmitQName: 'QM2'
```

```

AMQ8652: DSPMQRTE command has finished.
```

- The last operation that DSPMQRTE observed was a Send, therefore the channel is running. Now we must work out why we did not receive any more activity reports from queue manager QM2 (as identified in RemoteQMgrName).
- To check whether there is any activity information on the system queue, start DSPMQRTE on QM2 to try and collect more activity reports. Use the following command to start DSPMQRTE:

```
dspmqrte -m QM2 -q SYSTEM.ADMIN.ACTIVITY.QUEUE
 -i 414D51204C41524745512020202020A3C9154220001502 -v outline
```

where 414D51204C41524745512020202020A3C9154220001502 is the MsgId of the trace-route message that was put.

- DSPMQRTE then performs a sequence of MQGETs again, waiting for responses on the system activity queue related to the trace-route message with the specified identifier.
- DSPMQRTE gets one more activity report, which it displays. DSPMQRTE determines that the preceding activity reports are missing, and displays a message saying this. We already know about this part of the route, however.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM2
 -q SYSTEM.ADMIN.ACTIVITY.QUEUE
 -i 414D51204C41524745512020202020A3C915420001502 -v outline'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
```

```

Activity:
Activity information unavailable.
```

```

Activity:
ApplName: 'cann\output\bin\AMQRMPPA.EXE'

Operation:
OperationType: Receive
QMgrName: 'QM2'
RemoteQMgrName: 'QM1'
ChannelName: 'QM1.TO.QM2'
ChannelType: Receiver
```

```
Operation:
OperationType: Discard
QMgrName: 'QM2'
QName: 'TARGET.Q'
Feedback: NotDelivered
```

```

AMQ8652: DSPMQRTE command has finished.
```

- This activity report indicates that the route information is now complete. No problem occurred.
- Just because route information is unavailable, or because DSPMQRTE cannot display all of the route, this does not mean that the message was not delivered. For example, the queue manager attributes of different queue managers might be different, or a reply queue might not be defined to get the response back.

#### Example 4 - Diagnosing a channel problem:

Diagnose a problem in which the trace-route message does not reach the target queue

In this example the WebSphere MQ display route application connects to queue manager, QM1, generates a trace-route message, then attempts to deliver it to the target queue, TARGET.Q, on remote queue manager, QM2. In this example the trace-route message does not reach the target queue. The available

activity report is used to diagnose the problem.

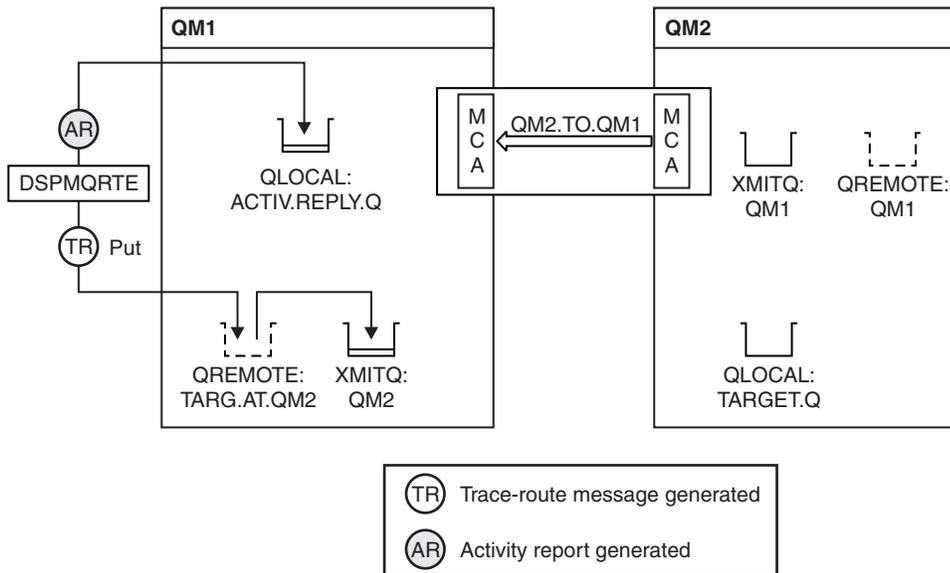


Figure 127. Diagnosing a channel problem

- In this example, the channel QM1.TO.QM2 is not running.
- DSPMQRTE puts a trace-route message (as in example 1) to the target queue and generates an activity report.
- There is no MCA to get the message from the transmission queue (QM2), therefore this is the only activity report that DSPMQRTE gets back from the reply queue. This time the fact that the route is not complete does indicate a problem. The administrator can use the transmission queue found in ResolvedQName to investigate why the transmission queue is not being serviced.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2
-rq ACTIV.REPLY.Q -v outline'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2',
queue manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
```

-----  
Activity:

App1Name: 'cann\output\bin\dspmqrte.exe'

Operation:

```
OperationType: Put
QMgrName: 'QM1'
QName: 'TARG.AT.QM2'
ResolvedQName: 'QM2'
RemoteQName: 'TARGET.Q'
RemoteQMgrName: 'QM2'
```

-----  
AMQ8652: DSPMQRTE command has finished.

## Activity report reference

Use this page to obtain an overview of the activity report message format. The activity report message data contains the parameters that describe the activity.

**Related concepts:**

“Activity report format”

**Related reference:**

“Activity report MQMD (message descriptor)” on page 928

“Activity report MQEPH (Embedded PCF header)” on page 933

“Activity report MQCFH (PCF header)” on page 934

“Activity report message data” on page 935

“Operation-specific activity report message data” on page 947

**Activity report format**

Activity reports are standard WebSphere MQ report messages containing a message descriptor and message data. Activity reports are PCF messages generated by applications that have performed an activity on behalf of a message as it has been routed through a queue manager network.

Activity reports contain the following information:

**A message descriptor**

An MQMD structure

**Message data**

Consists of the following:

- An embedded PCF header (MQEPH).
- Activity report message data.

Activity report message data consists of the *Activity* PCF group and, if generated for a trace-route message, the *TraceRoute* PCF group.

Table 106 on page 928 shows the structure of these reports, including parameters that are returned only under certain conditions.

Table 106. Activity report format

| Message descriptor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Message data                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQMD structure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Embedded PCF header MQEPH structure                                                                                                                                                                                                                                                                                                      | Activity report message data                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Structure identifier<br>Structure version<br>Report options<br>Message type<br>Expiration time<br>Feedback<br>Encoding<br>Coded character set ID<br>Message format<br>Priority<br>Persistence<br>Message identifier<br>Correlation identifier<br>Backout count<br>Reply-to queue<br>Reply-to queue manager<br>User identifier<br>Accounting token<br>Application identity data<br>Application type<br>Application name<br>Put date<br>Put time<br>Application origin data<br>Group identifier<br>Message sequence number<br>Offset<br>Message flags<br>Original length | Structure identifier<br>Structure version<br>Structure length<br>Encoding<br>Coded character set ID<br>Message format<br>Flags<br>PCF header (MQCFH)<br>Structure type<br>Structure length<br>Structure version<br>Command identifier<br>Message sequence number<br>Control options<br>Completion code<br>Reason code<br>Parameter count | Activity<br>Activity application name<br>Activity application type<br>Activity description<br>Operation<br>Operation type<br>Operation date<br>Operation time<br>Message<br>Message length<br>MQMD <sup>8</sup><br>EmbeddedMQMD<br>Queue manager name<br>Queue sharing group name<br>Queue name <sup>1 2 3 7</sup><br>Resolved queue name <sup>1 3 7</sup><br>Remote queue name <sup>3 7</sup><br>Remote queue manager name <sup>2 3 4 5 7</sup><br>Subscription level <sup>9</sup><br>Subscription identifier <sup>9</sup><br>Feedback <sup>2 10</sup><br>Channel name <sup>4 5</sup><br>Channel type <sup>4 5</sup><br>Transmission queue name <sup>5</sup><br>TraceRoute <sup>6</sup><br>Detail<br>Recorded activities<br>Unrecorded activities<br>Discontinuity count<br>Max activities<br>Accumulate<br>Deliver |

**Note:**

1. Returned for Get and Browse operations.
2. Returned for Discard operations.
3. Returned for Put, Put Reply, and Put Report operations.
4. Returned for Receive operations.
5. Returned for Send operations.
6. Returned for trace-route messages.
7. Not returned for Put operations to a topic, contained within Publish activities.
8. Not returned for Excluded Publish operations. For Publish and Discarded Publish operations, returned containing a subset of parameters.
9. Returned for Publish, Discarded Publish, and Excluded Publish operations.
10. Returned for Discarded Publish and Excluded Publish operations.

**Activity report MQMD (message descriptor)**

Use this page to view the values contained by the MQMD structure for an activity report

### *StrucId*

Description: Structure identifier.  
Data type: MQCHAR4.  
Value: MQMD\_STRUC\_ID.

### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: Copied from the original message descriptor. Possible values are:

**MQMD\_VERSION\_1**

Version-1 message descriptor structure, supported in all environments.

**MQMD\_VERSION\_2**

Version-2 message descriptor structure, supported on AIX, HP-UX, z/OS, HP OpenVMS, IBM i, Solaris, Linux, Windows, and all WebSphere MQ MQI clients connected to these systems.

### *Report*

Description: Options for further report messages.  
Data type: MQLONG.  
Value: If MQRO\_PASS\_DISCARD\_AND\_EXPIRY or MQRO\_DISCARD\_MSG were specified in the *Report* field of the original message descriptor:

**MQRO\_DISCARD**

The report is discarded if it cannot be delivered to the destination queue.

Otherwise:

**MQRO\_NONE**

No reports required.

### *MsgType*

Description: Indicates type of message.  
Data type: MQLONG.  
Value: MQMT\_REPORT.

### *Expiry*

Description: Report message lifetime.  
Data type: MQLONG.  
Value: If the *Report* field in the original message descriptor is specified as MQRO\_PASS\_DISCARD\_AND\_EXPIRY, the remaining expiry time from the original message is used.

Otherwise:

**MQEI\_UNLIMITED**

The report does not have an expiry time.

### *Feedback*

Description: Feedback or reason code.  
Data type: MQLONG.  
Value: **MQFB\_ACTIVITY**  
Activity report.

#### *Encoding*

Description: Numeric encoding of report message data.  
Data type: MQLONG.  
Value: MQENC\_NATIVE.

#### *CodedCharSetId*

Description: Character set identifier of report message data.  
Data type: MQLONG.  
Value: Set as appropriate.

#### *Format*

Description: Format name of report message data  
Data type: MQCHAR8.  
Value: **MQFMT\_EMBEDDED\_PCF**  
Embedded PCF message.

#### *Priority*

Description: Report message priority.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

#### *Persistence*

Description: Report message persistence.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

#### *MsgId*

Description: Message identifier.  
Data type: MQBYTE24.  
Values: If the *Report* field in the original message descriptor is specified as MQRO\_PASS\_MSG\_ID, the message identifier from the original message is used.  
  
Otherwise, a unique value will be generated by the queue manager.

#### *CorrelId*

Description: Correlation identifier.  
Data type: MQBYTE24.  
Value: If the *Report* field in the original message descriptor is specified as MQRO\_PASS\_CORREL\_ID, the correlation identifier from the original message is used. Otherwise, the message identifier is copied from the original message.

#### *BackoutCount*

Description: Backout counter.  
Data type: MQLONG.  
Value: 0.

#### *ReplyToQ*

Description: Name of reply queue.  
Data type: MQCHAR48.  
Values: Blank.

#### *ReplyToQMgr*

Description: Name of reply queue manager.  
Data type: MQCHAR48.  
Value: The queue manager name that generated the report message.

#### *UserIdentifier*

Description: The user identifier of the application that generated the report message.  
Data type: MQCHAR12.  
Value: Copied from the original message descriptor.

#### *AccountingToken*

Description: Accounting token that allows an application to charge for work done as a result of the message.  
Data type: MQBYTE32.  
Value: Copied from the original message descriptor.

#### *ApplIdentityData*

Description: Application data relating to identity.  
Data type: MQCHAR32.  
Values: Copied from the original message descriptor.

#### *PutApplType*

Description: Type of application that put the report message.  
Data type: MQLONG.  
Value: **MQAT\_QMGR**  
Queue manager generated message.

#### *PutAppName*

Description: Name of application that put the report message.  
Data type: MQCHAR28.  
Value: Either the first 28 bytes of the queue manager name, or the name of the MCA that generated the report message.

#### *PutDate*

Description: Date when message was put.  
Data type: MQCHAR8.  
Value: As generated by the queue manager.

#### *PutTime*

Description: Time when message was put.  
Data type: MQCHAR8.  
Value: As generated by the queue manager.

#### *ApplOriginData*

Description: Application data relating to origin.  
Data type: MQCHAR4.  
Value: Blank.

If *Version* is MQMD\_VERSION\_2, the following additional fields are present:

#### *GroupId*

Description: Identifies to which message group or logical message the physical message belongs.  
Data type: MQBYTE24.  
Value: Copied from the original message descriptor.

#### *MsgSeqNumber*

Description: Sequence number of logical message within group.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

#### *Offset*

Description: Offset of data in physical message from start of logical message.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

### *MsgFlags*

Description: Message flags that specify attributes of the message or control its processing.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

### *OriginalLength*

Description: Length of original message.  
Data type: MQLONG.  
Value: Copied from the original message descriptor.

## **Activity report MQEPH (Embedded PCF header)**

Use this page to view the values contained by the MQEPH structure for an activity report

The MQEPH structure contains a description of both the PCF information that accompanies the message data of an activity report, and the application message data that follows it.

For an activity report, the MQEPH structure contains the following values:

### *StrucId*

Description: Structure identifier.  
Data type: MQCHAR4.  
Value: MQEPH\_STRUC\_ID.

### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: MQEPH\_VERSION\_1.

### *StrucLength*

Description: Structure length.  
Data type: MQLONG.  
Value: Total length of the structure including the PCF parameter structures that follow it.

### *Encoding*

Description: Numeric encoding of the message data that follows the last PCF parameter structure.  
Data type: MQLONG.  
Value: If any data from the original application message data is included in the report message, the value will be copied from the *Encoding* field of the original message descriptor.  
  
Otherwise, 0.

### *CodedCharSetId*

Description: Character set identifier of the message data that follows the last PCF parameter structure.  
Data type: MQLONG.  
Value: If any data from the original application message data is included in the report message, the value will be copied from the *CodedCharSetId* field of the original message descriptor.  
Otherwise, MQCCSI\_UNDEFINED.

#### *Format*

Description: Format name of message data that follows the last PCF parameter structure.  
Data type: MQCHAR8.  
Value: If any data from the original application message data is included in the report message, the value will be copied from the *Format* field of the original message descriptor.  
Otherwise, MQFMT\_NONE.

#### *Flags*

Description: Flags that specify attributes of the structure or control its processing.  
Data type: MQLONG.  
Value: **MQEPH\_CCSID\_EMBEDDED**  
Specifies that the character set of the parameters containing character data is specified individually within the *CodedCharSetId* field in each structure.

#### *PCFHeader*

Description: Programmable Command Format Header  
Data type: MQCFH.  
Value: See "Activity report MQCFH (PCF header)."

### **Activity report MQCFH (PCF header)**

Use this page to view the PCF values contained by the MQCFH structure for an activity report

For an activity report, the MQCFH structure contains the following values:

#### *Type*

Description: Structure type that identifies the content of the report message.  
Data type: MQLONG.  
Value: **MQCFT\_REPORT**  
Message is a report.

#### *StrucLength*

Description: Structure length.  
Data type: MQLONG.  
Value: **MQCFH\_STRUC\_LENGTH**  
Length in bytes of MQCFH structure.

#### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: MQCFH\_VERSION\_3

#### *Command*

Description: Command identifier. This identifies the category of the message.  
Data type: MQLONG.  
Values: **MQCMD\_ACTIVITY\_MSG**  
Message activity.

#### *MsgSeqNumber*

Description: Message sequence number. This is the sequence number of the message within a group of related messages.  
Data type: MQLONG.  
Values: 1.

#### *Control*

Description: Control options.  
Data type: MQLONG.  
Values: MQCFC\_LAST.

#### *CompCode*

Description: Completion code.  
Data type: MQLONG.  
Values: MQCC\_OK.

#### *Reason*

Description: Reason code qualifying completion code.  
Data type: MQLONG.  
Values: MQRC\_NONE.

#### *ParameterCount*

Description: Count of parameter structures. This is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.  
Data type: MQLONG.  
Values: 1 or greater.

### **Activity report message data**

Use this page to view the parameters contained by the *Activity* PCF group in an activity report message. Some parameters are returned only when specific operations have been performed.

Activity report message data consists of the *Activity* PCF group and, if generated for a trace-route message, the *TraceRoute* PCF group. The *Activity* PCF group is detailed in this topic.

Some parameters, which are described as Operation-specific activity report message data, are returned only when specific operations have been performed.

For an activity report, the activity report message data contains the following parameters:

### *Activity*

Description: Grouped parameters describing the activity.  
Identifier: MQGACF\_ACTIVITY.  
Data type: MQCFGR.  
Included in PCF group: None.  
Parameters in PCF group: *ActivityApplName*  
*ActivityApplType*  
*ActivityDescription*  
*Operation*  
*TraceRoute*  
Returned: Always.

### *ActivityApplName*

Description: Name of application that performed the activity.  
Identifier: MQCACF\_APPL\_NAME.  
Data type: MQCFST.  
Included in PCF group: *Activity*.  
Maximum length: MQ\_APPL\_NAME\_LENGTH.  
Returned: Always.

### *ActivityApplType*

Description: Type of application that performed the activity.  
Identifier: MQIA\_APPL\_TYPE.  
Data type: MQCFIN.  
Included in PCF group: *Activity*.  
Returned: Always.

### *ActivityDescription*

Description: Description of activity performed by the application.  
Identifier: MQCACF\_ACTIVITY\_DESCRIPTION.  
Data type: MQCFST.  
Included in PCF group: *Activity*.  
Maximum length: 64  
Returned: Always.

### *Operation*

Description: Grouped parameters describing an operation of the activity.  
Identifier: MQGACF\_OPERATION.  
Data type: MQCFGR.  
Included in PCF group: *Activity*.

Parameters in PCF group:  
*OperationType*  
*OperationDate*  
*OperationTime*  
*Message*  
*QMgrName*  
*QSGName*

**Note:** Additional parameters are returned in this group depending on the operation type. These additional parameters are described as Operation-specific activity report message data.  
Returned: One *Operation* PCF group per operation in the activity.

### *OperationType*

Description: Type of operation performed.  
Identifier: MQIACF\_OPERATION\_TYPE.  
Data type: MQCFIN.  
Included in PCF group: *Operation*.  
Values: MQOPER\_\*.  
Returned: Always.

### *OperationDate*

Description: Date when the operation was performed.  
Identifier: MQCACF\_OPERATION\_DATE.  
Data type: MQCFST.  
Included in PCF group: *Operation*.  
Maximum length: MQ\_DATE\_LENGTH.  
Returned: Always.

### *OperationTime*

Description: Time when the operation was performed.  
Identifier: MQCACF\_OPERATION\_TIME.  
Data type: MQCFST.  
Included in PCF group: *Operation*.  
Maximum length: MQ\_TIME\_LENGTH.  
Returned: Always.

### *Message*

Description: Grouped parameters describing the message that caused the activity.  
Identifier: MQGACF\_MESSAGE.  
Data type: MQCFGR.  
Included in PCF group: *Operation.*  
Parameters in group: *MsgLength*  
*MQMD*  
*EmbeddedMQMD*  
Returned: Always, except for Excluded Publish operations.

### *MsgLength*

Description: Length of the message that caused the activity, before the activity occurred.  
Identifier: MQIACF\_MSG\_LENGTH.  
Data type: MQCFIN.  
Included in PCF group: *Message.*  
Returned: Always.

### *MQMD*

Description: Grouped parameters related to the message descriptor of the message that caused the activity.  
Identifier: MQGACF\_MQMD.  
Data type: MQCFGR.  
Included in PCF group: *Message.*

Parameters in group:

- StrucId*
- Version*
- Report*
- MsgType*
- Expiry*
- Feedback*
- Encoding*
- CodedCharSetId*
- Format*
- Priority*
- Persistence*
- MsgId*
- CorrelId*
- BackoutCount*
- ReplyToQ*
- ReplyToQMgr*
- UserIdentifier*
- AccountingToken*
- ApplIdentityData*
- PutApplType*
- PutApplName*
- PutDate*
- PutTime*
- ApplOriginData*
- GroupId*
- MsgSeqNumber*
- Offset*
- MsgFlags*
- OriginalLength*

Returned: Always, except for Excluded Publish operations.

*EmbeddedMQMD*

Description: Grouped parameters describing the message descriptor embedded within a message on a transmission queue.

Identifier: MQGACF\_EMBEDDED\_MQMD.

Data type: MQCFGR.

Included in PCF group: *Message*.

Parameters in group:

- StrucId*
- Version*
- Report*
- MsgType*
- Expiry*
- Feedback*
- Encoding*
- CodedCharSetId*
- Format*
- Priority*
- Persistence*
- MsgId*
- CorrelId*
- BackoutCount*
- ReplyToQ*
- ReplyToQMgr*
- UserIdentifier*
- AccountingToken*
- ApplIdentityData*
- PutApplType*
- PutApplName*
- PutDate*
- PutTime*
- ApplOriginData*
- GroupId*
- MsgSeqNumber*
- Offset*
- MsgFlags*
- OriginalLength*

Returned: For Get operations where the queue resolves to a transmission queue.

### *StrucId*

Description: Structure identifier  
 Identifier: MQCACF\_STRUC\_ID.  
 Data type: MQCFST.  
 Included in PCF group: MQMD or *EmbeddedMQMD*.  
 Maximum length: 4.  
 Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Version*

Description: Structure version number.  
Identifier: MQIACF\_VERSION.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Report*

Description: Options for report messages.  
Identifier: MQIACF\_REPORT.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *MsgType*

Description: Indicates type of message.  
Identifier: MQIACF\_MSG\_TYPE.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Expiry*

Description: Message lifetime.  
Identifier: MQIACF\_EXPIRY.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Feedback*

Description: Feedback or reason code.  
Identifier: MQIACF\_FEEDBACK.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Encoding*

Description: Numeric encoding of message data.  
Identifier: MQIACF\_ENCODING.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *CodedCharSetId*

Description: Character set identifier of message data.  
Identifier: MQIA\_CODED\_CHAR\_SET\_ID.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *Format*

Description: Format name of message data  
Identifier: MQCACH\_FORMAT\_NAME.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_FORMAT\_LENGTH.  
Returned: Always, except for Excluded Publish operations.

### *Priority*

Description: Message priority.  
Identifier: MQIACF\_PRIORITY.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations.

### *Persistence*

Description: Message persistence.  
Identifier: MQIACF\_PERSISTENCE.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish operations.

### *MsgId*

Description: Message identifier.  
Identifier: MQBACF\_MSG\_ID.  
Data type: MQCFBS.  
Included in PCF group: MQMD or EmbeddedMQMD.  
Maximum length: MQ\_MSG\_ID\_LENGTH.  
Returned: Always, except for Excluded Publish operations.

### *CorrelId*

Description: Correlation identifier.  
Identifier: MQBACF\_CORREL\_ID.  
Data type: MQCFBS.  
Included in PCF group: MQMD or EmbeddedMQMD.  
Maximum length: MQ\_CORREL\_ID\_LENGTH.  
Returned: Always, except for Excluded Publish operations.

### *BackoutCount*

Description: Backout counter.  
Identifier: MQIACF\_BACKOUT\_COUNT.  
Data type: MQCFIN.  
Included in PCF group: MQMD or EmbeddedMQMD.  
Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

### *ReplyToQ*

Description: Name of reply queue.  
Identifier: MQCACF\_REPLY\_TO\_QUEUE.  
Data type: MQCFST.  
Included in PCF group: MQMD or EmbeddedMQMD.  
Maximum length: MQ\_Q\_NAME\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish operations.

### *ReplyToQMgr*

Description: Name of reply queue manager.  
Identifier: MQCACF\_REPLY\_TO\_Q\_MGR.  
Data type: MQCFST.  
Included in PCF group: MQMD or EmbeddedMQMD.  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *UserIdentifier*

Description: The user identifier of the application that originated the message.  
Identifier: MQCACF\_USER\_IDENTIFIER.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_USER\_ID\_LENGTH.  
Returned: Always, except for Excluded Publish Operations.

### *AccountingToken*

Description: Accounting token that allows an application to charge for work done as a result of the message.  
Identifier: MQBACF\_ACCOUNTING\_TOKEN.  
Data type: MQCFBS.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_ACCOUNTING\_TOKEN\_LENGTH.  
Returned: Always, except for Excluded Publish Operations.

### *ApplIdentityData*

Description: Application data relating to identity.  
Identifier: MQCACF\_APPL\_IDENTITY\_DATA.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_APPL\_IDENTITY\_DATA\_LENGTH.  
Returned: Always, except for Excluded Publish Operations.

### *PutApplType*

Description: Type of application that put the message.  
Identifier: MQIA\_APPL\_TYPE.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *PutApplName*

Description: Name of application that put the message.  
Identifier: MQCACF\_APPL\_NAME.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_APPL\_NAME\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *PutDate*

Description: Date when message was put.  
Identifier: MQCACF\_PUT\_DATE.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_PUT\_DATE\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

#### *PutTime*

Description: Time when message was put.  
Identifier: MQCACF\_PUT\_TIME.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_PUT\_TIME\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

#### *ApplOriginData*

Description: Application data relating to origin.  
Identifier: MQCACF\_APPL\_ORIGIN\_DATA.  
Data type: MQCFST.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_APPL\_ORIGIN\_DATA\_LENGTH.  
Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

#### *GroupId*

Description: Identifies to which message group or logical message the physical message belongs.  
Identifier: MQBACF\_GROUP\_ID.  
Data type: MQCFBS.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Maximum length: MQ\_GROUP\_ID\_LENGTH.  
Returned: If the *Version* is specified as MQMD\_VERSION\_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

#### *MsgSeqNumber*

Description: Sequence number of logical message within group.  
Identifier: MQIACH\_MSG\_SEQUENCE\_NUMBER.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: If *Version* is specified as MQMD\_VERSION\_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

#### *Offset*

Description: Offset of data in physical message from start of logical message.  
Identifier: MQIACF\_OFFSET.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: If *Version* is specified as MQMD\_VERSION\_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *MsgFlags*

Description: Message flags that specify attributes of the message or control its processing.  
Identifier: MQIACF\_MSG\_FLAGS.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: If *Version* is specified as MQMD\_VERSION\_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *OriginalLength*

Description: Length of original message.  
Identifier: MQIACF\_ORIGINAL\_LENGTH.  
Data type: MQCFIN.  
Included in PCF group: MQMD or *EmbeddedMQMD*.  
Returned: If *Version* is specified as MQMD\_VERSION\_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

### *QMgrName*

Description: Name of the queue manager where the activity was performed.  
Identifier: MQCA\_Q\_MGR\_NAME.  
Data type: MQCFST.  
Included in PCF group: *Operation*.  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
Returned: Always.

### *QSGName*

Description: Name of the queue-sharing group to which the queue manager where the activity was performed belongs.  
Identifier: MQCA\_QSG\_NAME.  
Data type: MQCFST.  
Included in PCF group: *Operation*.  
Maximum length: MQ\_QSG\_NAME\_LENGTH  
Returned: If the activity was performed on a WebSphere MQ for z/OS queue manager.

### *TraceRoute*

Description: Grouped parameters specifying attributes of the trace-route message.  
 Identifier: MQGACF\_TRACE\_ROUTE.  
 Data type: MQCFGR.  
 Contained in PCF group: *Activity*.  
 Parameters in group: *Detail*  
                           *RecordedActivities*  
                           *UnrecordedActivities*  
                           *DiscontinuityCount*  
                           *MaxActivities*  
                           *Accumulate*  
                           *Forward*  
                           *Deliver*

Returned: If the activity was performed on behalf of the trace-route message.

The values of the parameters in the *TraceRoute* PCF group are those from the trace-route message at the time the activity report was generated.

### Operation-specific activity report message data

Use this page to view the additional PCF parameters that might be returned in the PCF group *Operation* in an activity report, depending on the value of the *OperationType* parameter

The additional parameters vary depending on the following operation types:

#### Related reference:

“Get/Browse (MQOPER\_GET/MQOPER\_BROWSE)”

“Discard (MQOPER\_DISCARD)” on page 948

“Publish/Discarded Publish/Excluded Publish (MQOPER\_PUBLISH/MQOPER\_DISCARDED\_PUBLISH/MQOPER\_EXCLUDED\_PUBLISH)” on page 949

“Put/Put Reply/Put Report (MQOPER\_PUT/MQOPER\_PUT\_REPLY/MQOPER\_PUT\_REPORT)” on page 950

“Receive (MQOPER\_RECEIVE)” on page 951

“Send (MQOPER\_SEND)” on page 952

#### Get/Browse (MQOPER\_GET/MQOPER\_BROWSE):

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Get/Browse (MQOPER\_GET/MQOPER\_BROWSE) operation type (a message on a queue was got, or browsed).

*QName*

Description: The name of the queue that was opened.  
 Identifier: MQCA\_Q\_NAME.  
 Data type: MQCFST.  
 Included in PCF group: *Operation*.  
 Maximum length: MQ\_Q\_NAME\_LENGTH  
 Returned: Always.

#### *ResolvedQName*

Description: The name that the opened queue resolves to.  
 Identifier: MQCACF\_RESOLVED\_Q\_NAME.  
 Data type: MQCFST.  
 Included in PCF group: *Operation*.  
 Maximum length: MQ\_Q\_NAME\_LENGTH  
 Returned: Always.

#### **Discard (MQOPER\_DISCARD):**

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Discard (MQOPER\_DISCARD) operation type (a message was discarded).

#### *Feedback*

Description: The reason for the message being discarded.  
 Identifier: MQIACF\_FEEDBACK.  
 Data type: MQCFIN.  
 Included in PCF group: *Operation*.  
 Returned: Always.

#### *QName*

Description: The name of the queue that was opened.  
 Identifier: MQCA\_Q\_NAME.  
 Data type: MQCFST.  
 Maximum length: MQ\_Q\_NAME\_LENGTH  
 Included in PCF group: *Operation*.  
 Returned: If the message was discarded because it was unsuccessfully put to a queue.

#### *RemoteQMgrName*

Description: The name of the queue manager to which the message was destined.  
 Identifier: MQCA\_REMOTE\_Q\_MGR\_NAME.  
 Data type: MQCFST.  
 Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
 Included in PCF group: *Operation*.  
 Returned: If the value of *Feedback* is MQFB\_NOT\_FORWARDED.

## **Publish/Discarded Publish/Excluded Publish (MQOPER\_PUBLISH/MQOPER\_DISCARDED\_PUBLISH/MQOPER\_EXCLUDED\_PUBLISH):**

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Publish/Discarded Publish/Excluded Publish (MQOPER\_PUBLISH/MQOPER\_DISCARDED\_PUBLISH/MQOPER\_EXCLUDED\_PUBLISH) operation type (a publish/subscribe message was delivered, discarded, or excluded).

### *SubId*

|                        |                              |
|------------------------|------------------------------|
| Description:           | The subscription identifier. |
| Identifier:            | MQBACF_SUB_ID.               |
| Data type:             | MQCFBS.                      |
| Included in PCF group: | <i>Operation</i> .           |
| Returned:              | Always.                      |

### *SubLevel*

|                        |                         |
|------------------------|-------------------------|
| Description:           | The subscription level. |
| Identifier:            | MQIACF_SUB_LEVEL.       |
| Data type:             | MQCFIN.                 |
| Included in PCF group: | <i>Operation</i> .      |
| Returned:              | Always.                 |

### *Feedback*

|                        |                                                                                                                                                  |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:           | The reason for discarding the message.                                                                                                           |
| Identifier:            | MQIACF_FEEDBACK.                                                                                                                                 |
| Data type:             | MQCFIN.                                                                                                                                          |
| Included in PCF group: | <i>Operation</i> .                                                                                                                               |
| Returned:              | If the message was discarded because it was not delivered to a subscriber, or the message was not delivered because the subscriber was excluded. |

The Publish operation MQOPER\_PUBLISH provides information about a message delivered to a particular subscriber. This operation describes the elements of the onward message that might have changed from the message described in the associated Put operation. Similarly to a Put operation, it contains a message group MQGACF\_MESSAGE and, inside that, an MQMD group MQGACF\_MQMD. However, this MQMD group contains only the following fields, which can be overridden by a subscriber: *Format, Priority, Persistence, MsgId, CorrelId, UserIdentifier, AccountingToken, ApplIdentityData*.

The *SubId* and *SubLevel* of the subscriber are included in the operation information. You can use the *SubID* with the MQCMD\_INQUIRE\_SUBSCRIBER PCF command to retrieve all other attributes for a subscriber.

The Discarded Publish operation MQOPER\_DISCARDED\_PUBLISH is analogous to the Discard operation that is used when a message is not delivered in point-to-point messaging. A message is not delivered to a subscriber if the message was explicitly requested not to be delivered to a local destination and this subscriber specifies a local destination. A message is also considered not delivered if there is a problem getting the message to the destination queue, for example, because the queue is full.

The information in a Discarded Publish operation is the same as for a Publish operation, with the addition of a *Feedback* field that gives the reasons why the message was not delivered. This feedback field contains MQFB\_\* or MQRC\_\* values that are common with the MQOPER\_DISCARD operation. The

reason for discarding a publish, as opposed to excluding it, are the same as the reasons for discarding a put.

The Excluded Publish operation `MQOPER_EXCLUDED_PUBLISH` provides information about a subscriber that was considered for delivery of the message, because the topic on which the subscriber is subscribing matches that of the associated Put operation, but the message was not delivered to the subscriber because other selection criteria do not match with the message that is being put to the topic. As with a Discarded Publish operation, the *Feedback* field provides information about the reason why this subscription was excluded. However, unlike the Discarded Publish operation, no message-related information is provided because no message was generated for this subscriber.

### **Put/Put Reply/Put Report (MQOPER\_PUT/MQOPER\_PUT\_REPLY/MQOPER\_PUT\_REPORT):**

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Put/Put Reply/Put Report (`MQOPER_PUT/MQOPER_PUT_REPLY/MQOPER_PUT_REPORT`) operation type (a message, reply message, or report message was put to a queue).

#### *QName*

|                        |                                                                                                                         |
|------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Description:           | The name of the queue that was opened.                                                                                  |
| Identifier:            | <code>MQCA_Q_NAME</code> .                                                                                              |
| Data type:             | <code>MQCFST</code> .                                                                                                   |
| Included in PCF group: | <i>Operation</i> .                                                                                                      |
| Maximum length:        | <code>MQ_Q_NAME_LENGTH</code>                                                                                           |
| Returned:              | Always, apart from one exception: not returned if the Put operation is to a topic, contained within a publish activity. |

#### *ResolvedQName*

|                        |                                                                                                                                |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Description:           | The name that the opened queue resolves to.                                                                                    |
| Identifier:            | <code>MQCACF_RESOLVED_Q_NAME</code> .                                                                                          |
| Data type:             | <code>MQCFST</code> .                                                                                                          |
| Included in PCF group: | <i>Operation</i> .                                                                                                             |
| Maximum length:        | <code>MQ_Q_NAME_LENGTH</code>                                                                                                  |
| Returned:              | When the opened queue could be resolved. Not returned if the Put operation is to a topic, contained within a publish activity. |

#### *RemoteQName*

|                        |                                                                                                                              |
|------------------------|------------------------------------------------------------------------------------------------------------------------------|
| Description:           | The name of the opened queue, as it is known on the remote queue manager.                                                    |
| Identifier:            | <code>MQCA_REMOTE_Q_NAME</code> .                                                                                            |
| Data type:             | <code>MQCFST</code> .                                                                                                        |
| Included in PCF group: | <i>Operation</i> .                                                                                                           |
| Maximum length:        | <code>MQ_Q_NAME_LENGTH</code>                                                                                                |
| Returned:              | If the opened queue is a remote queue. Not returned if the Put operation is to a topic, contained within a publish activity. |

#### *RemoteQMgrName*

Description: The name of the remote queue manager on which the remote queue is defined.  
 Identifier: MQCA\_REMOTE\_Q\_MGR\_NAME.  
 Data type: MQCFST.  
 Included in PCF group: *Operation*.  
 Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
 Returned: If the opened queue is a remote queue. Not returned if the Put operation is to a topic, contained within a publish activity.

### *TopicString*

Description: The full topic string to which the message is being put.  
 Identifier: MQCA\_TOPIC\_STRING.  
 Data type: MQCFST.  
 Included in PCF group: *Operation*.  
 Returned: If the Put operation is to a topic, contained within a publish activity.

### *Feedback*

Description: The reason for the message being put on the dead-letter queue.  
 Identifier: MQIACF\_FEEDBACK.  
 Data type: MQCFIN.  
 Included in PCF group: *Operation*.  
 Returned: If the message was put on the dead-letter queue.

## **Receive (MQOPER\_RECEIVE):**

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Receive (MQOPER\_RECEIVE) operation type (a message was received on a channel).

### *ChannelName*

Description: The name of the channel on which the message was received.  
 Identifier: MQCACH\_CHANNEL\_NAME.  
 Data type: MQCFST.  
 Included in PCF group: *Operation*.  
 Maximum length: MQ\_CHANNEL\_NAME\_LENGTH  
 Returned: Always.

### *ChannelType*

Description: The type of channel on which the message was received.  
 Identifier: MQIACH\_CHANNEL\_TYPE.  
 Data type: MQCFIN.  
 Included in PCF group: *Operation*.  
 Returned: Always.

### *RemoteQMgrName*

|                        |                                                                    |
|------------------------|--------------------------------------------------------------------|
| Description:           | The name of the queue manager from which the message was received. |
| Identifier:            | MQCA_REMOTE_Q_MGR_NAME.                                            |
| Data type:             | MQCFST.                                                            |
| Included in PCF group: | <i>Operation.</i>                                                  |
| Maximum length:        | MQ_Q_MGR_NAME_LENGTH                                               |
| Returned:              | Always.                                                            |

### Send (MQOPER\_SEND):

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Send (MQOPER\_SEND) operation type (a message was sent on a channel).

#### *ChannelName*

|                        |                                                     |
|------------------------|-----------------------------------------------------|
| Description:           | The name of the channel where the message was sent. |
| Identifier:            | MQCACH_CHANNEL_NAME.                                |
| Data type:             | MQCFST.                                             |
| Included in PCF group: | <i>Operation.</i>                                   |
| Maximum length:        | MQ_CHANNEL_NAME_LENGTH.                             |
| Returned:              | Always.                                             |

#### *ChannelType*

|                        |                                                 |
|------------------------|-------------------------------------------------|
| Description:           | The type of channel where the message was sent. |
| Identifier:            | MQIACH_CHANNEL_TYPE.                            |
| Data type:             | MQCFIN.                                         |
| Included in PCF group: | <i>Operation.</i>                               |
| Returned:              | Always.                                         |

#### *XmitQName*

|                        |                                                              |
|------------------------|--------------------------------------------------------------|
| Description:           | The transmission queue from which the message was retrieved. |
| Identifier:            | MQCACH_XMIT_Q_NAME.                                          |
| Data type:             | MQCFST.                                                      |
| Included in PCF group: | <i>Operation.</i>                                            |
| Maximum length:        | MQ_Q_NAME_LENGTH.                                            |
| Returned:              | Always.                                                      |

#### *RemoteQMgrName*

|                        |                                                                     |
|------------------------|---------------------------------------------------------------------|
| Description:           | The name of the remote queue manager to which the message was sent. |
| Identifier:            | MQCA_REMOTE_Q_MGR_NAME.                                             |
| Data type:             | MQCFST.                                                             |
| Included in PCF group: | <i>Operation.</i>                                                   |
| Maximum length:        | MQ_Q_MGR_NAME_LENGTH                                                |
| Returned:              | Always.                                                             |

## Trace-route message reference

Use this page to obtain an overview of the trace-route message format. The trace-route message data includes parameters that describe the activities that the trace-route message has caused

**Related concepts:**

“Trace-route message format”

**Related reference:**

“Trace-route message MQMD (message descriptor)” on page 954

“Trace-route message MQEPH (Embedded PCF header)” on page 958

“Trace-route message MQCFH (PCF header)” on page 959

“Trace-route message data” on page 961

**Trace-route message format**

Trace-route messages are standard WebSphere MQ messages containing a message descriptor and message data. The message data contains information about the activities performed on a trace-route message as it has been routed through a queue manager network.

Trace-route messages contain the following information:

**A message descriptor**

An MQMD structure, with the *Format* field set to MQFMT\_ADMIN or MQFMT\_EMBEDDED\_PCF.

**Message data**

Consists of either:

- A PCF header (MQCFH) and trace-route message data, if *Format* is set to MQFMT\_ADMIN, or
- An embedded PCF header (MQEPH), trace-route message data, and additional user-specified message data, if *Format* is set to MQFMT\_EMBEDDED\_PCF.

When using the WebSphere MQ display route application to generate a trace-route message, *Format* is set to MQFMT\_ADMIN.

The content of the trace-route message data is determined by the *Accumulate* parameter from the *TraceRoute* PCF group, as follows:

- If *Accumulate* is set to MQROUTE\_ACCUMULATE\_NONE, the trace-route message data contains the *TraceRoute* PCF group.
- If *Accumulate* is set to either MQROUTE\_ACCUMULATE\_IN\_MSG or MQROUTE\_ACCUMULATE\_AND\_REPLY, the trace-route message data contains the *TraceRoute* PCF group and zero or more *Activity* PCF groups.

Table 107 on page 954 shows the structure of a trace-route message.

Table 107. Trace-route message format

| Message descriptor        | Message data                        |                          |
|---------------------------|-------------------------------------|--------------------------|
| MQMD structure            | Embedded PCF header MQEPH structure | Trace-route message data |
| Structure identifier      | Structure identifier                | TraceRoute               |
| Structure version         | Structure version                   | Detail                   |
| Report options            | Structure length                    | Recorded activities      |
| Message type              | Encoding                            | Unrecorded activities    |
| Expiration time           | Coded character set ID              | Discontinuity count      |
| Feedback                  | Message format                      | Max activities           |
| Encoding                  | Flags                               | Accumulate               |
| Coded character set ID    | PCF header (MQCFH)                  | Deliver                  |
| Message format            | Structure type                      |                          |
| Priority                  | Structure length                    |                          |
| Persistence               | Structure version                   |                          |
| Message identifier        | Command identifier                  |                          |
| Correlation identifier    | Message sequence number             |                          |
| Backout count             | Control options                     |                          |
| Reply-to queue            | Completion code                     |                          |
| Reply-to queue manager    | Reason code                         |                          |
| User identifier           | Parameter count                     |                          |
| Accounting token          |                                     |                          |
| Application identity data |                                     |                          |
| Application type          |                                     |                          |
| Application name          |                                     |                          |
| Put date                  |                                     |                          |
| Put time                  |                                     |                          |
| Application origin data   |                                     |                          |
| Group identifier          |                                     |                          |
| Message sequence number   |                                     |                          |
| Offset                    |                                     |                          |
| Message flags             |                                     |                          |
| Original length           |                                     |                          |

### Trace-route message MQMD (message descriptor)

Use this page to view the values contained by the MQMD structure for a trace-route message

#### *StrucId*

Description: Structure identifier.  
 Data type: MQCHAR4.  
 Value: MQMD\_STRUC\_ID.

#### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: **MQMD\_VERSION\_1.**

### *Report*

Description: Options for report messages.  
Data type: MQLONG.  
Value: Set according to requirements. Common report options follow:

#### **MQRO\_DISCARD\_MSG**

The message is discarded on arrival to a local queue.

#### **MQRO\_PASS\_DISCARD\_AND\_EXPIRY**

Every response (activity reports or trace-route reply message) will have the report option **MQRO\_DISCARD\_MSG** set, and the remaining expiry passed on. This ensures that responses do not remain in the queue manager network indefinitely.

### *MsgType*

Description: Type of message.  
Data type: MQLONG.  
Value: If the *Accumulate* parameter in the TraceRoute group is specified as **MQROUTE\_ACCUMULATE\_AND\_REPLY**, then message type is **MQMT\_REQUEST**

Otherwise:

**MQMT\_DATAGRAM.**

### *Expiry*

Description: Message lifetime.  
Data type: MQLONG.  
Value: Set according to requirements. This parameter can be used to ensure trace-route messages are not left in a queue manager network indefinitely.

### *Feedback*

Description: Feedback or reason code.  
Data type: MQLONG.  
Value: **MQFB\_NONE.**

### *Encoding*

Description: Numeric encoding of message data.  
Data type: MQLONG.  
Value: Set as appropriate.

#### *CodedCharSetId*

Description: Character set identifier of message data.  
Data type: MQLONG.  
Value: Set as appropriate.

#### *Format*

Description: Format name of message data  
Data type: MQCHAR8.  
Value: **MQFMT\_ADMIN**  
Admin message. No user data follows the *TraceRoute* PCF group.  
**MQFMT\_EMBEDDED\_PCF**  
Embedded PCF message. User data follows the *TraceRoute* PCF group.

#### *Priority*

Description: Message priority.  
Data type: MQLONG.  
Value: Set according to requirements.

#### *Persistence*

Description: Message persistence.  
Data type: MQLONG.  
Value: Set according to requirements.

#### *MsgId*

Description: Message identifier.  
Data type: MQBYTE24.  
Value: Set according to requirements.

#### *CorrelId*

Description: Correlation identifier.  
Data type: MQBYTE24.  
Value: Set according to requirements.

#### *BackoutCount*

Description: Backout counter.  
Data type: MQLONG.  
Value: 0.

#### *ReplyToQ*

Description: Name of reply queue.  
Data type: MQCHAR48.  
Values: Set according to requirements.

If *MsgType* is set to MQMT\_REQUEST or if *Report* has any report generating options set, then this parameter must be non-blank.

#### *ReplyToQMgr*

Description: Name of reply queue manager.  
Data type: MQCHAR48.  
Value: Set according to requirements.

#### *UserIdentifier*

Description: The user identifier of the application that originated the message.  
Data type: MQCHAR12.  
Value: Set as normal.

#### *AccountingToken*

Description: Accounting token that allows an application to charge for work done as a result of the message.  
Data type: MQBYTE32.  
Value: Set as normal.

#### *ApplIdentityData*

Description: Application data relating to identity.  
Data type: MQCHAR32.  
Values: Set as normal.

#### *PutApplType*

Description: Type of application that put the message.  
Data type: MQLONG.  
Value: Set as normal.

#### *PutApplName*

Description: Name of application that put the message.  
Data type: MQCHAR28.  
Value: Set as normal.

#### *PutDate*

Description: Date when message was put.  
Data type: MQCHAR8.  
Value: Set as normal.

#### *PutTime*

Description: Time when message was put.  
Data type: MQCHAR8.  
Value: Set as normal.

#### *ApplOriginData*

Description: Application data relating to origin.  
Data type: MQCHAR4.  
Value: Set as normal..

### **Trace-route message MQEPH (Embedded PCF header)**

Use this page to view the values contained by the MQEPH structure for a trace-route message

The MQEPH structure contains a description of both the PCF information that accompanies the message data of a trace-route message, and the application message data that follows it. An MQEPH structure is used only if additional user message data follows the TraceRoute PCF group.

For a trace-route message, the MQEPH structure contains the following values:

#### *StrucId*

Description: Structure identifier.  
Data type: MQCHAR4.  
Value: MQEPH\_STRUC\_ID.

#### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: MQEPH\_VERSION\_1.

#### *StrucLength*

Description: Structure length.  
Data type: MQLONG.  
Value: Total length of the structure including the PCF parameter structures that follow it.

#### *Encoding*

Description: Numeric encoding of the message data that follows the last PCF parameter structure.  
Data type: MQLONG.  
Value: The encoding of the message data.

#### *CodedCharSetId*

Description: Character set identifier of the message data that follows the last PCF parameter structure.  
Data type: MQLONG.  
Value: The character set of the message data.

#### *Format*

Description: Format name of the message data that follows the last PCF parameter structure.  
Data type: MQCHAR8.  
Value: The format name of the message data.

#### *Flags*

Description: Flags that specify attributes of the structure or control its processing.  
Data type: MQLONG.  
Value: **MQEPH\_NONE**  
No flags specified.  
**MQEPH\_CCSID\_EMBEDDED**  
Specifies that the character set of the parameters containing character data is specified individually within the *CodedCharSetId* field in each structure.

#### *PCFHeader*

Description: Programmable Command Format Header  
Data type: MQCFH.  
Value: See "Trace-route message MQCFH (PCF header)."

### **Trace-route message MQCFH (PCF header)**

Use this page to view the PCF values contained by the MQCFH structure for a trace-route message

For a trace-route message, the MQCFH structure contains the following values:

#### *Type*

Description: Structure type that identifies the content of the message.  
Data type: MQLONG.  
Value: **MQCFT\_TRACE\_ROUTE**  
Message is a trace-route message.

### *StrucLength*

Description: Structure length.  
Data type: MQLONG.  
Value: **MQCFH\_STRUC\_LENGTH**  
Length in bytes of MQCFH structure.

### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: **MQCFH\_VERSION\_3**

### *Command*

Description: Command identifier. This identifies the category of the message.  
Data type: MQLONG.  
Values: **MQCMD\_TRACE\_ROUTE**  
Trace-route message.

### *MsgSeqNumber*

Description: Message sequence number. This is the sequence number of the message within a group of related messages.  
Data type: MQLONG.  
Values: 1.

### *Control*

Description: Control options.  
Data type: MQLONG.  
Values: **MQCFC\_LAST**.

### *CompCode*

Description: Completion code.  
Data type: MQLONG.  
Values: **MQCC\_OK**.

### *Reason*

Description: Reason code qualifying completion code.  
Data type: MQLONG.  
Values: MQRC\_NONE.

#### *ParameterCount*

Description: Count of parameter structures. This is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.  
Data type: MQLONG.  
Values: 1 or greater.

### **Trace-route message data**

Use this page to view the parameters that make up the *TraceRoute* PCF group part of trace-route message data

The content of trace-route message data depends on the *Accumulate* parameter from the *TraceRoute* PCF group. Trace-route message data consists of the *TraceRoute* PCF group, and zero or more *Activity* PCF groups. The *TraceRoute* PCF group is detailed in this topic. Refer to the related information for details of the *Activity* PCF group.

Trace-route message data contains the following parameters:

#### *TraceRoute*

Description: Grouped parameters specifying attributes of the trace-route message. For a trace-route message, some of these parameters can be altered to control how it is processed.  
Identifier: MQGACF\_TRACE\_ROUTE.  
Data type: MQCFGR.  
Contained in PCF group: None.  
Parameters in group: *Detail*  
*RecordedActivities*  
*UnrecordedActivities*  
*DiscontinuityCount*  
*MaxActivities*  
*Accumulate*  
*Forward*  
*Deliver*

#### *Detail*

Description: The detail level that will be recorded for the activity.  
Identifier: MQIACF\_ROUTE\_DETAIL.  
Data type: MQCFIN.  
Contained in PCF group: *TraceRoute*.

Values:

**MQRROUTE\_DETAIL\_LOW**

Activities performed by user-written application are recorded.

**MQRROUTE\_DETAIL\_MEDIUM**

Activities specified in MQRROUTE\_DETAIL\_LOW are recorded. Additionally, activities performed by MCAs are recorded.

**MQRROUTE\_DETAIL\_HIGH**

Activities specified in MQRROUTE\_DETAIL\_LOW, and MQRROUTE\_DETAIL\_MEDIUM are recorded. MCAs do not record any further activity information at this level of detail. This option is only available to user-written applications that are to record further activity information.

*RecordedActivities*

Description: The number of activities that the trace-route message has caused, where information was recorded.  
Identifier: MQIACF\_RECORDED\_ACTIVITIES.  
Data type: MQCFIN.  
Contained in PCF group: *TraceRoute*.

*UnrecordedActivities*

Description: The number of activities that the trace-route message has caused, where information was not recorded.  
Identifier: MQIACF\_UNRECORDED\_ACTIVITIES.  
Data type: MQCFIN.  
Contained in PCF group: *TraceRoute*.

*DiscontinuityCount*

Description: The number of times a trace-route message has been received from a queue manager that does not support trace-route messaging.  
Identifier: MQIACF\_DISCONTINUITY\_COUNT.  
Data type: MQCFIN.  
Contained in PCF group: *TraceRoute*.

*MaxActivities*

Description: The maximum number of activities the trace-route message can be involved in before it stops being processed.  
Identifier: MQIACF\_MAX\_ACTIVITIES.  
Data type: MQCFIN.  
Contained in PCF group: *TraceRoute*.  
Value:

**A positive integer**

The maximum number of activities.

**MQRROUTE\_UNLIMITED\_ACTIVITIES**

An unlimited number of activities.

*Accumulate*

Description: Specifies whether activity information is accumulated within the trace-route message, and whether a reply message containing the accumulated activity information is generated before the trace-route message is discarded or is put on a non-transmission queue.

Identifier: MQIACF\_ROUTE\_ACCUMULATION.

Data type: MQCFIN.

Contained in PCF group: *TraceRoute*.

Value:

**MQROUTE\_ACCUMULATE\_NONE**  
Activity information is not accumulated in the message data of the trace-route message.

**MQROUTE\_ACCUMULATE\_IN\_MSG**  
Activity information is accumulated in the message data of the trace-route message.

**MQROUTE\_ACCUMULATE\_AND\_REPLY**  
Activity information is accumulated in the message data of the trace-route message, and a trace-route reply message will be generated.

### *Forward*

Description: Specifies queue managers that the trace-route message can be forwarded to. When determining whether to forward a message to a remote queue manager, queue managers use the algorithm that is described in Forwarding.

Identifier: MQIACF\_ROUTE\_FORWARDING.

Data type: MQCFIN.

Contained in PCF group: *TraceRoute*.

Value:

**MQROUTE\_FORWARD\_IF\_SUPPORTED**  
The trace-route message is only forwarded to queue managers that will honor the value of the *Deliver* parameter from the *TraceRoute* group.

**MQROUTE\_FORWARD\_ALL**  
The trace-route message is forwarded to any queue manager, regardless of whether the value of the *Deliver* parameter will be honored.

### *Deliver*

Description: Specifies the action to be taken if the trace-route message arrives at the destination queue successfully.

Identifier: MQIACF\_ROUTE\_DELIVERY.

Data type: MQCFIN.

Contained in PCF group: *TraceRoute*.

Value:

**MQROUTE\_DELIVER\_YES**  
On arrival, the trace-route message is put on the target queue. Any application performing a destructive get on the target queue can receive the trace-route message.

**MQROUTE\_DELIVER\_NO**  
On arrival, the trace-route message is discarded.

## Trace-route reply message reference

Use this page to obtain an overview of the trace-route reply message format. The trace-route reply message data is a duplicate of the trace-route message data from the trace-route message for which it was generated

**Related concepts:**

“Trace-route reply message format”

“Trace-route reply message data” on page 966

**Related reference:**

“Trace-route reply message MQMD (message descriptor)” on page 965

“Trace-route reply message MQCFH (PCF header)” on page 966

**Trace-route reply message format**

Trace-route reply messages are standard WebSphere MQ messages containing a message descriptor and message data. The message data contains information about the activities performed on a trace-route message as it has been routed through a queue manager network.

Trace-route reply messages contain the following information:

**A message descriptor**

An MQMD structure

**Message data**

A PCF header (MQCFH) and trace-route reply message data

Trace-route reply message data consists of one or more *Activity* PCF groups.

When a trace-route message reaches its target queue, a trace-route reply message can be generated that contains a copy of the activity information from the trace-route message. The trace-route reply message will be delivered to a reply-to queue or to a system queue.

Table 108 on page 965 shows the structure of a trace-route reply message, including parameters that are only returned under certain conditions.

Table 108. Trace-route reply message format

| Message descriptor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Message data                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQMD structure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | PCF header MQCFH structure                                                                                                                                                                             | Trace-route reply message data                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Structure identifier<br>Structure version<br>Report options<br>Message type<br>Expiration time<br>Feedback<br>Encoding<br>Coded character set ID<br>Message format<br>Priority<br>Persistence<br>Message identifier<br>Correlation identifier<br>Backout count<br>Reply-to queue<br>Reply-to queue manager<br>User identifier<br>Accounting token<br>Application identity data<br>Application type<br>Application name<br>Put date<br>Put time<br>Application origin data<br>Group identifier<br>Message sequence number<br>Offset<br>Message flags<br>Original length | PCF header (MQCFH)<br>Structure type<br>Structure length<br>Structure version<br>Command identifier<br>Message sequence number<br>Control options<br>Completion code<br>Reason code<br>Parameter count | Activity<br>Activity application name<br>Activity application type<br>Activity description<br>Operation<br>Operation type<br>Operation date<br>Operation time<br>Message<br>Message length<br>MQMD<br>EmbeddedMQMD<br>Queue manager name<br>Queue sharing group name<br>Queue name <sup>1 2 3</sup><br>Resolved queue name <sup>1 3</sup><br>Remote queue name <sup>3</sup><br>Remote queue manager-name <sup>2 3 4 5</sup><br>Feedback <sup>2</sup><br>Channel name <sup>4 5</sup><br>Channel type <sup>4 5</sup><br>Transmission queue name <sup>5</sup><br>TraceRoute<br>Detail<br>Recorded activities<br>Unrecorded activities<br>Discontinuity count<br>Max activities<br>Accumulate<br>Deliver |
| <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. Returned for Get and Browse operations.</li> <li>2. Returned for Discard operations.</li> <li>3. Returned for Put, Put Reply, and Put Report operations.</li> <li>4. Returned for Receive operations.</li> <li>5. Returned for Send operations.</li> </ol>                                                                                                                                                                                                                                               |                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

### Trace-route reply message MQMD (message descriptor)

Use this page to view the values contained by the MQMD structure for a trace-route reply message

For a trace-route reply message, the MQMD structure contains the parameters described in Activity report message descriptor. Some of the parameter values in a trace-route reply message descriptor are different from those in an activity report message descriptor, as follows:

*MsgType*

Description: Type of message.  
Data type: MQLONG.  
Value: **MQMT\_REPLY**

#### *Feedback*

Description: Feedback or reason code.  
Data type: MQLONG.  
Value: **MQFB\_NONE**

#### *Encoding*

Description: Numeric encoding of message data.  
Data type: MQLONG.  
Value: Copied from trace-route message descriptor.

#### *CodedCharSetId*

Description: Character set identifier of message data.  
Data type: MQLONG.  
Value: Copied from trace-route message descriptor.

#### *Format*

Description: Format name of message data  
Data type: MQCHAR8.  
Value: **MQFMT\_ADMIN**  
Admin message.

### **Trace-route reply message MQCFH (PCF header)**

Use this page to view the PCF values contained by the MQCFH structure for a trace-route reply message

The PCF header (MQCFH) for a trace-route reply message is the same as for a trace-route message.

### **Trace-route reply message data**

The trace-route reply message data is a duplicate of the trace-route message data from the trace-route message for which it was generated

The trace-route reply message data contains one or more *Activity* groups. The parameters are described in "Activity report message data" on page 935.

---

## **Accounting and statistics messages**

Queue managers generate accounting and statistics messages to record information about the MQI operations performed by IBM WebSphere MQ applications, or to record information about the activities occurring in a IBM WebSphere MQ system.

### **Accounting messages**

Accounting messages are used to record information about the MQI operations performed by IBM WebSphere MQ applications, see "Accounting messages" on page 967.

## Statistics messages

Statistics messages are used to record information about the activities occurring in a IBM WebSphere MQ system, see “Statistics messages” on page 971.

Accounting messages and statistics messages as described here are not available on IBM WebSphere MQ for z/OS, but equivalent functionality is available through the System Management Facility (SMF).

Accounting and statistics messages are delivered to one of two system queues. User applications can retrieve the messages from these system queues and use the recorded information for various purposes:

- Account for application resource use.
- Record application activity.
- Capacity planning.
- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Improve the efficiency of your queue manager network.
- Familiarize yourself with the running of your queue manager network.
- Confirm that your queue manager network is running correctly.

### Related concepts:

“Accounting messages”

“Statistics messages” on page 971

“Displaying accounting and statistics information” on page 976

“Accounting and statistics message reference” on page 981

“Using System Management Facility” on page 1119

“Accounting messages”

“Statistics messages” on page 971

“Displaying accounting and statistics information” on page 976

“Accounting and statistics message reference” on page 981

“Using System Management Facility” on page 1119

## Accounting messages

Accounting messages record information about the MQI operations performed by WebSphere MQ applications. An accounting message is a PCF message that contains a number of PCF structures.

When an application disconnects from a queue manager, an accounting message is generated and delivered to the system accounting queue (SYSTEM.ADMIN.ACCOUNTING.QUEUE). For long running WebSphere MQ applications, intermediate accounting messages are generated as follows:

- When the time since the connection was established exceeds the configured interval.
- When the time since the last intermediate accounting message exceeds the configured interval.

Accounting messages are in the following categories:

### MQI accounting messages

MQI accounting messages contain information relating to the number of MQI calls made using a connection to a queue manager.

### Queue accounting messages

Queue accounting messages contain information relating to the number of MQI calls made using connections to a queue manager, grouped by queue.

Each queue accounting message can contain up to 100 records, with every record relating to an activity performed by the application with respect to a specific queue.

Accounting messages are recorded only for local queues. If an application makes an MQI call against an alias queue, the accounting data is recorded against the base queue, and, for a remote queue, the accounting data is recorded against the transmission queue.

**Related concepts:**

“Accounting message format”

“Accounting information collection”

**Related reference:**

“MQI accounting message data” on page 985

“Queue accounting message data” on page 996

## Accounting message format

Accounting messages comprise a set of PCF fields that consist of a message descriptor and message data.

### Message descriptor

- An accounting message MQMD (message descriptor)

### Accounting message data

- An accounting message MQCFH (PCF header)
- Accounting message data that is always returned
- Accounting message data that is returned if available

The accounting message MQCFH (PCF header) contains information about the application, and the interval for which the accounting data was recorded.

Accounting message data comprises PCF parameters that store the accounting information. The content of accounting messages depends on the message category as follows:

### MQI accounting message

MQI accounting message data consists of a number of PCF parameters, but no PCF groups.

### Queue accounting message

Queue accounting message data consists of a number of PCF parameters, and in the range 1 through 100 *QAccountingData* PCF groups.

There is one *QAccountingData* PCF group for every queue that had accounting data collected. If an application accesses more than 100 queues, multiple accounting messages are generated. Each message has the *SeqNumber* in the MQCFH (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the MQCFH specified as MQCFC\_LAST.

## Accounting information collection

Use queue and queue manager attributes to control the collection of accounting information. You can also use MQCONN options to control collection at the connection level.

**Related concepts:**

“MQI accounting information”

“Queue accounting information” on page 969

“MQCONN options” on page 969

“Accounting message generation” on page 970

### MQI accounting information:

Use the queue manager attribute ACCTMQI to control the collection of MQI accounting information

To change the value of this attribute, use the MQSC command, ALTER QMGR, and specify the parameter ACCTMQI. Accounting messages are generated only for connections that begin after accounting is enabled. The ACCTMQI parameter can have the following values:

**ON** MQI accounting information is collected for every connection to the queue manager.

**OFF** MQI accounting information is not collected. This is the default value.

For example, to enable MQI accounting information collection use the following MQSC command:

```
ALTER QMGR ACCTMQI(ON)
```

### **Queue accounting information:**

Use the queue attribute ACCTQ and the queue manager attribute ACCTQ to control the collection of queue accounting information.

To change the value of the queue attribute, use the MQSC command, ALTER QLOCAL and specify the parameter ACCTQ. Accounting messages are generated only for connections that begin after accounting is enabled. The queue attribute ACCTQ can have the following values:

**ON** Queue accounting information for this queue is collected for every connection to the queue manager that opens the queue.

**OFF** Queue accounting information for this queue is not collected.

### **QMGR**

The collection of queue accounting information for this queue is controlled according to the value of the queue manager attribute ACCTQ. This is the default value.

To change the value of the queue manager attribute, use the MQSC command, ALTER QMGR and specify the parameter ACCTQ. The queue manager attribute ACCTQ can have the following values:

**ON** Queue accounting information is collected for queues that have the queue attribute ACCTQ set as QMGR.

**OFF** Queue accounting information is not collected for queues that have the queue attribute ACCTQ set as QMGR. This is the default value.

### **NONE**

The collection of queue accounting information is disabled for all queues, regardless of the queue attribute ACCTQ.

If the queue manager attribute, ACCTQ, is set to NONE, the collection of queue accounting information is disabled for all queues, regardless of the queue attribute ACCTQ.

For example, to enable accounting information collection for the queue, Q1, use the following MQSC command:

```
ALTER QLOCAL(Q1) ACCTQ(ON)
```

To enable accounting information collection for all queues that specify the queue attribute ACCTQ as QMGR, use the following MQSC command:

```
ALTER QMGR ACCTQ(ON)
```

### **MQCONN options:**

Use the **ConnectOpts** parameter on the MQCONN call to modify the collection of both MQI and queue accounting information at the connection level by overriding the effective values of the queue manager attributes ACCTMQI and ACCTQ

The **ConnectOpts** parameter can have the following values:

#### **MQCNO\_ACCOUNTING\_MQI\_ENABLED**

If the value of the queue manager attribute ACCTMQI is specified as OFF, MQI accounting is enabled for this connection. This is equivalent of the queue manager attribute ACCTMQI being specified as ON.

If the value of the queue manager attribute ACCTMQI is not specified as OFF, this attribute has no effect.

#### **MQCNO\_ACCOUNTING\_MQI\_DISABLED**

If the value of the queue manager attribute ACCTMQI is specified as ON, MQI accounting is disabled for this connection. This is equivalent of the queue manager attribute ACCTMQI being specified as OFF.

If the value of the queue manager attribute ACCTMQI is not specified as ON, this attribute has no effect.

#### **MQCNO\_ACCOUNTING\_Q\_ENABLED**

If the value of the queue manager attribute ACCTQ is specified as OFF, queue accounting is enabled for this connection. All queues with ACCTQ specified as QMGR, are enabled for queue accounting. This is equivalent of the queue manager attribute ACCTQ being specified as ON.

If the value of the queue manager attribute ACCTQ is not specified as OFF, this attribute has no effect.

#### **MQCNO\_ACCOUNTING\_Q\_DISABLED**

If the value of the queue manager attribute ACCTQ is specified as ON, queue accounting is disabled for this connection. This is equivalent of the queue manager attribute ACCTQ being specified as OFF.

If the value of the queue manager attribute ACCTQ is not specified as ON, this attribute has no effect.

These overrides are by disabled by default. To enable them, set the queue manager attribute ACCTCONO to ENABLED. To enable accounting overrides for individual connections use the following MQSC command:

```
ALTER QMGR ACCTCONO(ENABLED)
```

#### **Accounting message generation:**

Accounting messages are generated when an application disconnects from the queue manager. Intermediate accounting messages are also written for long running WebSphere MQ applications.

Accounting messages are generated in either of the following ways when an application disconnects:

- The application issues an MQDISC call
- The queue manager recognises that the application has terminated

Intermediate accounting messages are written for long running WebSphere MQ applications when the interval since the connection was established or since the last intermediate accounting message that was written exceeds the configured interval. The queue manager attribute, ACCTINT, specifies the time, in seconds, after which intermediate accounting messages can be automatically written. Accounting messages are generated only when the application interacts with the queue manager, so applications that remain connected to the queue manager for long periods without executing MQI requests do not generate accounting messages until the execution of the first MQI request following the completion of the accounting interval.

The default accounting interval is 1800 seconds (30 minutes). For example, to change the accounting interval to 900 seconds (15 minutes) use the following MQSC command:

```
ALTER QMGR ACCTINT(900)
```

## Statistics messages

Statistics messages record information about the activities occurring in a WebSphere MQ system. An statistics messages is a PCF message that contains a number of PCF structures.

Statistics messages are delivered to the system queue (SYSTEM.ADMIN.STATISTICS.QUEUE) at configured intervals.

Statistics messages are in the following categories:

### MQI statistics messages

MQI statistics messages contain information relating to the number of MQI calls made during a configured interval. For example, the information can include the number of MQI calls issued by a queue manager.

### Queue statistics messages

Queue statistics messages contain information relating to the activity of a queue during a configured interval. The information includes the number of messages put on, and retrieved from, the queue, and the total number of bytes processed by a queue.

Each queue statistics message can contain up to 100 records, with each record relating to the activity per queue for which statistics were collected.

Statistics messages are recorded only for local queues. If an application makes an MQI call against an alias queue, the statistics data is recorded against the base queue, and, for a remote queue, the statistics data is recorded against the transmission queue.

### Channel statistics messages

Channel statistics messages contain information relating to the activity of a channel during a configured interval. For example the information might be the number of messages transferred by the channel, or the number of bytes transferred by the channel.

Each channel statistics message contains up to 100 records, with each record relating to the activity per channel for which statistics were collected.

### Related concepts:

“Statistics messages format”

“Statistics information collection” on page 972

### Related reference:

“MQI statistics information” on page 972

“Queue statistics information” on page 972

“Channel statistics information” on page 973

## Statistics messages format

Statistics messages comprise a set of PCF fields that consist of a message descriptor and message data.

### Message descriptor

- A statistics message MQMD (message descriptor)

### Accounting message data

- A statistics message MQCFH (PCF header)
- Statistics message data that is always returned
- Statistics message data that is returned if available

The statistics message MQCFH (PCF header) contains information about the interval for which the statistics data was recorded.

Statistics message data comprises PCF parameters that store the statistics information. The content of statistics messages depends on the message category as follows:

### **MQI statistics message**

MQI statistics message data consists of a number of PCF parameters, but no PCF groups.

### **Queue statistics message**

Queue statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *QStatisticsData* PCF groups.

There is one *QStatisticsData* PCF group for every queue was active in the interval. If more than 100 queues were active in the interval, multiple statistics messages are generated. Each message has the *SeqNumber* in the MQCFH (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the MQCFH specified as MQCFC\_LAST.

### **Channel statistics message**

Channel statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *ChlStatisticsData* PCF groups.

There is one *ChlStatisticsData* PCF group for every channel that was active in the interval. If more than 100 channels were active in the interval, multiple statistics messages are generated. Each message has the *SeqNumber* in the MQCFH (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the MQCFH specified as MQCFC\_LAST.

## **Statistics information collection**

Use queue, queue manager, and channel attributes to control the collection of statistics information

### **Related concepts:**

“Statistics message generation” on page 975

### **Related reference:**

“MQI statistics information”

“Queue statistics information”

“Channel statistics information” on page 973

### **MQI statistics information:**

Use the queue manager attribute STATMQI to control the collection of MQI statistics information

To change the value of this attribute, use the MQSC command, ALTER QMGR and specify the parameter STATMQI. Statistics messages are generated only for queues that are opened after statistics collection has been enabled. The STATMQI parameter can have the following values:

**ON** MQI statistics information is collected for every connection to the queue manager.

**OFF** MQI statistics information is not collected. This is the default value.

For example, to enable MQI statistics information collection use the following MQSC command:

```
ALTER QMGR STATMQI(ON)
```

### **Queue statistics information:**

Use the queue attribute STATQ and the queue manager attribute STATQ to control the collection of queue statistics information.

You can enable or disable queue statistics information collection for individual queues or for multiple queues. To control individual queues, set the queue attribute STATQ. You enable or disable queue statistics information collection at the queue manager level by using the queue manager attribute STATQ. For all queues that have the queue attribute STATQ specified with the value QMGR, queue statistics information collection is controlled at the queue manager level.

Queue statistics are incremented only for operations using IBM WebSphere MQ MQI Object Handles that were opened after statistics collection has been enabled.

Queue Statistics messages are generated only for queues for which statistics data has been collected in the previous time period.

The same queue can have several put operations and get operations through several Object Handles. Some Object Handles might have been opened before statistics collection was enabled, but others were opened afterwards. Therefore, it is possible for the queue statistics to record the activity of some put operations and get operations, and not all.

To ensure that the Queue Statistics are recording the activity of all applications, you must close and reopen new Object Handles on the queue, or queues, that you are monitoring. The best way to achieve this, is to end and restart all applications after enabling statistics collection.

To change the value of the queue attribute `STATQ`, use the MQSC command, `ALTER QLOCAL` and specify the parameter `STATQ`. The queue attribute `STATQ` can have the following values:

**ON** Queue statistics information is collected for every connection to the queue manager that opens the queue.

**OFF** Queue statistics information for this queue is not collected.

**QMGR**

The collection of queue statistics information for this queue is controlled according to the value of the queue manager attribute, `STATQ`. This is the default value.

To change the value of the queue manager attribute `STATQ`, use the MQSC command, `ALTER QMGR` and specify the parameter `STATQ`. The queue manager attribute `STATQ` can have the following values:

**ON** Queue statistics information is collected for queues that have the queue attribute `STATQ` set as `QMGR`

**OFF** Queue statistics information is not collected for queues that have the queue attribute `STATQ` set as `QMGR`. This is the default value.

**NONE**

The collection of queue statistics information is disabled for all queues, regardless of the queue attribute `STATQ`.

If the queue manager attribute `STATQ` is set to `NONE`, the collection of queue statistics information is disabled for all queues, regardless of the queue attribute `STATQ`.

For example, to enable statistics information collection for the queue, `Q1`, use the following MQSC command:

```
ALTER QLOCAL(Q1) STATQ(ON)
```

To enable statistics information collection for all queues that specify the queue attribute `STATQ` as `QMGR`, use the following MQSC command:

```
ALTER QMGR STATQ(ON)
```

**Channel statistics information:**

Use the channel attribute `STATCHL` to control the collection of channel statistics information. You can also set queue manager attributes to control information collection.

You can enable or disable channel statistics information collection for individual channels, or for multiple channels. To control individual channels, you must set the channel attribute `STATCHL` to enable or disable channel statistic information collection. To control many channels together, you enable or disable

channel statistics information collection at the queue manager level by using the queue manager attribute `STATCHL`. For all channels that have the channel attribute `STATCHL` specified with the value `QMGR`, channel statistics information collection is controlled at the queue manager level.

Automatically defined cluster-sender channels are not WebSphere MQ objects, so do not have attributes in the same way as channel objects. To control automatically defined cluster-sender channels, use the queue manager attribute `STATACLS`. This attribute determines whether automatically defined cluster-sender channels within a queue manager are enabled or disabled for channel statistics information collection.

You can set channel statistics information collection to one of the three monitoring levels: low, medium or high. You can set the monitoring level at either object level or at the queue manager level. The choice of which level to use is dependant on your system. Collecting statistics information data might require some instructions that are relatively expensive computationally, so to reduce the impact of channel statistics information collection, the medium and low monitoring options measure a sample of the data at regular intervals rather than collecting data all the time. Table 109 summarizes the levels available with channel statistics information collection:

*Table 109. Detail level of channel statistics information collection*

| Level  | Description                                               | Usage                                                                                                        |
|--------|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Low    | Measure a small sample of the data, at regular intervals. | For objects that process a high volume of messages.                                                          |
| Medium | Measure a sample of the data, at regular intervals.       | For most objects.                                                                                            |
| High   | Measure all data, at regular intervals.                   | For objects that process only a few messages per second, on which the most current information is important. |

To change the value of the channel attribute `STATCHL`, use the MQSC command, `ALTER CHANNEL` and specify the parameter `STATCHL`.

To change the value of the queue manager attribute `STATCHL`, use the MQSC command, `ALTER QMGR` and specify the parameter `STATCHL`.

To change the value of the queue manager attribute `STATACLS`, use the MQSC command, `ALTER QMGR` and specify the parameter `STATACLS`.

The channel attribute, `STATCHL`, can have the following values:

**LOW** Channel statistics information is collected with a low level of detail.

**MEDIUM**

Channel statistics information is collected with a medium level of detail.

**HIGH** Channel statistics information is collected with a high level of detail.

**OFF** Channel statistics information is not collected for this channel.

This is the default value.

**QMGR**

The channel attribute is set as `QMGR`. The collection of statistics information for this channel is controlled by the value of the queue manager attribute, `STATCHL`.

The queue manager attribute, `STATCHL`, can have the following values:

**LOW** Channel statistics information is collected with a low level of detail, for all channels that have the channel attribute `STATCHL` set as `QMGR`.

**MEDIUM**

Channel statistics information is collected with a medium level of detail, for all channels that have the channel attribute `STATCHL` set as `QMGR`.

**HIGH** Channel statistics information is collected with a high level of detail, for all channels that have the channel attribute `STATCHL` set as `QMGR`.

**OFF** Channel statistics information is not collected for all channels that have the channel attribute `STATCHL` set as `QMGR`.

This is the default value.

**NONE**

The collection of channel statistics information is disabled for all channel, regardless of the channel attribute `STATCHL`.

The queue manager attribute, `STATACLS`, can have the following values:

**LOW** Statistics information is collected with a low level of detail for automatically defined cluster-sender channels.

**MEDIUM**

Statistics information is collected with a medium level of detail for automatically defined cluster-sender channels.

**HIGH** Statistics information is collected with a high level of detail for automatically defined cluster-sender channels.

**OFF** Statistics information is not for automatically defined cluster-sender channels.

This is the default value.

**QMGR**

The collection of statistics information for automatically defined cluster-sender channels is controlled by the value of the queue manager attribute, `STATCHL`.

For example, to enable statistics information collection, with a medium level of detail, for the sender channel `QM1.TO.QM2`, use the following MQSC command:

```
ALTER CHANNEL(QM1.TO.QM2) CHLTYPE(SDR) STATCHL(MEDIUM)
```

To enable statistics information collection, at a medium level of detail, for all channels that specify the channel attribute `STATCHL` as `QMGR`, use the following MQSC command:

```
ALTER QMGR STATCHL(MEDIUM)
```

To enable statistics information collection, at a medium level of detail, for all automatically defined cluster-sender channels, use the following MQSC command:

```
ALTER QMGR STATACLS(MEDIUM)
```

**Statistics message generation:**

Statistics messages are generated at configured intervals, and when a queue manager shuts down in a controlled fashion.

The configured interval is controlled by the `STATINT` queue manager attribute, which specifies the interval, in seconds, between the generation of statistics messages. The default statistics interval is 1800 seconds (30 minutes). To change the statistics interval, use the MQSC command `ALTER QMGR` and specify the `STATINT` parameter. For example, to change the statistics interval to 900 seconds (15 minutes) use the following MQSC command:

ALTER QMGR STATINT(900)

To write the currently collected statistics data to the statistics queue before the statistics collection interval is due to expire, use the MQSC command RESET QMGR TYPE(STATISTICS). Issuing this command causes the collected statistics data to be written to the statistics queue and a new statistics data collection interval to begin.

## Displaying accounting and statistics information

To use the information recorded in accounting and statistics messages, run an application such as the **amqsmon** sample program to transform the recorded information into a suitable format

Accounting and statistics messages are written to the system accounting and statistics queues. **amqsmon** is a sample program supplied with WebSphere MQ that processes messages from the accounting and statistics queues and displays the information to the screen in a readable form.

Because **amqsmon** is a sample program, you can use the supplied source code as template for writing your own application to process accounting or statistics messages, or modify the **amqsmon** source code to meet your own particular requirements.

### Related reference:

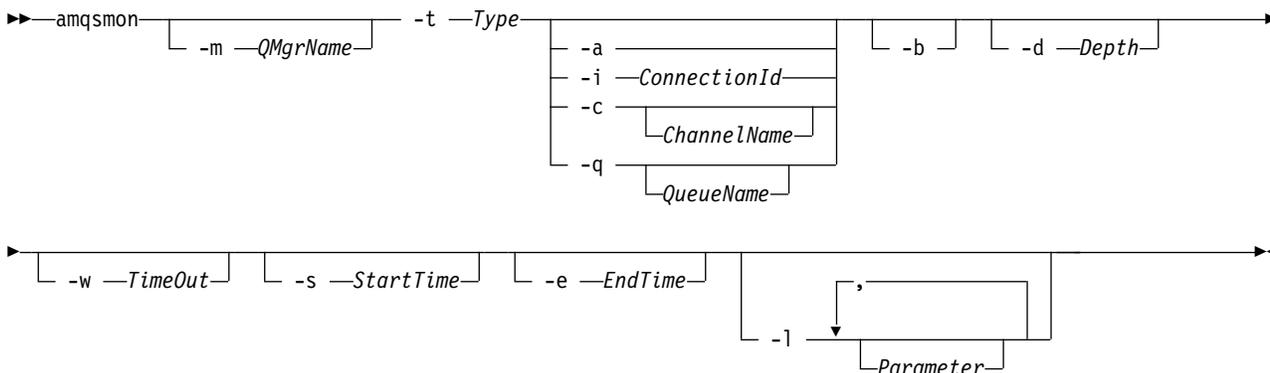
“amqsmon (Display formatted monitoring information)”

“amqsmon examples” on page 978

## amqsmon (Display formatted monitoring information)

Use the **amqsmon** sample program to display in a readable format the information contained within accounting and statistics messages. The **amqsmon** program reads accounting messages from the accounting queue, SYSTEM.ADMIN.ACCOUNTING.QUEUE. and reads statistics messages from the statistics queue, SYSTEM.ADMIN.STATISTICS.QUEUE.

### Syntax



### Required parameters

#### -t *Type*

The type of messages to process. Specify *Type* as one of the following:

#### accounting

Accounting records are processed. Messages are read from the system queue, SYSTEM.ADMIN.ACCOUNTING.QUEUE.

#### statistics

Statistics records are processed. Messages are read from the system queue, SYSTEM.ADMIN.STATISTICS.QUEUE.

## Optional Parameters

### **-m** *QMgrName*

The name of the queue manager from which accounting or statistics messages are to be processed.

If you do not specify this parameter, the default queue manager is used.

### **-a** Process messages containing MQI records only.

Only display MQI records. Messages not containing MQI records will always be left on the queue they were read from.

### **-q** *QueueName*

*QueueName* is an optional parameter.

If *QueueName* is not supplied:

Displays queue accounting and queue statistics records only.

If *QueueName* is supplied:

Displays queue accounting and queue statistics records for the queue specified by *QueueName* only.

If *-b* is not specified then the accounting and statistics messages from which the records came are discarded. Since accounting and statistics messages can also contain records from other queues, if *-b* is not specified then unseen records can be discarded.

### **-c** *ChannelName*

*ChannelName* is an optional parameter.

If *ChannelName* is not supplied:

Displays channel statistics records only.

If *ChannelName* is supplied:

Displays channel statistics records for the channel specified by *ChannelName* only.

If *-b* is not specified then the statistics messages from which the records came are discarded. Since statistics messages can also contain records from other channels, if *-b* is not specified then unseen records can be discarded.

This parameter is available when displaying statistics messages only, (*-t statistics*).

### **-i** *ConnectionId*

Displays records related to the connection identifier specified by *ConnectionId* only.

This parameter is available when displaying accounting messages only, (*-t accounting*).

If *-b* is not specified then the statistics messages from which the records came are discarded. Since statistics messages can also contain records from other channels, if *-b* is not specified then unseen records can be discarded.

### **-b** Browse messages.

Messages are retrieved non-destructively.

### **-d** *Depth*

The maximum number of messages that can be processed.

If you do not specify this parameter, then an unlimited number of messages can be processed.

### **-w** *TimeOut*

Time maximum number of seconds to wait for a message to become available.

If you do not specify this parameter, **amqsmon** will end once there are no more messages to process.

### **-s** *StartTime*

Process messages put after the specified *StartTime* only.

*StartTime* is specified in the format yyyy-mm-dd hh.mm.ss. If a date is specified without a time, then the time will default to 00.00.00 on the date specified. Times are in GMT.

For the effect of not specifying this parameter, see Note 1.

**-e EndTime**

Process messages put before the specified *EndTime* only.

The *EndTime* is specified in the format yyyy-mm-dd hh.mm.ss. If a date is specified without a time, then the time will default to 23.59.59 on the date specified. Times are in GMT.

For the effect of not specifying this parameter, see Note 1.

**-l Parameter**

Only display the selected fields from the records processed. *Parameter* is a comma-separated list of integer values, with each integer value mapping to the numeric constant of a field, see amqsmon example 5.

If you do not specify this parameter, then all available fields are displayed.

**Note:**

1. If you do not specify *-s StartTime* or *-e EndTime*, the messages that can be processed are not restricted by put time.

**amqsmon examples**

Use this page to view examples of running the amqsmon (Display formatted monitoring information) sample program

1.

The following command displays all MQI statistics messages from queue manager saturn.queue.manager:

```
amqsmon -m saturn.queue.manager -t statistics -a
```

The output from this command follows:

```
RecordType: MQIStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ConnCount: 23
ConnFailCount: 0
ConnHighwater: 8
DiscCount: [17, 0, 0]
OpenCount: [0, 80, 1, 0, 0, 3, 0, 0, 0, 0, 0, 0, 0]
OpenFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseCount: [0, 73, 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0]
CloseFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
InqCount: [4, 2102, 0, 0, 0, 46, 0, 0, 0, 0, 0, 0, 0]
InqFailCount: [0, 31, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetCount: [0, 1, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0]
SetFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
PutCount: [26, 1]
PutFailCount: 0
Put1Count: [40, 0]
Put1FailCount: 0
PutBytes: [57064, 12320]
GetCount: [18, 1]
GetBytes: [52, 12320]
GetFailCount: 2254
```

```
BrowseCount: [18, 60]
BrowseBytes: [23784, 30760]
BrowseFailCount: 9
CommitCount: 0
CommitFailCount: 0
BackCount: 0
ExpiredMsgCount: 0
PurgeCount: 0
```

2.

The following command displays all queue statistics messages for queue LOCALQ on queue manager saturn.queue.manager:

```
amqsmon -m saturn.queue.manager -t statistics -q LOCALQ
```

The output from this command follows:

```
RecordType: QueueStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ObjectCount: 3
QueueStatistics:
 QueueName: 'LOCALQ'
 CreateDate: '2005-03-08'
 CreateTime: '17.07.02'
 QueueType: Predefined
 QueueDefinitionType: Local
 QMinDepth: 0
 QMaxDepth: 18
 AverageQueueTime: [29827281, 0]
 PutCount: [26, 0]
 PutFailCount: 0
 Put1Count: [0, 0]
 Put1FailCount: 0
 PutBytes: [88, 0]
 GetCount: [18, 0]
 GetBytes: [52, 0]
 GetFailCount: 0
 BrowseCount: [0, 0]
 BrowseBytes: [0, 0]
 BrowseFailCount: 1
 NonQueuedMsgCount: 0
 ExpiredMsgCount: 0
 PurgedMsgCount: 0
```

3. The following command displays all of the statistics messages recorded since 15:30 on 30 April 2005 from queue manager saturn.queue.manager.

```
amqsmon -m saturn.queue.manager -t statistics -s "2005-04-30 15.30.00"
```

The output from this command follows:

```
RecordType: MQIStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
```

```
ConnCount: 23
ConnFailCount: 0
ConnHighwater: 8
DiscCount: [17, 0, 0]
OpenCount: [0, 80, 1, 0, 0, 3, 0, 0, 0, 0, 0, 0, 0]
```

```
...
RecordType: QueueStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ObjectCount: 3
QueueStatistics: 0
 QueueName: 'LOCALQ'
 CreateDate: '2005-03-08'
 CreateTime: '17.07.02'
 QueueType: Predefined
```

```
...
QueueStatistics: 1
 QueueName: 'SAMPLEQ'
 CreateDate: '2005-03-08'
 CreateTime: '17.07.02'
 QueueType: Predefined
...
```

4. The following command displays all accounting messages recorded on 30 April 2005 from queue manager saturn.queue.manager:

```
amqsmon -m saturn.queue.manager -t accounting -s "2005-04-30" -e "2005-04-30"
```

The output from this command follows:

```
RecordType: MQIAccounting
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.29'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.09.30'
CommandLevel: 600
ConnectionId: x'414d514354524556312020202020208d0b3742010a0020'
SeqNumber: 0
ApplicationName: 'amqsput'
ApplicationPid: 8572
ApplicationTid: 1
UserId: 'admin'
ConnDate: '2005-03-16'
ConnTime: '15.09.29'
DiscDate: '2005-03-16'
DiscTime: '15.09.30'
DiscType: Normal
OpenCount: [0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
OpenFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseCount: [0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
PutCount: [1, 0]
PutFailCount: 0
PutBytes: [4, 0]
GetCount: [0, 0]
GetFailCount: 0
GetBytes: [0, 0]
```

```

BrowseCount: [0, 0]
BrowseFailCount: 0
BrowseBytes: [0, 0]
CommitCount: 0
CommitFailCount: 0
BackCount: 0
InqCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
InqFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]

RecordType: MQIAccounting
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-03-16'
IntervalStartTime: '15.16.22'
IntervalEndDate: '2005-03-16'
IntervalEndTime: '15.16.22'
CommandLevel: 600
ConnectionId: x'414d514354524556312020202020208d0b3742010c0020'
SeqNumber: 0
ApplicationName: 'runmqsc'
ApplicationPid: 8615
ApplicationTid: 1
...

```

5. The following command browses the accounting queue and displays the application name and connection identifier of every application for which MQI accounting information is available:  
`amqsmn -m saturn.queue.manager -t accounting -b -a -l 7006,3024`

The output from this command follows:

```

ConnectionId: x'414d514354524556312020202020208d0b374203090020'
ApplicationName: 'runmqsc'

ConnectionId: x'414d514354524556312020202020208d0b3742010a0020'
ApplicationName: 'amqsput'

ConnectionId: x'414d514354524556312020202020208d0b3742010c0020'
ApplicationName: 'runmqsc'

ConnectionId: x'414d514354524556312020202020208d0b3742010d0020'
ApplicationName: 'amqsput'

ConnectionId: x'414d514354524556312020202020208d0b3742150d0020'
ApplicationName: 'amqsget'

```

5 Records Processed.

## Accounting and statistics message reference

Use this page to obtain an overview of the format of accounting and statistics messages and the information returned in these messages

Accounting and statistics message messages are standard WebSphere MQ messages containing a message descriptor and message data. The message data contains information about the MQI operations performed by WebSphere MQ applications, or information about the activities occurring in a WebSphere MQ system.

### Message descriptor

- An MQMD structure

**Message data**

- A PCF header (MQCFH)
- Accounting or statistics message data that is always returned
- Accounting or statistics message data that is returned if available

**Related concepts:**

“Accounting and statistics message format”

“Message data in accounting and statistics messages” on page 984

**Related reference:**

“Accounting and statistics message MQMD (message descriptor)” on page 983

“MQI accounting message data” on page 985

“Queue accounting message data” on page 996

“MQI statistics message data” on page 1006

“Queue statistics message data” on page 1017

“Channel statistics message data” on page 1024

“Reference notes” on page 1029

**Accounting and statistics message format**

Use this page as an example of the structure of an MQI accounting message

Table 110. MQI accounting message structure

| Message descriptor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Message data                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQMD structure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Accounting message header<br>MQCFH structure                                                                                                                                     | MQI accounting message data <sup>1</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Structure identifier<br>Structure version<br>Report options<br>Message type<br>Expiration time<br>Feedback code<br>Encoding<br>Coded character set ID<br>Message format<br>Message priority<br>Persistence<br>Message identifier<br>Correlation identifier<br>Backout count<br>Reply-to queue<br>Reply-to queue manager<br>User identifier<br>Accounting token<br>Application identity data<br>Application type<br>Application name<br>Put date<br>Put time<br>Application origin data<br>Group identifier<br>Message sequence number<br>Offset<br>Message flags<br>Original length | Structure type<br>Structure length<br>Structure version<br>Command identifier<br>Message sequence number<br>Control options<br>Completion code<br>Reason code<br>Parameter count | Queue manager<br>Interval start date<br>Interval start time<br>Interval end date<br>Interval end time<br>Command level<br>Connection identifier<br>Sequence number<br>Application name<br>Application process identifier<br>Application thread identifier<br>User identifier<br>Connection date<br>Connection time<br>Connection name<br>Channel name<br>Disconnect date<br>Disconnect time<br>Disconnect type<br>Open count<br>Open fail count<br>Close count<br>Close fail count<br>Put count<br>Put fail count<br>Put1 count<br>Put1 fail count<br>Put bytes<br>Get count<br>Get fail count<br>Get bytes<br>Browse count<br>Browse fail count<br>Browse bytes<br>Commit count<br>Commit fail count<br>Backout count<br>Inquire count<br>Inquire fail count<br>Set count<br>Set fail count |
| <p><b>Note:</b></p> <p>1. The parameters shown are those returned for an MQI accounting message. The actual accounting or statistics message data depends on the message category.</p>                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

### Accounting and statistics message MQMD (message descriptor)

Use this page to understand the differences between the message descriptor of accounting and statistics messages and the message descriptor of event messages

The parameters and values in the message descriptor of accounting and statistics message are the same as in the message descriptor of event messages, with the following exception:

#### *Format*

Description: Format name of message data.  
Data type: MQCHAR8.  
Value: **MQFMT\_ADMIN**  
Admin message.

Some of the parameters contained in the message descriptor of accounting and statistics message contain fixed data supplied by the queue manager that generated the message.

The MQMD also specifies the name of the queue manager (truncated to 28 characters) that put the message, and the date and time when the message was put on the accounting, or statistics, queue.

### **Message data in accounting and statistics messages**

The message data in accounting and statistics messages is based on the programmable command format (PCF), which is used in PCF command inquiries and responses. The message data in accounting and statistics messages consists of a PCF header (MQCFH) and an accounting or statistics report.

### **Accounting and statistics message MQCFH (PCF header)**

The message header of accounting and statistics messages is an MQCFH structure. The parameters and values in the message header of accounting and statistics message are the same as in the message header of event messages, with the following exceptions:

#### *Command*

Description: Command identifier. This identifies the accounting or statistics message category.  
Data type: MQLONG.  
Values: **MQCMD\_ACCOUNTING\_MQI**  
MQI accounting message.  
**MQCMD\_ACCOUNTING\_Q**  
Queue accounting message.  
**MQCMD\_STATISTICS\_MQI**  
MQI statistics message.  
**MQCMD\_STATISTICS\_Q**  
Queue statistics message.  
**MQCMD\_STATISTICS\_CHANNEL**  
Channel statistics message.

#### *Version*

Description: Structure version number.  
 Data type: MQLONG.  
 Value: **MQCFH\_VERSION\_3**  
 Version-3 for accounting and statistics messages.

## Accounting and statistics message data

The content of accounting and statistics message data is dependent on the category of the accounting or statistics message, as follows:

### MQI accounting message

MQI accounting message data consists of a number of PCF parameters, but no PCF groups.

### Queue accounting message

Queue accounting message data consists of a number of PCF parameters, and in the range 1 through 100 *QAccountingData* PCF groups.

### MQI statistics message

MQI statistics message data consists of a number of PCF parameters, but no PCF groups.

### Queue statistics message

Queue statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *QStatisticsData* PCF groups.

### Channel statistics message

Channel statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *ChlStatisticsData* PCF groups.

## MQI accounting message data

Use this page to view the structure of an MQI accounting message

|               |                                    |
|---------------|------------------------------------|
| Message name: | MQI accounting message.            |
| Platforms:    | All, except WebSphere MQ for z/OS. |
| System queue: | SYSTEM.ADMIN.ACCOUNTING.QUEUE.     |

### *QueueManager*

Description: The name of the queue manager  
 Identifier: MQCA\_Q\_MGR\_NAME  
 Data type: MQCFST  
 Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
 Returned: Always

### *IntervalStartDate*

Description: The date of the start of the monitoring period  
 Identifier: MQCAMO\_START\_DATE  
 Data type: MQCFST  
 Maximum length: MQ\_DATE\_LENGTH  
 Returned: Always

### *IntervalStartTime*

Description: The time of the start of the monitoring period  
Identifier: MQCAMO\_START\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

#### *IntervalEndDate*

Description: The date of the end of the monitoring period  
Identifier: MQCAMO\_END\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH  
Returned: Always

#### *IntervalEndTime*

Description: The time of the end of the monitoring period  
Identifier: MQCAMO\_END\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

#### *CommandLevel*

Description: The queue manager command level  
Identifier: MQIA\_COMMAND\_LEVEL  
Data type: MQCFIN  
Returned: Always

#### *ConnectionId*

Description: The connection identifier for the WebSphere MQ connection  
Identifier: MQBACF\_CONNECTION\_ID  
Data type: MQCFBS  
Maximum length: MQ\_CONNECTION\_ID\_LENGTH  
Returned: Always

#### *SeqNumber*

Description: The sequence number. This value is incremented for each subsequent record for long running connections.  
Identifier: MQIACF\_SEQUENCE\_NUMBER  
Data type: MQCFIN  
Returned: Always

#### *ApplicationName*

Description: The name of the application. The contents of this field are equivalent to the contents of the *PutApplName* field in the message descriptor.  
Identifier: MQCACF\_APPL\_NAME  
Data type: MQCFST  
Maximum length: MQ\_APPL\_NAME\_LENGTH  
Returned: Always

#### *ApplicationPid*

Description: The operating system process identifier of the application  
Identifier: MQIACF\_PROCESS\_ID  
Data type: MQCFIN  
Returned: Always

#### *ApplicationTid*

Description: The WebSphere MQ thread identifier of the connection in the application  
Identifier: MQIACF\_THREAD\_ID  
Data type: MQCFIN  
Returned: Always

#### *UserId*

Description: The user identifier context of the application  
Identifier: MQCACF\_USER\_IDENTIFIER  
Data type: MQCFST  
Maximum length: MQ\_USER\_ID\_LENGTH  
Returned: Always

#### *ConnDate*

Description: Date of MQCONN operation  
Identifier: MQCAMO\_CONN\_DATE  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: When available

#### *ConnTime*

Description: Time of MQCONN operation  
Identifier: MQCAMO\_CONN\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: When available

#### *ConnName*

Description: Connection name for client connection  
Identifier: MQCACH\_CONNECTION\_NAME  
Data type: MQCFST  
Maximum length: MQ\_CONN\_NAME\_LENGTH  
Returned: When available

#### *ChannelName*

Description: Channel name for client connection  
Identifier: MQCACH\_CHANNEL\_NAME  
Data type: MQCFST  
Maximum length: MQ\_CHANNEL\_NAME\_LENGTH  
Returned: When available

#### *DiscDate*

Description: Date of MQDISC operation  
Identifier: MQCAMO\_DISC\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH  
Returned: When available

#### *DiscTime*

Description: Time of MQDISC operation  
Identifier: MQCAMO\_DISC\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: When available

#### *DiscType*

Description: Type of disconnect  
Identifier: MQIAMO\_DISC\_TYPE  
Data type: MQCFIN  
Values: The possible values are:  
**MQDISCONNECT\_NORMAL**  
Requested by application  
**MQDISCONNECT\_IMPLICIT**  
Abnormal application termination  
**MQDISCONNECT\_Q\_MGR**  
Connection broken by queue manager  
Returned: When available

#### *OpenCount*

Description: The number of objects opened. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_OPENS  
Data type: MQCFIL  
Returned: When available

#### *OpenFailCount*

Description: The number of unsuccessful attempts to open an object. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_OPENS\_FAILED  
Data type: MQCFIL  
Returned: When available

#### *CloseCount*

Description: The number of objects closed. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_CLOSES  
Data type: MQCFIL  
Returned: When available

#### *CloseFailCount*

Description: The number of unsuccessful attempts to close an object. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_CLOSES\_FAILED  
Data type: MQCFIL  
Returned: When available

#### *PutCount*

Description: The number persistent and nonpersistent messages successfully put to a queue, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_PUTS  
Data type: MQCFIL  
Returned: When available

#### *PutFailCount*

Description: The number of unsuccessful attempts to put a message  
Identifier: MQIAMO\_PUTS\_FAILED  
Data type: MQCFIN  
Returned: When available

#### *Put1Count*

Description: The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_PUT1S  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *Put1FailCount*

Description: The number of unsuccessful attempts to put a message using MQPUT1 calls  
Identifier: MQIAMO\_PUT1S\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *PutBytes*

Description: The number bytes written using put calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_PUT\_BYTES  
Data type: MQCFIL64  
Returned: When available

#### *GetCount*

Description: The number of successful destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_GETS  
Data type: MQCFIL  
Returned: When available

#### *GetFailCount*

Description: The number of failed destructive MQGET calls  
Identifier: MQIAMO\_GETS\_FAILED  
Data type: MQCFIN  
Returned: When available

#### *GetBytes*

Description: Total number of bytes retrieved for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_GET\_BYTES  
Data type: MQCFIL64  
Returned: When available

#### *BrowseCount*

Description: The number of successful non-destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_BROWSES  
Data type: MQCFIL  
Returned: When available

#### *BrowseFailCount*

Description: The number of unsuccessful non-destructive MQGET calls  
Identifier: MQIAMO\_BROWSES\_FAILED  
Data type: MQCFIN  
Returned: When available

#### *BrowseBytes*

Description: Total number of bytes browsed for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_BROWSE\_BYTES  
Data type: MQCFIL64  
Returned: When available

#### *CommitCount*

Description: The number of successful transactions. This number includes those transactions committed implicitly by the connected application. Commit requests where there is no outstanding work are included in this count.  
Identifier: MQIAMO\_COMMITS  
Data type: MQCFIN  
Returned: When available

#### *CommitFailCount*

Description: The number of unsuccessful attempts to complete a transaction  
Identifier: MQIAMO\_COMMITS\_FAILED  
Data type: MQCFIN  
Returned: When available

#### *BackCount*

Description: The number of backouts processed, including implicit backouts due to abnormal disconnection  
Identifier: MQIAMO\_BACKOUTS  
Data type: MQCFIN  
Returned: When available

#### *InqCount*

Description: The number of successful objects inquired upon. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_INQS  
Data type: MQCFIL  
Returned: When available

#### *InqFailCount*

Description: The number of unsuccessful object inquire attempts. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_INQS\_FAILED  
Data type: MQCFIL  
Returned: When available

#### *SetCount*

Description: The number of successful MQSET calls. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_SETS  
Data type: MQCFIL  
Returned: When available

#### *SetFailCount*

Description: The number of unsuccessful MQSET calls. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_SETS\_FAILED  
Data type: MQCFIL  
Returned: When available

#### *SubCountDur*

Description: The number of successful subscribe requests which created, altered or resumed durable subscriptions. This is an array of values indexed by the type of operation  
  
0 = The number of subscriptions created  
  
1 = The number of subscriptions altered  
  
2 = The number of subscriptions resumed  
Identifier: MQIAMO\_SUBS\_DUR  
Data type: MQCFIL  
Returned: When available.

#### *SubCountNDur*

Description: The number of successful subscribe requests which created, altered or resumed non-durable subscriptions. This is an array of values indexed by the type of operation

0 = The number of subscriptions created

1 = The number of subscriptions altered

2 = The number of subscriptions resumed

Identifier: MQIAMO\_SUBS\_NDUR

Data type: MQCFIL

Returned: When available.

#### *SubFailCount*

Description: The number of unsuccessful Subscribe requests.

Identifier: MQIAMO\_SUBS\_FAILED

Data type: MQCFIN

Returned: When available.

#### *UnsubCountDur*

Description: The number of successful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation

0 – The subscription was closed but not removed

1 – The subscription was closed and removed

Identifier: MQIAMO\_UNSUBS\_DUR

Data type: MQCFIL

Returned: When available.

#### *UnsubCountNDur*

Description: The number of successful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation

0 – The subscription was closed but not removed

1 – The subscription was closed and removed

Identifier: MQIAMO\_UNSUBS\_NDUR

Data type: MQCFIL

Returned: When available.

#### *UnsubFailCount*

Description: The number of unsuccessful unsubscribe requests.

Identifier: MQIAMO\_UNSUBS\_FAILED

Data type: MQCFIN

Returned: When available.

#### *SubRqCount*

Description: The number of successful MQSUBRQ requests.  
Identifier: MQIAMO\_SUBRQS  
Data type: MQCFIN  
Returned: When available.

#### *SubRqFailCount*

Description: The number of unsuccessful MQSUB requests.  
Identifier: MQIAMO\_SUBRQS\_FAILED  
Data type: MQCFIN  
Returned: When available.

#### *CBCount*

Description: The number of successful MQCB requests. This is an array of values indexed by the type of operation

- 0 – A callback was created or altered
- 1 – A callback was removed
- 2 – A callback was resumed
- 3 – A callback was suspended

Identifier: MQIAMO\_CBS  
Data type: MQCFIN  
Returned: When available.

#### *CBFailCount*

Description: The number of unsuccessful MQCB requests.  
Identifier: MQIAMO\_CBS\_FAILED  
Data type: MQCFIN  
Returned: When available.

#### *CtlCount*

Description: The number of successful MQCTL requests. This is an array of values indexed by the type of operation

- 0 – The connection was started
- 1 – The connection was stopped
- 2 – The connection was resumed
- 3 – The connection was suspended

Identifier: MQIAMO\_CTL  
Data type: MQCFIL  
Returned: When available.

#### *CtlFailCount*

Description: The number of unsuccessful MQCTL requests.  
Identifier: MQIAMO\_CTLIS\_FAILED  
Data type: MQCFIN  
Returned: When available.

#### *StatCount*

Description: The number of successful MQSTAT requests.  
Identifier: MQIAMO\_STATS.  
Data type: MQCFIN  
Returned: When available.

#### *StatFailCount*

Description: The number of unsuccessful MQSTAT requests.  
Identifier: MQIAMO\_STATS\_FAILED  
Data type: MQCFIN  
Returned: When available.

#### *PutTopicCount*

Description: The number persistent and nonpersistent messages successfully put to a topic, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see Reference note 2.

Note: Messages put using a queue alias which resolve to a topic are included in this value.

Identifier: MQIAMO\_TOPIC\_PUTS  
Data type: MQCFIL  
Returned: When available.

#### *PutTopicFailCount*

Description: The number of unsuccessful attempts to put a message to a topic.  
Identifier: MQIAMO\_TOPIC\_PUTS\_FAILED  
Data type: MQCFIN  
Returned: When available.

#### *Put1TopicCount*

Description: The number of persistent and nonpersistent messages successfully put to a topic using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2.

Note: Messages put using a queue alias which resolve to a topic are included in this value.

Identifier: MQIAMO\_TOPIC\_PUT1S  
Data type: MQCFIL  
Returned: When available.

#### *Put1TopicFailCount*

Description: The number of unsuccessful attempts to put a message to a topic using MQPUT1 calls.  
 Identifier: MQIAMO\_TOPIC\_PUT1S\_FAILED  
 Data type: MQCFIN  
 Returned: When available.

*PutTopicBytes*

Description: The number bytes written using put calls for persistent and nonpersistent messages which resolve to a publish operation. This is number of bytes put by the application and not the resultant number of bytes delivered to subscribers. This parameter is an integer list indexed by persistence value, see Reference note 2.  
 Identifier: MQIAMO64\_TOPIC\_PUT\_BYTES  
 Data type: MQCFIL64  
 Returned: When available.

**Queue accounting message data**

Use this page to view the structure of a queue accounting message

|               |                                    |
|---------------|------------------------------------|
| Message name: | Queue accounting message.          |
| Platforms:    | All, except WebSphere MQ for z/OS. |
| System queue: | SYSTEM.ADMIN.ACCOUNTING.QUEUE.     |

*QueueManager*

Description: The name of the queue manager  
 Identifier: MQCA\_Q\_MGR\_NAME  
 Data type: MQCFST  
 Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
 Returned: Always

*IntervalStartDate*

Description: The date of the start of the monitoring period  
 Identifier: MQCAMO\_START\_DATE  
 Data type: MQCFST  
 Maximum length: MQ\_DATE\_LENGTH  
 Returned: Always

*IntervalStartTime*

Description: The time of the start of the monitoring period  
 Identifier: MQCAMO\_START\_TIME  
 Data type: MQCFST  
 Maximum length: MQ\_TIME\_LENGTH  
 Returned: Always

*IntervalEndDate*

Description: The date of the end of the monitoring period  
Identifier: MQCAMO\_END\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH  
Returned: Always

#### *IntervalEndTime*

Description: The time of the end of the monitoring period  
Identifier: MQCAMO\_END\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

#### *CommandLevel*

Description: The queue manager command level  
Identifier: MQIA\_COMMAND\_LEVEL  
Data type: MQCFIN  
Returned: Always

#### *ConnectionId*

Description: The connection identifier for the WebSphere MQ connection  
Identifier: MQBACF\_CONNECTION\_ID  
Data type: MQCFBS  
Maximum length: MQ\_CONNECTION\_ID\_LENGTH  
Returned: Always

#### *SeqNumber*

Description: The sequence number. This value is incremented for each subsequent record for long running connections.  
Identifier: MQIACF\_SEQUENCE\_NUMBER  
Data type: MQCFIN  
Returned: Always

#### *ApplicationName*

Description: The name of the application. The contents of this field are equivalent to the contents of the PutApplName field in the message descriptor.  
Identifier: MQCACF\_APPL\_NAME  
Data type: MQCFST  
Maximum length: MQ\_APPL\_NAME\_LENGTH  
Returned: Always

#### *ApplicationPid*

Description: The operating system process identifier of the application  
Identifier: MQIACF\_PROCESS\_ID  
Data type: MQCFIN  
Returned: Always

#### *ApplicationTid*

Description: The WebSphere MQ thread identifier of the connection in the application  
Identifier: MQIACF\_THREAD\_ID  
Data type: MQCFIN  
Returned: Always

#### *UserId*

Description: The user identifier context of the application  
Identifier: MQCACF\_USER\_IDENTIFIER  
Data type: MQCFST  
Maximum length: MQ\_USER\_ID\_LENGTH  
Returned: Always

#### *ObjectCount*

Description: The number of queues accessed in the interval for which accounting data has been recorded. This value is set to the number of *QAccountingData* PCF groups contained in the message.  
Identifier: MQIAMO\_OBJECT\_COUNT  
Data type: MQCFIN  
Returned: Always

#### *QAccountingData*

Description: Grouped parameters specifying accounting details for a queue  
Identifier: MQGACF\_Q\_ACCOUNTING\_DATA  
Data type: MQCFGR

Parameters in group:

- QName*
- CreateDate*
- CreateTime*
- QType*
- QDefinitionType*
- OpenCount*
- OpenDate*
- OpenTime*
- CloseDate*
- CloseTime*
- PutCount*
- PutFailCount*
- Put1Count*
- Put1FailCount*
- PutBytes*
- PutMinBytes*
- PutMaxBytes*
- GetCount*
- GetFailCount*
- GetBytes*
- GetMinBytes*
- GetMaxBytes*
- BrowseCount*
- BrowseFailCount*
- BrowseBytes*
- BrowseMinBytes*
- BrowseMaxBytes*
- TimeOnQMin*
- TimeOnQAvg*
- TimeOnQMax*

Returned: Always

### *QName*

Description: The name of the queue  
 Identifier: MQCA\_Q\_NAME  
 Data type: MQCFST  
 Included in PCF group: *QAccountingData*  
 Maximum length: MQ\_Q\_NAME\_LENGTH  
 Returned: When available

### *CreateDate*

Description: The date the queue was created  
Identifier: MQCA\_CREATION\_DATE  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Maximum length: MQ\_DATE\_LENGTH  
Returned: When available

#### *CreateTime*

Description: The time the queue was created  
Identifier: MQCA\_CREATION\_TIME  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Maximum length: MQ\_TIME\_LENGTH  
Returned: When available

#### *QType*

Description: The type of the queue  
Identifier: MQIA\_Q\_TYPE  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Value: MQQT\_LOCAL  
Returned: When available

#### *QDefinitionType*

Description: The queue definition type  
Identifier: MQIA\_DEFINITION\_TYPE  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Values: Possible values are:  
**MQQDT\_PREDEFINED**  
**MQQDT\_PERMANENT\_DYNAMIC**  
**MQQDT\_TEMPORARY\_DYNAMIC**  
Returned: When available

#### *OpenCount*

Description: The number of times this queue was opened by the application in this interval  
Identifier: MQIAMO\_OPENS  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *OpenDate*

Description: The date the queue was first opened in this recording interval. If the queue was already open at the start of this interval, this value reflects the date the queue was originally opened.  
Identifier: MQCAMO\_OPEN\_DATE  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *OpenTime*

Description: The time the queue was first opened in this recording interval. If the queue was already open at the start of this interval, this value reflects the time the queue was originally opened.  
Identifier: MQCAMO\_OPEN\_TIME  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *CloseDate*

Description: The date of the final close of the queue in this recording interval. If the queue is still open then the value is not returned.  
Identifier: MQCAMO\_CLOSE\_DATE  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *CloseTime*

Description: The time of final close of the queue in this recording interval. If the queue is still open then the value is not returned.  
Identifier: MQCAMO\_CLOSE\_TIME  
Data type: MQCFST  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *PutCount*

Description: The number of persistent and nonpersistent messages successfully put to the queue, with the exception of MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_PUTS  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *PutFailCount*

Description: The number of unsuccessful attempts to put a message, with the exception of MQPUT1 calls  
Identifier: MQIAMO\_PUTS\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *Put1Count*

Description: The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_PUT1S  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *Put1FailCount*

Description: The number of unsuccessful attempts to put a message using MQPUT1 calls  
Identifier: MQIAMO\_PUT1S\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *PutBytes*

Description: The total number of bytes put for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_PUT\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *PutMinBytes*

Description: The smallest persistent and nonpersistent message size placed on the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
 Identifier: MQIAMO\_PUT\_MIN\_BYTES  
 Data type: MQCFIL  
 Included in PCF group: *QAccountingData*  
 Returned: When available

### *PutMaxBytes*

Description: The largest persistent and nonpersistent message size placed on the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
 Identifier: MQIAMO\_PUT\_MAX\_BYTES  
 Data type: MQCFIL  
 Included in PCF group: *QAccountingData*  
 Returned: When available

### *GeneratedMsgCount*

Description: The number of generated messages. Generated messages are
 

- Queue Depth Hi Events
- Queue Depth Low Events

 Identifier: MQIAMO\_GENERATED\_MSGS  
 Data type: MQCFIN  
 Included in PCF group: *QAccountingData*  
 Returned: When available

### *GetCount*

Description: The number of successful destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
 Identifier: MQIAMO\_GETS  
 Data type: MQCFIL  
 Included in PCF group: *QAccountingData*  
 Returned: When available

### *GetFailCount*

Description: The number of failed destructive MQGET calls  
 Identifier: MQIAMO\_GETS\_FAILED  
 Data type: MQCFIN  
 Included in PCF group: *QAccountingData*  
 Returned: When available

### *GetBytes*

Description: The number of bytes read in destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_GET\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *GetMinBytes*

Description: The size of the smallest persistent and nonpersistent message retrieved from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_GET\_MIN\_BYTES  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *GetMaxBytes*

Description: The size of the largest persistent and nonpersistent message retrieved from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_GET\_MAX\_BYTES  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *BrowseCount*

Description: The number of successful non-destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_BROWSES  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *BrowseFailCount*

Description: The number of unsuccessful non-destructive MQGET calls  
Identifier: MQIAMO\_BROWSES\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QAccountingData*  
Returned: When available

#### *BrowseBytes*

Description: The number of bytes read in non-destructive MQGET calls that returned persistent messages  
Identifier: MQIAMO64\_BROWSE\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QAccountingData*  
Returned: When available

### *BrowseMinBytes*

Description: The size of the smallest persistent and nonpersistent message browsed from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_BROWSE\_MIN\_BYTES  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

### *BrowseMaxBytes*

Description: The size of the largest persistent and nonpersistent message browsed from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_BROWSE\_MAX\_BYTES  
Data type: MQCFIL  
Included in PCF group: *QAccountingData*  
Returned: When available

### *CBCount*

Description: The number of successful MQCB requests. This is an array of values indexed by the type of operation

- 0 – A callback was created or altered
- 1 – A callback was removed
- 2 – A callback was resumed
- 3 – A callback was suspended

Identifier: MQIAMO\_CBS  
Data type: MQCFIN  
Returned: When available.

### *CBFailCount*

Description: The number of unsuccessful MQCB requests.  
Identifier: MQIAMO\_CBS\_FAILED  
Data type: MQCFIN  
Returned: When available.

### *TimeOnQMin*

**Description:** The shortest time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2.

**Identifier:** MQIAMO64\_Q\_TIME\_MIN

**Data type:** MQCFIL64

**Included in PCF group:** *QAccountingData*

**Returned:** When available

#### *TimeOnQAvg*

**Description:** The average time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2.

**Identifier:** MQIAMO64\_Q\_TIME\_AVG

**Data type:** MQCFIL64

**Included in PCF group:** *QAccountingData*

**Returned:** When available

#### *TimeOnQMax*

**Description:** The longest time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2.

**Identifier:** MQIAMO64\_Q\_TIME\_MAX

**Data type:** MQCFIL64

**Included in PCF group:** *QAccountingData*

**Returned:** When available

## **MQI statistics message data**

Use this page to view the structure of an MQI statistics message

|                      |                                    |
|----------------------|------------------------------------|
| <b>Message name:</b> | MQI statistics message.            |
| <b>Platforms:</b>    | All, except WebSphere MQ for z/OS. |
| <b>System queue:</b> | SYSTEM.ADMIN.STATISTICS.QUEUE.     |

#### *QueueManager*

**Description:** Name of the queue manager.

**Identifier:** MQCA\_Q\_MGR\_NAME.

**Data type:** MQCFST.

**Maximum length:** MQ\_Q\_MGR\_NAME\_LENGTH.

**Returned:** Always.

#### *IntervalStartDate*

Description: The date at the start of the monitoring period.  
Identifier: MQCAMO\_START\_DATE.  
Data type: MQCFST.  
Maximum length: MQ\_DATE\_LENGTH  
Returned: Always.

#### *IntervalStartTime*

Description: The time at the start of the monitoring period.  
Identifier: MQCAMO\_START\_TIME.  
Data type: MQCFST.  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always.

#### *IntervalEndDate*

Description: The date at the end of the monitoring period.  
Identifier: MQCAMO\_END\_DATE.  
Data type: MQCFST.  
Maximum length: MQ\_DATE\_LENGTH  
Returned: Always.

#### *IntervalEndTime*

Description: The time at the end of the monitoring period.  
Identifier: MQCAMO\_END\_TIME.  
Data type: MQCFST.  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always.

#### *CommandLevel*

Description: The queue manager command level.  
Identifier: MQIA\_COMMAND\_LEVEL.  
Data type: MQCFIN.  
Returned: Always.

#### *ConnCount*

Description: The number of successful connections to the queue manager.  
Identifier: MQIAMO\_CONNS.  
Data type: MQCFIN.  
Returned: When available.

#### *ConnFailCount*

Description: The number of unsuccessful connection attempts.  
Identifier: MQIAMO\_CONNS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *ConnsMax*

Description: The maximum number of concurrent connections in the recording interval.  
Identifier: MQIAMO\_CONNS\_MAX.  
Data type: MQCFIN.  
Returned: When available.

#### *DiscCount*

Description: The number of disconnects from the queue manager. This is an integer array, indexed by the following constants:

- MQDISCONNECT\_NORMAL
- MQDISCONNECT\_IMPLICIT
- MQDISCONNECT\_Q\_MGR

Identifier: MQIAMO\_DISCS.  
Data type: MQCFIL.  
Returned: When available.

#### *OpenCount*

Description: The number of objects successfully opened. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_OPENS.  
Data type: MQCFIL.  
Returned: When available.

#### *OpenFailCount*

Description: The number of unsuccessful open object attempts. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_OPENS\_FAILED.  
Data type: MQCFIL.  
Returned: When available.

#### *CloseCount*

Description: The number of objects successfully closed. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_CLOSES.  
Data type: MQCFIL.  
Returned: When available.

#### *CloseFailCount*

Description: The number of unsuccessful close object attempts. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_CLOSES\_FAILED.  
Data type: MQCFIL.  
Returned: When available.

#### *InqCount*

Description: The number of objects successfully inquired upon. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_INQS.  
Data type: MQCFIL.  
Returned: When available.

#### *InqFailCount*

Description: The number of unsuccessful object inquire attempts. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_INQS\_FAILED.  
Data type: MQCFIL.  
Returned: When available.

#### *SetCount*

Description: The number of objects successfully updated (SET). This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_SETS.  
Data type: MQCFIL.  
Returned: When available.

#### *SetFailCount*

Description: The number of unsuccessful SET attempts. This parameter is an integer list indexed by object type, see Reference note 1.  
Identifier: MQIAMO\_SETS\_FAILED.  
Data type: MQCFIL.  
Returned: When available.

#### *PutCount*

Description: The number of persistent and nonpersistent messages successfully put to a queue, with the exception of MQPUT1 requests. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO\_PUTS.  
Data type: MQCFIL.  
Returned: When available.

#### *PutFailCount*

Description: The number of unsuccessful put message attempts.  
Identifier: MQIAMO\_PUTS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *Put1Count*

Description: The number of persistent and nonpersistent messages successfully put to a queue using MQPUT1 requests. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO\_PUT1S.  
Data type: MQCFIL.  
Returned: When available.

#### *Put1FailCount*

Description: The number of unsuccessful attempts to put a persistent and nonpersistent message to a queue using MQPUT1 requests. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO\_PUT1S\_FAILED.  
Data type: MQCFIL.  
Returned: When available.

#### *PutBytes*

Description: The number bytes for persistent and nonpersistent messages written in using put requests. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO64\_PUT\_BYTES.  
Data type: MQCFIL64.  
Returned: When available.

#### *GetCount*

Description: The number of successful destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO\_GETS.  
Data type: MQCFIL.  
Returned: When available.

#### *GetFailCount*

Description: The number of unsuccessful destructive get requests.  
Identifier: MQIAMO\_GETS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *GetBytes*

Description: The number of bytes read in destructive gets requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO64\_GET\_BYTES.  
Data type: MQCFIL64.  
Returned: When available.

#### *BrowseCount*

Description: The number of successful non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO\_BROWSES.  
Data type: MQCFIL.  
Returned: When available.

#### *BrowseFailCount*

Description: The number of unsuccessful non-destructive get requests.  
Identifier: MQIAMO\_BROWSES\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *BrowseBytes*

Description: The number of bytes read in non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2  
Identifier: MQIAMO64\_BROWSE\_BYTES.  
Data type: MQCFIL64.  
Returned: When available.

#### *CommitCount*

Description: The number of transactions successfully completed. This number includes transactions committed implicitly by the application disconnecting, and commit requests where there is no outstanding work.  
Identifier: MQIAMO\_COMMITS.  
Data type: MQCFIN.  
Returned: When available.

#### *CommitFailCount*

Description: The number of unsuccessful attempts to complete a transaction.  
Identifier: MQIAMO\_COMMITS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *BackCount*

Description: The number of backouts processed, including implicit backout upon abnormal disconnect.  
Identifier: MQIAMO\_BACKOUTS.  
Data type: MQCFIN.  
Returned: When available.

#### *ExpiredMsgCount*

Description: The number of persistent and nonpersistent messages that were discarded because they had expired, before they could be retrieved.  
Identifier: MQIAMO\_MSGS\_EXPIRED.  
Data type: MQCFIN.  
Returned: When available.

#### *PurgeCount*

Description: The number of times the queue has been cleared.  
Identifier: MQIAMO\_MSGS\_PURGED.  
Data type: MQCFIN.  
Returned: When available.

#### *SubCountDur*

Description: The number of successful Subscribe requests which created, altered or resumed durable subscriptions. This is an array of values indexed by the type of operation  
  
0 = The number of subscriptions created  
  
1 = The number of subscriptions altered  
  
2 = The number of subscriptions resumed  
Identifier: MQIAMO\_SUBS\_DUR.  
Data type: MQCFIL  
Returned: When available.

#### *SubCountNDur*

Description: The number of successful Subscribe requests which created, altered or resumed non-durable subscriptions. This is an array of values indexed by the type of operation  
  
0 = The number of subscriptions created  
  
1 = The number of subscriptions altered  
  
2 = The number of subscriptions resumed  
Identifier: MQIAMO\_SUBS\_NDUR.  
Data type: MQCFIL.  
Returned: When available.

#### *SubFailCount*

Description: The number of unsuccessful Subscribe requests.  
Identifier: MQIAMO\_SUBS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *UnsubCountDur*

Description: The number of succesful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation  
0 – The subscription was closed but not removed  
1 – The subscription was closed and removed  
Identifier: MQIAMO\_UNSUBS\_DUR.  
Data type: MQCFIL.  
Returned: When available.

#### *UnsubCountNDur*

Description: The number of succesful unsubscribe requests for non-durable subscriptions. This is an array of values indexed by the type of operation  
0 – The subscription was closed but not removed  
1 – The subscription was closed and removed  
Identifier: MQIAMO\_UNSUBS\_NDUR.  
Data type: MQCFIL.  
Returned: When available.

#### *UnsubFailCount*

Description: The number of failed unsubscribe requests.  
Identifier: MQIAMO\_UNSUBS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *SubRqCount*

Description: The number of successful MQSUBRQ requests.  
Identifier: MQIAMO\_SUBRQS  
Data type: MQCFIN  
Returned: When available.

#### *SubRqFailCount*

Description: The number of unsuccessful MQSUBRQ requests.  
Identifier: MQIAMO\_SUBRQS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *CBCCount*

Description: The number of successful MQCB requests. This is an array of values indexed by the type of operation

- 0 – A callback was created or altered
- 1 – A callback was removed
- 2 – A callback was resumed
- 3 – A callback was suspended

Identifier: MQIAMO\_CBS.  
Data type: MQCFIL.  
Returned: When available.

#### *CBFailCount*

Description: The number of unsuccessful MQCB requests.  
Identifier: MQIAMO\_CBS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *CtlCount*

Description: The number of unsuccessful MQCB requests.

- 0 – The connection was started
- 1 – The connection was stopped
- 2 – The connection was resumed
- 3 – The connection was suspended

Identifier: MQIAMO\_CTL.  
Data type: MQCFIL.  
Returned: When available.

#### *CtlFailCount*

Description: The number of unsuccessful MQCTL requests.  
Identifier: MQIAMO\_CTL\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *StatCount*

Description: The number of successful MQSTAT requests.  
Identifier: MQIAMO\_STATS.  
Data type: MQCFIN.  
Returned: When available.

#### *StatFailCount*

Description: The number of unsuccessful MQSTAT requests.  
Identifier: MQIAMO\_STATS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *SubCountDurHighWater*

Description: The high-water mark on the number of durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE

- 0 – The high-water mark for all durable subscriptions in the system
- 1 – The high-water mark for durable application subscriptions (MQSUBTYPE\_API)
- 2 – The high-water mark for durable admin subscription (MQSUBTYPE\_ADMIN)
- 3 – The high-water mark for durable proxy subscriptions (MQSUBTYPE\_PROXY)

Identifier: MQIAMO\_SUB\_DUR\_HIGHWATER  
Data type: MQCFIL.  
Returned: When available.

#### *SubCountDurLowWater*

Description: The low-water mark on the number of durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE.

- 0 – The low-water mark for all durable subscriptions in the system
- 1 – The low-water mark for durable application subscriptions (MQSUBTYPE\_API)
- 2 – The low-water mark for durable admin subscriptions (MQSUBTYPE\_ADMIN)
- 3 – The low-water mark for durable proxy subscriptions (MQSUBTYPE\_PROXY)

Identifier: MQIAMO\_SUB\_DUR\_LOWWATER  
Data type: MQCFIL.  
Returned: When available.

#### *SubCountNDurHighWater*

Description: The high-water mark on the number of non-durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE

- 0 – The high-water mark for all non-durable subscriptions in the system
- 1 – The high-water mark for non-durable application subscriptions (MQSUBTYPE\_API)
- 2 – The high-water mark for non-durable admin subscription (MQSUBTYPE\_ADMIN)
- 3 – The high-water mark for non-durable proxy subscriptions (MQSUBTYPE\_PROXY)

Identifier: MQIAMO\_SUB\_NDUR\_HIGHWATER  
Data type: MQCFIL.  
Returned: When available.

#### *SubCountNDurLowWater*

Description: The low-water mark on the number of non-durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE.

- 0 – The low-water mark for all non-durable subscriptions in the system
- 1 – The low-water mark for non-durable application subscriptions (MQSUBTYPE\_API)
- 2 – The low-water mark for non-durable admin subscriptions (MQSUBTYPE\_ADMIN)
- 3 – The low-water mark for non-durable proxy subscriptions (MQSUBTYPE\_PROXY)

Identifier: MQIAMO\_SUB\_NDUR\_LOWWATER  
Data type: MQCFIL.  
Returned: When available.

#### *PutTopicCount*

Description: The number persistent and nonpersistent messages successfully put to a topic, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see Reference note 2.

Note: Messages put using a queue alias which resolve to a topic are included in this value.

Identifier: MQIAMO\_TOPIC\_PUTS.  
Data type: MQCFIL.  
Returned: When available.

#### *PutTopicFailCount*

Description: The number of unsuccessful attempts to put a message to a topic.  
Identifier: MQIAMO\_TOPIC\_PUTS\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *Put1TopicCount*

Description: The number of persistent and nonpersistent messages successfully put to a topic using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2.

Note: Messages put using a queue alias which resolve to a topic are included in this value.

Identifier: MQIAMO\_TOPIC\_PUT1S.  
Data type: MQCFIL.  
Returned: When available.

#### *Put1TopicFailCount*

Description: The number of unsuccessful attempts to put a message to a topic using MQPUT1 calls.  
Identifier: MQIAMO\_TOPIC\_PUT1S\_FAILED.  
Data type: MQCFIN.  
Returned: When available.

#### *PutTopicBytes*

Description: The number bytes written using put calls for persistent and nonpersistent messages which resolve to a publish operation. This is number of bytes put by the application and not the resultant number of bytes delivered to subscribers, see PublishMsgBytes for this value. This parameter is an integer list indexed by persistence value, see Reference note 2.

Identifier: MQIAMO64\_TOPIC\_PUT\_BYTES.

Data type: MQCFIL64.

Returned: When available.

#### *PublishMsgCount*

Description: The number of messages delivered to subscriptions in the time interval. This parameter is an integer list indexed by persistence value, see Reference note 2.

Identifier: MQIAMO64\_PUBLISH\_MSG\_COUNT

Data type: MQCFIL.

Returned: When available.

#### *PublishMsgBytes*

Description: The number of bytes delivered to subscriptions in the time interval. This parameter is an integer list indexed by persistence value, see Reference note 2.

Identifier: MQIAMO64\_PUBLISH\_MSG\_BYTES

Data type: MQCFIL64.

Returned: When available.

### **Queue statistics message data**

Use this page to view the structure of a queue statistics message

|               |                                    |
|---------------|------------------------------------|
| Message name: | Queue statistics message.          |
| Platforms:    | All, except WebSphere MQ for z/OS. |
| System queue: | SYSTEM.ADMIN.STATISTICS.QUEUE.     |

#### *QueueManager*

Description: Name of the queue manager

Identifier: MQCA\_Q\_MGR\_NAME

Data type: MQCFST

Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

Returned: Always

#### *IntervalStartDate*

Description: The date at the start of the monitoring period

Identifier: MQCAMO\_START\_DATE

Data type: MQCFST

Maximum length: MQ\_DATE\_LENGTH

Returned: Always

#### *IntervalStartTime*

Description: The time at the start of the monitoring period  
Identifier: MQCAMO\_START\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

#### *IntervalEndDate*

Description: The date at the end of the monitoring period  
Identifier: MQCAMO\_END\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH  
Returned: Always

#### *IntervalEndTime*

Description: The time at the end of the monitoring period  
Identifier: MQCAMO\_END\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

#### *CommandLevel*

Description: The queue manager command level  
Identifier: MQIA\_COMMAND\_LEVEL  
Data type: MQCFIN  
Returned: Always

#### *ObjectCount*

Description: The number of queue objects accessed in the interval for which statistics data has been recorded. This value is set to the number of QStatisticsData PCF groups contained in the message.  
Identifier: MQIAMO\_OBJECT\_COUNT  
Data type: MQCFIN  
Returned: Always

#### *QStatisticsData*

Description: Grouped parameters specifying statistics details for a queue  
Identifier: MQGACF\_Q\_STATISTICS\_DATA  
Data type: MQCFGR

Parameters in group:

- QName*
- CreateDate*
- CreateTime*
- QType*
- QDefinitionType*
- QMinDepth*
- QMaxDepth*
- AvgTimeOnQ*
- PutCount*
- PutFailCount*
- Put1Count*
- Put1FailCount*
- PutBytes*
- GetCount*
- GetFailCount*
- GetBytes*
- BrowseCount*
- BrowseFailCount*
- BrowseBytes*
- NonQueuedMsgCount*
- ExpiredMsgCount*
- PurgeCount*

Returned: Always

*QName*

Description: The name of the queue  
 Identifier: MQCA\_Q\_NAME  
 Data type: MQCFST  
 Maximum length: MQ\_Q\_NAME\_LENGTH  
 Returned: Always

*CreateDate*

Description: The date when the queue was created  
 Identifier: MQCA\_CREATION\_DATE  
 Data type: MQCFST  
 Maximum length: MQ\_DATE\_LENGTH  
 Returned: Always

*CreateTime*

Description: The time when the queue was created  
Identifier: MQCA\_CREATION\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always

### *QType*

Description: The type of the queue  
Identifier: MQIA\_Q\_TYPE  
Data type: MQCFIN  
Value: MQOT\_LOCAL  
Returned: Always

### *QDefinitionType*

Description: The queue definition type  
Identifier: MQIA\_DEFINITION\_TYPE  
Data type: MQCFIN  
Values: Possible values are

- MQQDT\_PREDEFINED
- MQQDT\_PERMANENT\_DYNAMIC
- MQQDT\_TEMPORARY\_DYNAMIC

Returned: When available

### *QMinDepth*

Description: The minimum queue depth during the monitoring period  
Identifier: MQIAMO\_Q\_MIN\_DEPTH  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *QMaxDepth*

Description: The maximum queue depth during the monitoring period  
Identifier: MQIAMO\_Q\_MAX\_DEPTH  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *AvgTimeOnQ*

Description: The average latency, in microseconds, of messages destructively retrieved from the queue during the monitoring period. This parameter is an integer list indexed by persistence value, see Reference note 2.  
Identifier: MQIAMO64\_AVG\_Q\_TIME  
Data type: MQCFIL64  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *PutCount*

Description: The number of persistent and nonpersistent messages successfully put to the queue, with exception of MQPUT1 requests. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO\_PUTS  
Data type: MQCFIL  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *PutFailCount*

Description: The number of unsuccessful attempts to put a message to the queue  
Identifier: MQIAMO\_PUTS\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *Put1Count*

Description: The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO\_PUT1S  
Data type: MQCFIL  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *Put1FailCount*

Description: The number of unsuccessful attempts to put a message using MQPUT1 calls  
Identifier: MQIAMO\_PUT1S\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *PutBytes*

Description: The number of bytes written in put requests to the queue  
Identifier: MQIAMO64\_PUT\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *GetCount*

Description: The number of successful destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO\_GETS  
Data type: MQCFIL  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *GetFailCount*

Description: The number of unsuccessful destructive get requests  
Identifier: MQIAMO\_GETS\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *GetBytes*

Description: The number of bytes read in destructive put requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO64\_GET\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *BrowseCount*

Description: The number of successful non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO\_BROWSES  
Data type: MQCFIL  
Included in PCF group: *QStatisticsData*  
Returned: When available

#### *BrowseFailCount*

Description: The number of unsuccessful non-destructive get requests  
Identifier: MQIAMO\_BROWSES\_FAILED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *BrowseBytes*

Description: The number of bytes read in non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2.  
Identifier: MQIAMO64\_BROWSE\_BYTES  
Data type: MQCFIL64  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *NonQueuedMsgCount*

Description: The number of messages that bypassed the queue and were transferred directly to a waiting application.  
  
Bypassing a queue can only occur in certain circumstances. This number represents how many times WebSphere MQ was able to bypass the queue, and not the number of times an application was waiting.  
Identifier: MQIAMO\_MSGS\_NOT\_QUEUED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *ExpiredMsgCount*

Description: The number of persistent and nonpersistent messages that were discarded because they had expired before they could be retrieved.  
Identifier: MQIAMO\_MSGS\_EXPIRED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *PurgeCount*

Description: The number of messages purged.  
Identifier: MQIAMO\_MSGS\_PURGED  
Data type: MQCFIN  
Included in PCF group: *QStatisticsData*  
Returned: When available

### *CBCount*

Description: The number of successful MQCB requests. This is an array of values indexed by the type of operation

- 0 – A callback was created or altered
- 1 – A callback was removed
- 2 – A callback was resumed
- 3 – A callback was suspended

Identifier: MQIAMO\_CBS  
 Data type: MQCFIN  
 Returned: When available.

*CBFailCount*

Description: The number of unsuccessful MQCB requests.  
 Identifier: MQIAMO\_CBS\_FAILED  
 Data type: MQCFIN  
 Returned: When available.

**Channel statistics message data**

Use this page to view the structure of a channel statistics message

|               |                                    |
|---------------|------------------------------------|
| Message name: | Channel statistics message.        |
| Platforms:    | All, except WebSphere MQ for z/OS. |
| System queue: | SYSTEM.ADMIN.STATISTICS.QUEUE.     |

*QueueManager*

Description: The name of the queue manager.  
 Identifier: MQCA\_Q\_MGR\_NAME.  
 Data type: MQCFST.  
 Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH.  
 Returned: Always.

*IntervalStartDate*

Description: The date at the start of the monitoring period.  
 Identifier: MQCAMO\_START\_DATE.  
 Data type: MQCFST.  
 Maximum length: MQ\_DATE\_LENGTH.  
 Returned: Always.

*IntervalStartTime*

Description: The time at the start of the monitoring period.  
Identifier: MQCAMO\_START\_TIME.  
Data type: MQCFST.  
Maximum length: MQ\_TIME\_LENGTH.  
Returned: Always.

#### *IntervalEndDate*

Description: The date at the end of the monitoring period  
Identifier: MQCAMO\_END\_DATE.  
Data type: MQCFST.  
Maximum length: MQ\_DATE\_LENGTH.  
Returned: Always.

#### *IntervalEndTime*

Description: The time at the end of the monitoring period  
Identifier: MQCAMO\_END\_TIME.  
Data type: MQCFST.  
Maximum length: MQ\_TIME\_LENGTH  
Returned: Always.

#### *CommandLevel*

Description: The queue manager command level.  
Identifier: MQIA\_COMMAND\_LEVEL.  
Data type: MQCFIN.  
Returned: Always.

#### *ObjectCount*

Description: The number of Channel objects accessed in the interval for which statistics data has been recorded. This value is set to the number of ChlStatisticsData PCF groups contained in the message.  
Identifier: MQIAMO\_OBJECT\_COUNT  
Data type: MQCFIN.  
Returned: Always.

#### *ChlStatisticsData*

Description: Grouped parameters specifying statistics details for a channel.  
Identifier: MQGACF\_CHL\_STATISTICS\_DATA.  
Data type: MQCFGR.

Parameters in group:

- ChannelName*
- ChannelType*
- RemoteQmgr*
- ConnectionName*
- MsgCount*
- TotalBytes*
- NetTimeMin*
- NetTimeAvg*
- NetTimeMax*
- ExitTimeMin*
- ExitTimeAvg*
- ExitTimeMax*
- FullBatchCount*
- IncplBatchCount*
- AverageBatchSize*
- PutRetryCount*

Returned: Always.

### *ChannelName*

Description: The name of the channel.  
Identifier: MQCACH\_CHANNEL\_NAME.  
Data type: MQCFST.  
Maximum length: MQ\_CHANNEL\_NAME\_LENGTH.  
Returned: Always.

### *ChannelType*

Description: The channel type.  
Identifier: MQIACH\_CHANNEL\_TYPE.  
Data type: MQCFIN.  
Values: Possible values are:  
**MQCHT\_SENDER**  
Sender channel.  
**MQCHT\_SERVER**  
Server channel.  
**MQCHT\_RECEIVER**  
Receiver channel.  
**MQCHT\_REQUESTER**  
Requester channel.  
**MQCHT\_CLUSRCVR**  
Cluster receiver channel.  
**MQCHT\_CLUSSDR**  
Cluster sender channel.  
Returned: Always.

### *RemoteQmgr*

Description: The name of the remote queue manager.  
Identifier: MQCA\_REMOTE\_Q\_MGR\_NAME.  
Data type: MQCFST.  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH  
Returned: When available.

#### *ConnectionName*

Description: Connection name of remote queue manager.  
Identifier: MQCACH\_CONNECTION\_NAME.  
Data type: MQCFST  
Maximum length: MQ\_CONN\_NAME\_LENGTH  
Returned: When available.

#### *MsgCount*

Description: The number of persistent and nonpersistent messages sent or received.  
Identifier: MQIAMO\_MSGS.  
Data type: MQCFIN  
Returned: When available.

#### *TotalBytes*

Description: The number of bytes sent or received for persistent and nonpersistent messages.  
Identifier: MQIAMO64\_BYTES.  
Data type: MQCFIN64.  
Returned: When available.

#### *NetTimeMin*

Description: The shortest recorded channel round trip measured in the recording interval, in microseconds.  
Identifier: MQIAMO\_NET\_TIME\_MIN.  
Data type: MQCFIN.  
Returned: When available.

#### *NetTimeAvg*

Description: The average recorded channel round trip measured in the recording interval, in microseconds.  
Identifier: MQIAMO\_NET\_TIME\_AVG.  
Data type: MQCFIN.  
Returned: When available.

#### *NetTimeMax*

Description: The longest recorded channel round trip measured in the recording interval, in microseconds.  
Identifier: MQIAMO\_NET\_TIME\_MAX.  
Data type: MQCFIN.  
Returned: When available.

#### *ExitTimeMin*

Description: The shortest recorded time, in microseconds, spent executing a user exit in the recording interval,  
Identifier: MQIAMO\_EXIT\_TIME\_MIN.  
Data type: MQCFIN.  
Returned: When available.

#### *ExitTimeAvg*

Description: The average recorded time, in microseconds, spent executing a user exit in the recording interval. Measured in microseconds.  
Identifier: MQIAMO\_EXIT\_TIME\_AVG.  
Data type: MQCFIN.  
Returned: When available.

#### *ExitTimeMax*

Description: The longest recorded time, in microseconds, spent executing a user exit in the recording interval. Measured in microseconds.  
Identifier: MQIAMO\_EXIT\_TIME\_MAX.  
Data type: MQCFIN.  
Returned: When available.

#### *FullBatchCount*

Description: The number of batches processed by the channel that were sent because the value of the channel attributes BATCHSZ or BATCHLIM was reached.  
Identifier: MQIAMO\_FULL\_BATCHES.  
Data type: MQCFIN.  
Returned: When available.

#### *IncplBatchCount*

Description: The number of batches processed by the channel, that were sent without the value of the channel attribute BATCHSZ being reached.  
Identifier: MQIAMO\_INCOMPLETE\_BATCHES.  
Data type: MQCFIN.  
Returned: When available.

#### *AverageBatchSize*

Description: The average batch size of batches processed by the channel.  
 Identifier: MQIAMO\_AVG\_BATCH\_SIZE.  
 Data type: MQCFIN.  
 Returned: When available.

### *PutRetryCount*

Description: The number of times in the time interval that a message failed to be put, and entered a retry loop.  
 Identifier: MQIAMO\_PUT\_RETRIES.  
 Data type: MQCFIN.  
 Returned: When available.

## Reference notes

Use this page to view the notes to which descriptions of the structure of accounting and statistics messages refer

The following message data descriptions refer to these notes:

- “MQI accounting message data” on page 985
  - “Queue accounting message data” on page 996
  - “MQI statistics message data” on page 1006
  - “Queue statistics message data” on page 1017
  - “Channel statistics message data” on page 1024
1. This parameter relates to WebSphere MQ objects. This parameter is an array of values (MQCFIL or MQCFIL64) indexed by the following constants:

*Table 111. Array indexed by object type*

| Object type        | Value context                                                      |
|--------------------|--------------------------------------------------------------------|
| MQOT_Q (1)         | Contains the value relating to queue objects.                      |
| MQOT_NAMELIST (2)  | Contains the value relating to namelist objects.                   |
| MQOT_PROCESS (3)   | Contains the value relating to process objects.                    |
| MQOT_Q_MGR (5)     | Contains the value relating to queue manager objects.              |
| MQOT_CHANNEL (6)   | Contains the value relating to channel objects.                    |
| MQOT_AUTH_INFO (7) | Contains the value relating to authentication information objects. |
| MQOT_TOPIC (8)     | Contains the value relating to topic objects.                      |

**Note:** An array of 13 MQCFIL or MQCFIL64 values are returned but only those listed are meaningful.

2. This parameter relates to WebSphere MQ messages. This parameter is an array of values (MQCFIL or MQCFIL64) indexed by the following constants:

Table 112. Array indexed by persistence value

| Constant | Value                                          |
|----------|------------------------------------------------|
| 1        | Contains the value for nonpersistent messages. |
| 2        | Contains the value for persistent messages.    |

**Note:** The index for each of these arrays starts at zero, so an index of 1 refers to the second row of the array. Elements of these arrays not listed in these tables contain no accounting or statistics information.

## Application activity trace

Application activity trace produces detailed information about the behavior of applications connected to a queue manager. It traces the behavior of an application and provides a detailed view of the parameters used by an application as it interacts with IBM WebSphere MQ resources. It also shows the sequence of MQI calls issued by an application.

Use Application activity trace when you require more information than is provided by Event monitoring, Message monitoring, Accounting and statistics messages, and Real-time monitoring.

Note that activity trace is not supported by IBM WebSphere MQ for z/OS.

## Collecting application activity trace information

An application activity trace message is a PCF message. You configure activity trace using a configuration file. To collect application activity trace information you set the ACTVTRC queue manager attribute. You can override this setting at connection level using MQCONN options, or at application stanza level using the activity trace configuration file.

### About this task

Activity trace messages are composed of an MQMD structure: a PCF (MQCFH) header structure, followed by a number of PCF parameters. A sequence of ApplicationTraceData PCF groups follows the PCF parameters. These PCF groups collect information about the MQI operations that an application performs while connected to a queue manager. You configure activity trace using a configuration file called `mqat.ini`.

To control whether or not application activity trace information is collected, you configure one or more of the following settings:

1. The ACTVTRC queue manager attribute.
2. The ACTVCONO settings (in the MQCNO structure passed in MQCONN).
3. The matching stanza for the application in the activity trace configuration file `mqat.ini`.

The previous sequence is significant. The ACTVTRC attribute is overridden by the ACTVCONO settings, which are overridden by the settings in the `mqat.ini` file.

Trace entries are written after each operation has completed, unless otherwise stated. These entries are first written to the system queue `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`, then written to application activity trace messages when the application disconnects from the queue manager. For long running applications, intermediate messages are written if any of the following events occurs:

- The lifetime of the connection reaches a defined timeout value.
- The number of operations reaches a specified number.
- The amount of data collected in memory reaches the maximum message length allowed for the queue.

You set the timeout value using the **ActivityInterval** parameter. You set the number of operations using the **ActivityCount** parameter. Both parameters are specified in the activity trace configuration file `mqat.ini`.

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See “Tuning the performance impact of application activity trace” on page 1036.

## Procedure

1. “Setting **ACTVTRC** to control collection of activity trace information.”
2. “Setting **MQCONN** options to control collection of activity trace information.”
3. “Configuring activity trace behavior using `mqat.ini`” on page 1032.
4. “Tuning the performance impact of application activity trace” on page 1036.

## Setting **ACTVTRC** to control collection of activity trace information

Use the queue manager attribute **ACTVTRC** to control the collection of MQI application activity trace information

### About this task

Application activity trace messages are generated only for connections that begin after application activity trace is enabled. The **ACTVTRC** parameter can have the following values:

**ON** API activity trace collection is switched on

**OFF**

API activity trace collection is switched off

**Note:** The **ACTVTRC** setting can be overridden by the queue manager **ACTVCONO** parameter. If you set the **ACTVCONO** parameter to **ENABLED**, then the **ACTVTRC** setting can be overridden for a given connection using the **Options** field in the **MQCNO** structure. See “Setting **MQCONN** options to control collection of activity trace information.”

### Example

To change the value of the **ACTVTRC** parameter, you use the MQSC command **ALTER QMGR**. For example, to enable MQI application activity trace information collection use the following MQSC command:

```
ALTER QMGR ACTVTRC(ON)
```

### What to do next

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See “Tuning the performance impact of application activity trace” on page 1036.

## Setting **MQCONN** options to control collection of activity trace information

If the queue manager attribute **ACTVCONO** is set to **ENABLED**, you can use the **ConnectOpts** parameter on the **MQCONN** call to enable or disable application activity reports on a per connection basis. These options override the activity trace behavior defined by the queue manager attribute **ACTVTRC**, and can be overridden by settings in the activity trace configuration file `mqat.ini`.

## Procedure

1. Set the queue manager attribute **ACTVCONO** to **ENABLED**.

**Note:** If an application attempts to modify the accounting behavior of an application using the **ConnectOpts** parameter, and the QMGR attribute **ACTVCONO** is set to **DISABLED**, then no error is returned to the application, and activity trace collection is defined by the queue manager attributes or the activity trace configuration file `mqat.ini`.

2. Set the **ConnectOpts** parameter on the MQCONNX call to `MQCNO_ACTIVITY_TRACE_ENABLED`.  
The **ConnectOpts** parameter on the MQCONNX call can have the following values:

**MQCNO\_ACTIVITY\_TRACE\_DISABLED**

Activity trace is switched off for the connection.

**MQCNO\_ACTIVITY\_TRACE\_ENABLED**

Activity trace is switched on for the connection.

**Note:** If an application selects both `MQCNO_ACTIVITY_TRACE_ENABLED` and `MQCNO_ACTIVITY_TRACE_DISABLED` for MQCONNX, the call fails with a reason code of `MQRC_OPTIONS_ERROR`.

3. Check that these activity trace settings are not being overridden by settings in the activity trace configuration file `mqat.ini`.  
See “Configuring activity trace behavior using `mqat.ini`.”

## What to do next

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See “Tuning the performance impact of application activity trace” on page 1036.

## Configuring activity trace behavior using `mqat.ini`

The activity trace behavior is configured using a configuration file called `mqat.ini`. This file follows the same stanza key and parameter-value pair format as the `mq5.ini` and `qm.ini` files.

## About this task

On UNIX and Linux systems, `mqat.ini` is located in the queue manager data directory, which is the same location as the `qm.ini` file.

On Windows systems, `mqat.ini` is located in the queue manager data directory `C:\Program Files\IBM\WebSphere MQ\qmgrs\queue_manager_name`. Users running applications to be traced need permission to read this file.

The syntax rules for the format of the file are:

- Text beginning with a hash or semi-colon is considered to be a comment which extends to the end of the line.
- The first significant (non-comment) line must be a stanza key.
- A stanza key consists of the name of the stanza followed by a colon.
- A parameter-value pair consists of the name of a parameter followed by an equals sign and then the value.
- Only a single parameter-value pair can appear on a line. (A parameter-value must not wrap onto another line).
- Leading and trailing whitespace is ignored. There is no limit on the amount of white space between stanza names, parameter names and values, or parameter/value pairs. Line breaks are significant and not ignored.
- The maximum length for any line is 2048 characters.
- The stanza keys, parameter names, and constant parameter values are not case-sensitive, but the variable parameter values (**ApplName** and **DebugPath**) are case-sensitive.

## Stanza keys

Two stanza key types are allowed in the configuration file: the AllActivityTrace stanza, and the ApplicationTrace stanza

### AllActivityTrace stanza

The AllActivityTrace stanza defines settings for the activity trace that is applied to all IBM WebSphere MQ connections unless overridden.

Individual values in the AllActivityTrace stanza can be overridden by more specific information in an ApplicationTrace stanza.

If more than one AllActivityTrace stanza is specified then the values in the last stanza is used. Parameters missing from the chosen AllActivityTrace take default values. Parameters and values from previous AllActivityTrace stanzas are ignored

### ApplicationTrace stanza

The ApplicationTrace stanza defines settings which can be applied to a specific name, type or both of IBM WebSphere MQ connection.

This stanza includes ApplName and ApplClass values which are used according to the matching rules defined in Connection Matching Rules to determine whether the stanza applies to a particular connection.

### Parameter/Value Pairs

The following table lists the parameter/value pairs which may be used in the activity trace configuration file.

| Name             | Stanza Type                          | Values (default in bold type) | Description                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------|--------------------------------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trace            | ApplicationTrace                     | ON / OFF                      | Activity trace switch. This switch can be used in the application-specific stanza to determine whether activity trace is active for the scope of the current application stanza. Note that this value overrides ACTVTRC and ACTVCONO settings for the queue manager.                                                                                                                                                     |
| ActivityInterval | AllActivityTrace<br>ApplicationTrace | 0-99999999 (0= <b>off</b> )   | Time interval in seconds between trace messages. Activity trace does not use a timer thread, so the trace message will not be written at the exact instant that the time elapses – rather it will be written when the first MQI operation is executed after the time interval has elapsed. If this value is 0 then the trace message is written when the connection disconnects (or when the activity count is reached). |
| ActivityCount    | AllActivityTrace<br>ApplicationTrace | 0-99999999 (0= <b>off</b> )   | Number of MQI or XA operations between trace messages. If this value is 0 then the trace message is written when the connection disconnects (or when the activity interval has elapsed).                                                                                                                                                                                                                                 |
| TraceLevel       | AllActivityTrace<br>ApplicationTrace | LOW / MEDIUM / HIGH           | Amount of parameter detail traced for each operation. The description of individual operations details which parameters are included for each trace level.                                                                                                                                                                                                                                                               |

| Name             | Stanza Type                          | Values (default in bold type)                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|--------------------------------------|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TraceMessageData | AllActivityTrace<br>ApplicationTrace | 0- 104 857 600 (100Mb)                                | Amount of message data traced in bytes for MQGET, MQPUT, MQPUT1, and Callback operations                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| ApplName         | ApplicationTrace                     | Character string<br>(Required parameter - no default) | This value is used to determine which applications the ApplicationTrace stanza applies to. It is matched to the ApplName value from the API exit context structure (which is equivalent to the MQMD.PutApplName). The content of the ApplName value varies according to the application environment. For distributed platforms, only the filename portion of the MQAXC.ApplName is matched to the value in the stanza. Characters to the left of the rightmost path separator are ignored when the comparison is made. For z/OS applications, the entire MQAXC.ApplName is matched to the value in the stanza. A single wildcard character (*) can be used at the end of the ApplName value to match any number of characters after that point are. If the ApplName value is set to a single wildcard character (*) then the ApplName value matches all applications. |
| ApplFunction     | ApplicationTrace<br>(IBM i only)     | Character string<br>(default *)                       | This value is used to qualify which application programs the ApplicationTrace stanza and ApplName value applies to. The stanza is optional, but is only valid for IBM i queue managers. A single wildcard character (*) can be used at the end of the ApplName value to match any number of characters. For example, an ApplicationTrace stanza specifying ApplName=* and ApplFunction=AMQSPUTO applies to all invocations of the AMQSPUTO program from any job.                                                                                                                                                                                                                                                                                                                                                                                                      |
| ApplClass        | ApplicationTrace                     | <b>USER</b> / MCA /<br>INTERNAL / ALL                 | The class of application. See the following table for an explanation of how the AppType values correspond to IBM WebSphere MQ connections                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

The following table shows how the AppClass values correspond to the APICallerType and APIEnvironment fields in the connection API exit context structure.

| APPLCLASS | API Caller Type: | API Environment:                                  | Description                                                              |
|-----------|------------------|---------------------------------------------------|--------------------------------------------------------------------------|
| USER      | MQXACT_EXTERNAL  | MQXE_OTHER                                        | Only user applications are traced                                        |
| MCA       | (Any value)      | MQXE_MCA<br>MQXE_MCA_CLNTCONN<br>MQXE_MCA_SVRCONN | Clients and channels (amqrmppa)                                          |
| INTERNAL  | MQXACT_EXTERNAL  | MQXE_COMMAND_SERVER<br>MQXE_MQSC                  | 'runmqsc' and command server                                             |
| INTERNAL  | MQXACT_INTERNAL  | (Any value)                                       | "trusted" and internal applications and processes – for example amqzdmaa |
| ALL       | (Any value)      | (Any value)                                       | All user and internal connections are traced                             |

## Connection Matching Rules

The queue manager applies the following rules to determine which stanzas settings to use for a connection.

1. A value specified in the AllActivityTrace stanza is used for the connection unless the value also occurs in an ApplicationTrace stanza and the stanza fulfills the matching criteria for the connection described in points 2, 3, and 4.
2. The ApplClass is matched against the type of the IBM WebSphere MQ connection. If the ApplClass does not match the connection type then the stanza is ignored for this connection.
3. The ApplName value in the stanza is matched against the file name portion of the ApplName field from the API exit context structure (MQAXC) for the connection. The file name portion is derived from the characters to the right of the final path separator (/ or \) character. If the stanza ApplName includes a wildcard (\*) then only the characters to the left of the wildcard are compared with the equivalent number of characters from the connections ApplName. For example, if a stanza value of "FRE\*" is specified then only the first three characters are used in the comparison, so "path/FREEDOM" and "path\FREDDY" match, but "path/FRIEND" does not. If the stanzas ApplName value does not match the connection ApplName then the stanza is ignored for this connection.
4. If more than one stanza matches the connections ApplName and ApplClass, then the stanza with the most specific ApplName is used. The most specific ApplName is defined as the one which uses the most characters to match the connections ApplName. For example, if the ini file contains a stanza with ApplName="FRE\*" and another stanza with ApplName="FREE\*" then the stanza with ApplName="FREE\*" is chosen as the best match for a connection with ApplName="path/FREEDOM" because it matches four characters (whereas ApplName="FRE\*" matches only three).
5. If after applying the rules in points 2, 3, and 4, there is more than one stanza that matches the connections ApplName and ApplClass, then the values from the last matching will be used and all other stanzas will be ignored.

## Application Activity Trace File Example

The following example shows how the configuration data is specified in the Activity Trace ini file. This example is shipped as a sample called `mqat.ini` in the C samples directory (the same directory as the `amqsact.c` file)

```

AllActivityTrace:
 ActivityInterval=0 # Time interval between trace messages
 # Values: 0-99999999 (0=off)
 # Default: 0
 ActivityCount=0 # Number of operations between trace msgs
 # Values: 0-99999999 (0=off)
 # Default: 0
 TraceLevel=MEDIUM # Amount of data traced for each operation
 # Values: LOW | MEDIUM | HIGH
 # Default: MEDIUM
 TraceMessageData=0 # Amount of message data traced
 # Values: 0-100000000
 # Default: 0

ApplicationTrace:
 ApplClass=USER # Application type
 # Values: (USER | MCA | INTERNAL | ALL)
 # Default: USER
 ApplName=AppName* # Application name (may be wildcarded)
 # (matched to app name without path)
 # Default: *
 Trace=OFF # Activity trace switch for application
 # Values: (ON | OFF)
 # Default: OFF
 ActivityInterval=0 # Time interval between trace messages
 # Values: 0-99999999 (0=off)
 # Default: 0
 ActivityCount=0 # Number of operations between trace msgs
 # Values: 0-99999999 (0=off)
 # Default: 0
 TraceLevel=MEDIUM # Amount of data traced for each operation
 # Values: LOW | MEDIUM | HIGH
 # Default: MEDIUM
 TraceMessageData=0 # Amount of message data traced
 # Values: 0-100000000
 # Default: 0

```

## What to do next

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See “Tuning the performance impact of application activity trace.”

## Tuning the performance impact of application activity trace

Enabling application activity trace can incur a performance penalty. This can be reduced by only tracing the applications that you need, by increasing the number of applications draining the queue, and by tuning **ActivityInterval**, **ActivityCount** and **TraceLevel** in `mqat.ini`.

## About this task

Enabling application activity trace selectively for an application or for all queue manager applications can result in additional messaging activity, and in the queue manager requiring additional storage space. In environments where messaging performance is critical, for example, in high workload applications or where a service level agreement (SLA) requires a minimum response time from the messaging provider, it might not be appropriate to collect application activity trace or it might be necessary to adjust the detail or frequency of trace activity messages that are produced. The preset values of **ActivityInterval**, **ActivityCount** and **TraceLevel** in the `mqat.ini` file give a default balance of detail and performance. However, you can tune these values to meet the precise functional and performance requirements of your system.

## Procedure

- Only trace the applications that you need.

Do this by creating an ApplicationTrace application-specific stanza in `mqat.ini`, or by changing the application to specify `MQCNO_ACTIVITY_TRACE_ENABLED` in the options field on the **MQCNO** structure on an MQCONN call. See “Configuring activity trace behavior using `mqat.ini`” on page 1032 and “Setting MQCONN options to control collection of activity trace information” on page 1031.

- Before starting trace, check that at least one application is running and is ready to retrieve the activity trace message data from the `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`.
- Keep the queue depth as low as possible, by increasing the number of applications draining the queue.
- Set the **TraceLevel** value in the `mqat.ini` file to collect the minimum amount of data required.  
TraceLevel=LOW has the lowest impact to messaging performance. See “Configuring activity trace behavior using `mqat.ini`” on page 1032.
- Tune the **ActivityCount** and **ActivityInterval** values in `mqat.ini`, to adjust how often activity trace messages are generated.

If you are tracing multiple applications, the activity trace messages might be being produced faster than they can be removed from the `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`. However, when you reduce how often activity trace messages are generated, you are also increasing the storage space required by the queue manager and the size of the messages when they are written to the queue.

## Application activity trace message reference

Use this page to obtain an overview of the format of application activity trace messages and the information returned in these messages

Application activity trace messages are standard IBM WebSphere MQ messages containing a message descriptor and message data. The message data contains information about the MQI operations performed by IBM WebSphere MQ applications, or information about the activities occurring in a IBM WebSphere MQ system.

### Message descriptor

- An MQMD structure

### Message data

- A PCF header (MQCFH)
- Application activity trace message data that is always returned
- Application activity trace message data that is operation-specific

### Related concepts:

“Variable parameters for application activity MQI operations” on page 1042

“Variable Parameters for Application Activity XA Operations” on page 1093

### Related reference:

“Application activity trace message MQMD (message descriptor)”

“MQCFH (PCF Header)” on page 1038

“Application activity trace message data” on page 1039

## Application activity trace message MQMD (message descriptor)

Use this page to understand the differences between the message descriptor of application activity trace messages and the message descriptor of event messages

The parameters and values in the message descriptor of application activity trace message are the same as in the message descriptor of event messages, with the following exception:

*Format*

Description: Format name of message data.  
Value: **MQFMT\_ADMIN**  
Admin message.

### *CorrelId*

Description: Correlation identifier.  
Value: Initialized with the ConnectionId of the application

## **MQCFH (PCF Header)**

Use this page to view the PCF values contained by the MQCFH structure for an activity trace message

For an activity trace message, the MQCFH structure contains the following values:

### *Type*

Description: Structure type that identifies the content of the message.  
Data type: MQLONG.  
Value: MQCFT\_APP\_ACTIVITY

### *StrucLength*

Description: Length in bytes of MQCFH structure.  
Data type: MQLONG.  
Value: MQCFH\_STRUC\_LENGTH

### *Version*

Description: Structure version number.  
Data type: MQLONG.  
Values: MQCFH\_VERSION\_3

### *Command*

Description: Command identifier. This field identifies the category of the message.  
Data type: MQLONG.  
Values: MQCMD\_ACTIVITY\_TRACE

### *MsgSeqNumber*

Description: Message sequence number. This field is the sequence number of the message within a group of related messages.  
Data type: MQLONG.  
Values: 1

### *Control*

Description: Control options.  
Data type: MQLONG.  
Values: MQCFC\_LAST.

#### *CompCode*

Description: Completion code.  
Data type: MQLONG.  
Values: MQCC\_OK.

#### *Reason*

Description: Reason code qualifying completion code.  
Data type: MQLONG.  
Values: MQRC\_NONE.

#### *ParameterCount*

Description: Count of parameter structures. This field is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.  
Data type: MQLONG.  
Values: 1 or greater

### **Application activity trace message data**

Immediately following the PCF header is a set of parameters describing the time interval for the activity trace. These parameters also indicate the sequence of messages in the event of messages being written. The order and number of fields following the header is not guaranteed, allowing additional information to be added in the future.

Message name: Activity trace message.  
System queue: SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE.

---

#### *QueueManager*

Description: The name of the queue manager  
Identifier: MQCA\_Q\_MGR\_NAME  
Data type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *QSGName*

Description: The name of QSG that the Queue Manager is a member of (z/OS only)  
Identifier: MQCA\_QSG\_NAME  
Data type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *HostName*

Description: The host name of the machine the Queue Manager is running on  
Identifier: MQCACF\_HOST\_NAME  
Data type: MQCFST

#### *IntervalStartDate*

Description: The date of the start of the monitoring period  
Identifier: MQCAMO\_START\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH

#### *IntervalStartTime*

Description: The time of the start of the monitoring period  
Identifier: MQCAMO\_START\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH

#### *IntervalEndDate*

Description: The date of the end of the monitoring period  
Identifier: MQCAMO\_END\_DATE  
Data type: MQCFST  
Maximum length: MQ\_DATE\_LENGTH

#### *IntervalEndTime*

Description: The time of the end of the monitoring period  
Identifier: MQCAMO\_END\_TIME  
Data type: MQCFST  
Maximum length: MQ\_TIME\_LENGTH

#### *CommandLevel*

Description: The IBM WebSphere MQ command level  
Identifier: MQIA\_COMMAND\_LEVEL  
Data type: MQCFIN

#### *SeqNumber*

Description: The sequence number normally zero. This value is incremented for each subsequent record for long running connections.  
Identifier: MQIACF\_SEQUENCE\_NUMBER  
Data type: MQCFIN

#### *ApplicationName*

Description: The name of the application. (program name)  
Identifier: MQCACF\_APPL\_NAME  
Data type: MQCFST  
Maximum length: MQ\_APPL\_NAME\_LENGTH

#### *ApplClass*

Description: Type of application that performed the activity. Possible values: MQAT\_\*  
Identifier: MQIA\_APPL\_TYPE  
Data type: MQCFIN

#### *ApplicationPid*

Description: The operating system Process ID of the application This is likely to be the TCB on z/OS – still TBD  
Identifier: MQIACF\_PROCESS\_ID  
Data type: MQCFIN

#### *UserId*

Description: The user identifier context of the application  
Identifier: MQCACF\_USER\_IDENTIFIER  
Data type: MQCFST  
Maximum length: MQ\_USER\_ID\_LENGTH

#### *APICallerType*

Description: The type of the application. Possible values: MQXACT\_EXTERNAL or MQXACT\_INTERNAL  
Identifier: MQIACF\_API\_CALLER\_TYPE  
Data type: MQCFIN

#### *Environment*

Description: The runtime environment of the application. Possible values: MQXE\_OTHER MQXE\_MCA MQXE\_MCA\_SVRCONN MQXE\_COMMAND\_SERVER MQXE\_MQSC  
Identifier: MQIACF\_API\_ENVIRONMENT  
Data type: MQCFIN

#### *Detail*

Description: The detail level that is recorded for the connection. Possible values: 1=LOW 2=MEDIUM 3=HIGH  
Identifier: MQIACF\_TRACE\_DETAIL  
Data type: MQCFIN

#### *TraceDataLength*

Description: The length of message data (in bytes) that is traced for this connection.  
Identifier: MQIACF\_TRACE\_DATA\_LENGTH  
Data type: MQCFIN

#### *Pointer Size*

Description: The length (in bytes) of pointers on the platform the application is running (to assist in interpretation of binary structures )  
Identifier: MQIACF\_POINTER\_SIZE  
Data type: MQCFIN

#### *Platform*

Description: The platform on which the queue manager is running. Value is one of the MQPL\_\* values.  
Identifier: MQIA\_PLATFORM  
Data type: MQCFIN

### **Variable parameters for application activity MQI operations**

The application activity data MQCFGR structure is followed by the set of PCF parameters which corresponds to the operation being performed . The parameters for each operation are defined in the following section.

The trace level indicates the level of trace granularity that is required for the parameters to be included in the trace. The possible trace level values are:

#### 1. Low

The parameter is included when “low”, “medium” or “high” activity tracing is configured for an application. This setting means that a parameter is always included in the AppActivityData group for the operation. This set of parameters is sufficient to trace the MQI calls an application makes, and to see if they are successful.

#### 2. Medium

The parameter is only included in the AppActivityData group for the operation when “medium” or “high” activity tracing is configured for an application. This set of parameters adds information about the resources, for example, queue and topic names used by the application.

#### 3. High

The parameter is only included in the AppActivityData group for the operation when “high” activity tracing is configured for an application. This set of parameters includes memory dumps of the structures passed to the MQI and XA functions. For this reason, it contains more information about the parameters used in MQI and XA calls. The structure memory dumps are shallow copies of the structures. To avoid erroneous attempts to dereference pointers, the pointer values in the structures are set to NULL.

**Note:** The version of the structure that is dumped is not necessarily identical to the version used by an application. The structure can be modified by an API crossing exit, by the activity trace code, or by the queue manager. A queue manager can modify a structure to a later version, but the queue manager never changes it to an earlier version of the structure. To do so, would risk losing data.

## **MQBACK:**

Application has started the MQBACK MQI function

### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

## **MQBEGIN:**

Application has started the MQBEGIN MQI function

### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

### *MQBO*

Description: The MQBEGIN options structure. This parameter is not included if a NULL pointer is used on the MQBEGIN call.  
PCF Parameter: MQBACF\_MQBO\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQBO structure.

## MQCALLBACK:

Application has started the MQCALLBACK function

### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

### *CallType*

Description: Why function has been called. One of the MQCBCT\_\* values  
PCF Parameter: MQIACF\_CALL\_TYPE  
Trace level: 1  
Type: MQCFIN

### *MsgBuffer*

Description: Message data.  
PCF Parameter: MQBACF\_MESSAGE\_DATA  
Trace level: 1  
Type: MQCFBS  
Length: Length is governed by the TRACEDATA() parameter set in the APPTTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.

### *MsgLength*

Description: Length of the message. (Taken from the DataLength field in the MQCBC structure).  
PCF Parameter: MQIACF\_MSG\_LENGTH  
Trace level: 1  
Type: MQCFIN

### *HighResTime*

Description: Time of operation in microseconds since midnight, January 1st 1970 (UTC)  
**Note:** The accuracy of this timer varies according to platform support for high a resolution timer  
PCF Parameter: MQIAMO64\_HIGHRES\_TIME  
Trace level: 2  
Type: MQCFIN64

### *ReportOptions*

Description: Options for report messages  
PCF Parameter: MQIACF\_REPORT  
Trace level: 2  
Type: MQCFIN

### *MsgType*

Description: Type of message  
PCF Parameter: MQIACF\_MSG\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *Expiry*

Description: Message lifetime  
PCF Parameter: MQIACF\_EXPIRY  
Trace level: 2  
Type: MQCFIN

#### *Format*

Description: Format name of message data  
PCF Parameter: MQCACH\_FORMAT\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_FORMAT\_LENGTH

#### *Priority*

Description: Message priority  
PCF Parameter: MQIACF\_PRIORITY  
Trace level: 2  
Type: MQCFIN

#### *Persistence*

Description: Message persistence  
PCF Parameter: MQIACF\_PERSISTENCE  
Trace level: 2  
Type: MQCFIN

#### *MsgId*

Description: Message identifier  
PCF Parameter: MQBACF\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *CorrelId*

Description: Correlation identifier  
PCF Parameter: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *ObjectName*

Description: The name of the opened object.  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQName*

Description: The local name of the queue from which the message was retrieved.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ReplyToQueue*

Description: MQ\_Q\_NAME\_LENGTH  
PCF Parameter: MQCACF\_REPLY\_TO\_Q  
Trace level: 2  
Type: MQCFST

#### *ReplyToQMgr*

Description: MQ\_Q\_MGR\_NAME\_LENGTH  
PCF Parameter: MQCACF\_REPLY\_TO\_Q\_MGR  
Trace level: 2  
Type: MQCFST

#### *CodedCharSetId*

Description: Character set identifier of message data  
PCF Parameter: MQIA\_CODED\_CHAR\_SET\_ID  
Trace level: 2  
Type: MQCFIN

#### *Encoding*

Description: Numeric encoding of message data.  
PCF Parameter: MQIACF\_ENCODING  
Trace level: 2  
Type: MQCFIN

#### *PutDate*

Description: MQ\_PUT\_DATE\_LENGTH  
PCF Parameter: MQCACF\_PUT\_DATE  
Trace level: 2  
Type: MQCFST

#### *PutTime*

Description: MQ\_PUT\_TIME\_LENGTH  
PCF Parameter: MQCACF\_PUT\_TIME  
Trace level: 2  
Type: MQCFST

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *PolicyName*

Description: The policy name that was applied to this message.  
**Note:** AMS protected messages only  
PCF Parameter: MQCA\_POLICY\_NAME  
Trace level: 2  
Type MQCFST  
Length: MQ\_OBJECT\_NAME\_LENGTH

#### *XmitqMsgId*

Description: The message ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_MSG\_ID  
Trace level: 2  
Type MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *XmitqCorrelId*

Description: The correlation ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_CORREL\_ID  
Trace level: 2  
Type MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *XmitqPutTime*

Description: The put time of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_TIME  
Trace level: 2  
Type MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *XmitqPutDate*

Description: The put date of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_DATE  
Trace level: 2  
Type MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *XmitqRemoteQName*

Description: The remote queue destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_Name  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *XmitqRemoteQMgr*

Description: The message ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_MSG\_ID\_LENGTH

#### *MsgDescStructure*

Description: The MQMD structure. This parameter is omitted if a version 4 MQGMO was used to request that a Message Handle be returned instead of an MQMD  
PCF Parameter: MQBACF\_MQMD\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQMD structure (actual size is dependent on structure version)

#### *GetMsgOptsStructure*

Description: The MQGMO structure.  
PCF Parameter: MQBACF\_MQGMO\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQGMO structure (actual size is dependent on structure version)

#### *MQCBCContextStructure*

Description: The MQCBC structure.  
PCF Parameter: MQBACF\_MQCBC\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQCBC structure (actual size is dependent on structure version)

### **MQCB:**

Application has started the manage callback MQI function

#### *CallbackOperation*

Description: The manage callback function operation. Set to one of the MQOP\_\* values  
PCF Parameter: MQIACF\_MQCB\_OPERATION  
Trace level: 1  
Type: MQCFIN

#### *CallbackType*

Description: The type of the callback function (CallbackType field from the MQCBD structure). Set to one of the MQCBT\_\* values  
PCF Parameter: MQIACF\_MQCB\_TYPE  
Trace level: 1  
Type: MQCFIN

#### *CallbackOptions*

Description: The callback options. Set to one of the MQCBDO\_\* values  
PCF Parameter: MQIACF\_MQCB\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *CallbackFunction*

Description: The pointer to the callback function if started as a function call.  
PCF Parameter: MQBACF\_MQCB\_FUNCTION  
Trace level: 1  
Type: MQCFBS  
Length: Size of MQPTR

#### *CallbackName*

Description: The name of the callback function if started as a dynamically linked program.  
PCF Parameter: MQCACF\_MQCB\_NAME  
Trace level: 1  
Type: MQCFST  
Length: Size of MQCHAR128

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *MaxMsgLength*

Description: Maximum message length. Set to an integer, or the special value MQCBD\_FULL\_MSG\_LENGTH  
PCF Parameter: MQIACH\_MAX\_MSG\_LENGTH  
Trace level: 2  
Type: MQCFIN

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *Callback DescriptorStructure*

Description: The MQCBD structure. This parameter is omitted if a NULL MQCBC value is passed to the MQCB call.

PCF Parameter: MQBACF\_MQCBD\_STRUCT

Trace level: 3

Type MQCFBS

Length: The length in bytes of the MQCBC structure

### *MsgDescStructure*

Description: The MQMD structure. The MsgDescStructure parameter is omitted if a NULL MQMD value is passed to the MQCB call.

PCF Parameter: MQBACF\_MQMD\_STRUCT

Trace level: 3

Type MQCFBS

Length: The length in bytes of the MQMD structure (actual size depends on structure version)

### *GetMsgOptsStructure*

Description: The MQGMO structure. This parameter is omitted if a NULL MQGMO value is passed to the MQCB call.

PCF Parameter: MQBACF\_MQGMO\_STRUCT

Trace level: 3

Type MQCFBS

Length: The length in bytes of the MQGMO structure (actual size depends on structure version)

## **MQCLOSE:**

Application has started the MQCLOSE MQI function

### *ObjectHandle*

Description: The object handle

PCF Parameter: MQIACF\_HOBJ

Trace level: 1

Type MQCFIN

### *CloseOptions*

Description: Close options

PCF Parameter: MQIACF\_CLOSE\_OPTIONS

Trace level: 1

Type MQCFIN

### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

### **MQCMIT:**

Application has started the MQCMIT MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

### **MQCONN and MQCONNX:**

Application has started the MQCONN or MQCONNX MQI function

#### *ConnectionId*

Description: The Connection ID if available or MQCONNID\_NONE if not  
PCF Parameter: MQBACF\_CONNECTION\_ID  
Trace level: 1  
Type: MQCFBS  
Maximum length: MQ\_CONNECTION\_ID\_LENGTH

#### *QueueManagerName*

Description: The (unresolved) name of the queue manager used in the MQCONN(X) call  
PCF Parameter: MQCA\_Q\_MGR\_NAME  
Trace level: 1  
Type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *ConnectOptions*

Description: Connect Options Derived from MQCNO\_\* values  
**Note:** MQCONNX only  
PCF Parameter: MQIACF\_CONNECT\_OPTIONS  
Trace level: 2  
Type: MQCFIN

#### *ConnectionOptionsStructure*

Description: The MQCNO structure.  
**Note:** MQCONNX only)  
PCF Parameter: MQBACF\_MQCNO\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Maximum length: The length in bytes of the MQCNO structure (actual size depends on structure version)

#### *ChannelDefinitionStructure*

Description: The MQCD structure.  
**Note:** Client connections only  
PCF Parameter: MQBACF\_MQCD\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Maximum length: The length in bytes of the MQCD structure (actual size depends on structure version)

#### **MQCTL:**

Application has started the MQCTL MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *CtlOperation*

Description: One of MQOP\_\* values  
PCF Parameter: MQIACF\_CTL\_OPERATION  
Trace level: 1  
Type: MQCFIN

#### **MQDISC:**

Application has started the MQDISC MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

### **MQGET:**

Application has started the MQGET MQI function

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *GetOptions*

Description: The get options from MQGMO.Options  
PCF Parameter: MQIACF\_GET\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *MsgBuffer*

Description: Message data. If TRACEDATA=NONE then this parameter is omitted  
PCF Parameter: MQBACF\_MESSAGE\_DATA  
Trace level: 1  
Type: MQCFBS  
Maximum length: Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration.  
(Included in the trace message as MQIACF\_TRACE\_DATA\_LENGTH).

#### *MsgLength*

Description: Length of the message.  
PCF Parameter: MQIACF\_MSG\_LENGTH  
Trace level: 1  
Type: MQCFIN

#### *HighResTime*

Description: Time of operation in microseconds since midnight, January 1 1970 (UTC)  
**Note:** The accuracy of this timer varies according to platform support for high a resolution timer  
PCF Parameter: MQIAMO64\_HIGHRES\_TIME  
Trace level: 2  
Type: MQCFIN64

#### *BufferLength*

Description: Length of the buffer provided by the application  
PCF Parameter: MQIACF\_BUFFER\_LENGTH  
Trace level: 2  
Type: MQCFIN

#### *ObjectName*

Description: The name of the opened object  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQName*

Description: The local name of the queue from which the message was retrieved.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ReportOptions*

Description: Message report options  
PCF Parameter: MQIACF\_REPORT  
Trace level: 2  
Type: MQCFIN

#### *MsgType*

Description: Type of message  
PCF Parameter: MQIACF\_MSG\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *Expiry*

Description: Message lifetime  
PCF Parameter: MQIACF\_EXPIRY  
Trace level: 2  
Type: MQCFIN

#### *Format*

Description: Format name of message data  
PCF Parameter: MQCACH\_FORMAT\_NAME  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_FORMAT\_LENGTH

#### *Priority*

Description: Message priority  
PCF Parameter: MQIACF\_PRIORITY  
Trace level: 2  
Type: MQCFIN

#### *Persistence*

Description: Message persistence  
PCF Parameter: MQIACF\_PERSISTENCE  
Trace level: 2  
Type: MQCFIN

#### *MsgId*

Description: Message identifier  
PCF Parameter: MQBACF\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Maximum length: MQ\_MSG\_ID\_LENGTH

#### *CorrelId*

Description: Correlation identifier  
PCF Parameter: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Maximum length: MQ\_CORREL\_ID\_LENGTH

#### *ReplyToQueue*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ReplyToQMgr*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *CodedCharSetId*

Description: Character set identifier of message data  
PCF Parameter: MQIA\_CODED\_CHAR\_SET\_ID  
Trace level: 2  
Type: MQCFIN

#### *Encoding*

Description: Numeric encoding of message data.  
PCF Parameter: MQIACF\_ENCODING  
Trace level: 2  
Type: MQCFIN

#### *PutDate*

Description:  
PCF Parameter: MQCACF\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_PUT\_DATE\_LENGTH

#### *PutTime*

Description:  
PCF Parameter: MQCACF\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_PUT\_TIME\_LENGTH

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *PolicyName*

Description: The policy name that was applied to this message.  
**Note:** AMS protected messages only  
PCF Parameter: MQCA\_POLICY\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_OBJECT\_NAME\_LENGTH

#### *XmitqMsgId*

Description: The message ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *XmitqCorrelId*

Description: The correlation ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *XmitqPutTime*

Description: The put time of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *XmitqPutDate*

Description: The put date of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *XmitqRemoteQName*

Description: The remote queue destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *XmitqRemoteQMgr*

Description: The remote queue manager destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *MsgDescStructure*

Description: The MQMD structure.  
PCF Parameter: MQBACF\_MQMD\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Maximum length: The length in bytes of the MQMD structure (actual size depends on structure version)

#### *GetMsgOptsStructure*

Description: The MQGMO structure.  
PCF Parameter: MQBACF\_MQGMO\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Maximum length: The length in bytes of the MQGMO structure (actual size depends on structure version)

### **MQINQ:**

Application has started the MQINQ MQI function

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *SelectorCount*

Description: The count of selectors that are supplied in the Selectors array.  
PCF Parameter: MQIACF\_SELECTOR\_COUNT  
Trace level: 2  
Type: MQCFIN

#### *Selectors*

Description: The list of attributes (integer or character) whose values must be returned by MQINQ.  
PCF Parameter: MQIACF\_SELECTORS  
Trace level: 2  
Type: MQCFIL

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Maximum length: Length varies

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *IntAttrCount*

Description: The number of integer attributes returned by the inquire operation  
PCF Parameter: MQIACF\_INTATTR\_COUNT  
Trace level: 3  
Type: MQCFIN

#### *IntAttrs*

Description: The integer attribute values returned by the inquire operation. This parameter is only present if IntAttrCount is > 0 when MQINQ returns.  
PCF Parameter: MQIACF\_INT\_ATTRS  
Trace level: 3  
Type: MQCFIL

#### *CharAttrs*

Description: The character attributes returned by the inquire operation. The values are concatenated together. This parameter is only included if CharAttrLength is > 0 when MQINQ returns.  
PCF Parameter: MQCACF\_CHAR\_ATTRS  
Trace level: 3  
Type: MQCFST

## **MQOPEN:**

Application has started the MQOPEN MQI function

### *ObjectType*

Description: The object type passed in MQOT.ObjectType  
PCF Parameter: MQIACF\_OBJECT\_TYPE  
Trace level: 1  
Type: MQCFIN

### *ObjectName*

Description: The name of the object passed to the MQI call before any queue name resolution is attempted.  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 1  
Type: MQCFST  
Maximum length: MQ\_Q\_NAME\_LENGTH

### *ObjectQMgrName*

Description: The name of the object queue manager passed to the MQI call before any queue name resolution is attempted.  
PCF Parameter: MQCACF\_OBJECT\_Q\_MGR\_NAME  
Trace level: 1  
Type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *OpenOptions*

Description: Options used to open the object  
PCF Parameter: MQIACF\_OPEN\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *AlternateUserId*

Description: Only included if MQOO\_ALTERNATE\_USER\_AUTHORITY is specified  
PCF Parameter: MQCACF\_ALTERNATE\_USERID  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_USER\_ID\_LENGTH

#### *RecsPresent*

Description: The number of object name records present. Only included if MQOD Version >= MQOD\_VERSION\_2  
PCF Parameter: MQIACF\_RECS\_PRESENT  
Trace level: 1  
Type: MQCFIN

#### *KnownDestCount*

Description: Number of local queues opened successfully Only included if MQOD Version >= MQOD\_VERSION\_2  
PCF Parameter: MQIACF\_KNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *UnknownDestCount*

Description: Number of remote queues opened successfully Only included if MQOD Version >= MQOD\_VERSION\_2  
PCF Parameter: MQIACF\_UNKNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *InvalidDestCount*

Description: Number of queues that failed to open Only included if MQOD Version >= MQOD\_VERSION\_2  
PCF Parameter: MQIACF\_INVALID\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *DynamicQName*

Description: The dynamic queue name passed as input to the MQOPEN call.  
PCF Parameter: MQCACF\_DYNAMIC\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedLocalQName*<sup>2324</sup>

Description: Contains the local queue name after name resolution has been carried out. (e.g. for remote queues this will be the name of the transmit queue)  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Range: If MQOD.Version is less than MQOD\_VERSION\_3 this contains the value of the MQOD.ObjectName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD\_VERSION\_3 this contains the value in the MQOD. ResolvedQName field.  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedLocalQMgrName*<sup>2324</sup>

Description: The local queue manager name after name resolution has been performed.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Range: Only if MQOD.Version >= MQOD\_VERSION\_3  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *ResolvedQName*<sup>2324</sup>

Description: The queue name after name resolution has been carried out.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Range: If MQOD.Version is less than MQOD\_VERSION\_3 this contains the value of the MQOD.ObjectName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD\_VERSION\_3 this contains the value in the MQOD. ResolvedQName field.  
Maximum length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQMgrName*<sup>2324</sup>

Description: Contains the queue manager name after name resolution has been carried out. If MQOD.Version is less than MQOD\_VERSION\_3 this contains the value of the MQOD.ObjectQMgrName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD\_VERSION\_3 this contains the value in the MQOD.ResolvedQMgrName field.

PCF Parameter: MQCACF\_RESOLVED\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Maximum length: MQ\_Q\_MGR\_NAME\_LENGTH

### *AlternateSecurityId*

Description: Alternative security identifier. Only present if MQOD.Version is equal or greater than MQOD\_VERSION\_3, MQOO\_ALTERNATE\_USER\_AUTHORITY is specified, and MQOD.AlternateSecurityId is not equal to MQSID\_NONE.

PCF Parameter: MQBACF\_ALTERNATE\_SECURITYID  
Trace level: 2  
Type: MQCFBS  
Maximum length: MQ\_SECURITY\_ID\_LENGTH

### *ObjectString*

Description: Long object name. Only included if MQOD.Version is equal or greater than MQOD\_VERSION\_4 and the VSLength field of MQOD.ObjectString is MQVS\_NULL\_TERMINATED or greater than zero.

PCF Parameter: MQCACF\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Maximum length: Length varies.

### *SelectionString*

Description: Selection string. Only included if MQOD.Version is equal or greater than MQOD\_VERSION\_4 and the VSLength field of MQOD.SelectionString is MQVS\_NULL\_TERMINATED or greater than zero.

PCF Parameter: MQCACF\_SELECTION\_STRING  
Trace level: 2  
Type: MQCFST  
Maximum length: Length varies.

### *ResObjectString*

Description: The long object name after the queue manager resolves the name provided in the ObjectName field. Only included for topics and queue aliases that reference a topic object if MQOD.Version is equal or greater than MQOD\_VERSION\_4 and VSLength is MQVS\_NULL\_TERMINATED or greater than zero.

PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Maximum length: Length varies.

### *ResolvedType*

Description: The type of the resolved (base) object being opened. Only included if MQOD.Version is equal or greater than MQOD\_VERSION\_4. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.

PCF Parameter: MQIACF\_RESOLVED\_TYPE

Trace level: 2

Type: MQCFIN

*Application Activity Distribution List PCF Group Header Structure:*

If the MQOPEN function opens a distribution list, then the MQOPEN parameters includes one AppActivityDistList PCF group for each of the queues in the distribution list up to the number of structures numbered in RecsPresent. The AppActivityDistList PCF group combines information from the MQOR, and MQRR structures to identify the queue name, and indicate the result of the open operation on the queue. An AppActivityDistList group always starts with the following MQCFGR structure:

| MQCFGR field   | Value                                   | Description                                                                                             |
|----------------|-----------------------------------------|---------------------------------------------------------------------------------------------------------|
| Type           | MQCFT_GROUP                             |                                                                                                         |
| StrucLength    | Length in bytes of the MQCFGR structure |                                                                                                         |
| Parameter      | MQGACF_APP_DIST_LIST                    | Distribution list group parameter                                                                       |
| ParameterCount | 4                                       | The number of parameter structures following the MQCFGR structure that are contained within this group. |

*ObjectName*

Description: The name of a queue in the distribution list MQ\_Q\_NAME\_LENGTH. Only included if MQOR structures are provided.

PCF Parameter: MQCACF\_OBJECT\_NAME

Trace level: 2

Type: MQCFST

Length: MQ\_Q\_NAME\_LENGTH. Only included if MQOR structures are provided.

*ObjectQMgrName*

Description: The name of the queue manager on which the queue named in ObjectName is defined.

PCF Parameter: MQCACF\_OBJECT\_Q\_MGR\_NAME

Trace level: 2

Type: MQCFST

Length: MQ\_Q\_MGR\_NAME\_LENGTH. Only included if MQOR structures are provided.

*CompCode*

23. This parameter is only included if the object being opened resolves to a queue, and the queue is opened for MQOO\_INPUT\_\*, MQOO\_OUTPUT, or MQOO\_BROWSE

24. The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

Description: The completion code indicating the result of the open for this object. Only included if MQRR structures are provided and the reason code for the MQOPEN is MQRC\_MULTIPLE\_REASONS  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 2  
Type: MQCFIN

#### *Reason*

Description: The reason code indicating the result of the open for this object. Only included if MQRR structures are provided and the reason code for the MQOPEN is MQRC\_MULTIPLE\_REASONS  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 2  
Type: MQCFIN

### **MQPUT:**

Application has started the MQPUT MQI function.

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *PutOptions*

Description: The put options from MQPMO.Options  
PCF Parameter: MQIACF\_PUT\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *MsgBuffer*

Description: Message data.  
PCF Parameter: MQBACF\_MESSAGE\_DATA  
Trace level: 1  
Type: MQCFBS  
Length: Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.

### *MsgLength*

Description: Length of the message.  
PCF Parameter: MQIACF\_MSG\_LENGTH  
Trace level: 1  
Type: MQCFIN

### *RecsPresent*

Description: The number of put message records or response records present. Only included if MQPMO Version >= MQPMO\_VERSION\_2  
PCF Parameter: MQIACF\_RECS\_PRESENT  
Trace level: 1  
Type: MQCFIN

### *KnownDestCount*

Description: Number of messages sent successfully to local queues  
PCF Parameter: MQIACF\_KNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

### *UnknownDestCount*

Description: Number of messages sent successfully to remote queues  
PCF Parameter: MQIACF\_UNKNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

### *InvalidDestCount*

Description: Number of messages that could not be sent  
PCF Parameter: MQIACF\_INVALID\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

### *HighResTime*

Description: Time of operation in microseconds since midnight, January 1st 1970 (UTC)  
**Note:** The accuracy of this timer varies according to platform support for high a resolution timer.  
PCF Parameter: MQIAMO64\_HIGHRES\_TIME  
Trace level: 2  
Type: MQCFIN64

#### *ObjectName*

Description: The name of the opened object.  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQName*

Description: The name of the queue after queue name resolution has been performed.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQMgrName*

Description: The queue manager name after name resolution has been performed.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *ResolvedLocalQName*<sup>25</sup>

Description: Contains the local queue name after name resolution has been carried out.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST

#### *ResolvedLocalQMgrName*<sup>25</sup>

Description: Contains the local queue manager name after name resolution has been carried out.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *ReportOptions*

Description: Message report options  
PCF Parameter: MQIACF\_REPORT  
Trace level: 2  
Type: MQCFIN

#### *MsgType*

Description: Type of message  
PCF Parameter: MQIACF\_MSG\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *Expiry*

Description: Message lifetime  
PCF Parameter: MQIACF\_EXPIRY  
Trace level: 2  
Type: MQCFIN

#### *Format*

Description: Format name of message data  
PCF Parameter: MQCACH\_FORMAT\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_FORMAT\_LENGTH

#### *Priority*

Description: Message priority  
PCF Parameter: MQIACF\_PRIORITY  
Trace level: 2  
Type: MQCFIN

#### *Persistence*

Description: Message persistence  
PCF Parameter: MQIACF\_PERSISTENCE  
Trace level: 2  
Type: MQCFIN

#### *MsgId*

Description: Message identifier  
PCF Parameter: MQBACF\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *CorrelId*

Description: Correlation identifier  
PCF Parameter: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *ReplyToQueue*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ReplyToQMgr*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *CodedCharSetId*

Description: Character set identifier of message data  
PCF Parameter: MQIA\_CODED\_CHAR\_SET\_ID  
Trace level: 2  
Type: MQCFIN

#### *Encoding*

Description: Numeric encoding of message data.  
PCF Parameter: MQIACF\_ENCODING  
Trace level: 2  
Type: MQCFIN

#### *PutDate*

Description:  
PCF Parameter: MQCACF\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *PutTime*

Description:  
PCF Parameter: MQCACF\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *PolicyName*

Description: The policy name that was applied to this message.  
**Note:** AMS protected messages only  
PCF Parameter: MQCA\_POLICY\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_OBJECT\_NAME\_LENGTH

#### *XmitqMsgId*

Description: The message ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *XmitqCorrelId*

Description: The correlation ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *XmitqPutTime*

Description: The put time of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *XmitqPutDate*

Description: The put date of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *XmitqRemoteQName*

Description: The remote queue destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *XmitqRemoteQMgr*

Description: The remote queue manager destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *PutMsgOptsStructure*

|                |                                                                                       |
|----------------|---------------------------------------------------------------------------------------|
| Description:   | The MQPMO structure.                                                                  |
| PCF Parameter: | MQBACF_MQPMO_STRUCT                                                                   |
| Trace level:   | 3                                                                                     |
| Type:          | MQCFBS                                                                                |
| Length:        | The length in bytes of the MQPMO structure (actual size depends on structure version) |

*MQPUT Application Activity Distribution List PCF Group Header Structure:*

If the MQPUT function is putting to a distribution list, then the MQPUT parameters include one AppActivityDistList PCF group. For each of the queues in the distribution list, see “Application Activity Distribution List PCF Group Header Structure” on page 1068. The AppActivityDistList PCF group combines information from the MQPMR, and MQRR structures to identify the PUT parameters, and indicate the result of the PUT operation on each queue. For MQPUT operations the AppActivityDistList group contains some or all of the following parameters (the CompCode and Reason is present if the reason code is MQRC\_MULTIPLE\_REASONS and the other parameters are determined by the MQPMO.PutMsgRecFields field):

*CompCode*

|                |                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:   | The completion code indicating the result of the operation. Only included if MQRR structures are provided and the reason code for the MQPUT is MQRC_MULTIPLE_REASONS |
| PCF Parameter: | MQIACF_COMP_CODE                                                                                                                                                     |
| Trace level:   | 2                                                                                                                                                                    |
| Type:          | MQCFIN                                                                                                                                                               |

*Reason*

|                |                                                                                                                                                                            |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:   | The reason code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT is MQRC_MULTIPLE_REASONS |
| PCF Parameter: | MQIACF_REASON_CODE                                                                                                                                                         |
| Trace level:   | 2                                                                                                                                                                          |
| Type:          | MQCFIN                                                                                                                                                                     |

*MsgId*

|                |                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------|
| Description:   | Message identifier. Only included if MQPMR structures are provided and PutMsgRecFields includes MQPMRF_MSG_ID |
| PCF Parameter: | MQBACF_MSG_ID                                                                                                 |
| Trace level:   | 2                                                                                                             |
| Type:          | MQCFBS                                                                                                        |
| Length:        | MQ_MSG_ID_LENGTH                                                                                              |

*CorrelId*

---

25. The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

Description: Correlation identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_CORREL\_ID  
PCF Parameter: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *GroupId*

Description: Group identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_GROUP\_ID  
PCF Parameter: MQBACF\_GROUP\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_GROUP\_ID\_LENGTH

#### *Feedback*

Description: Feedback. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_FEEDBACK  
PCF Parameter: MQIACF\_FEEDBACK  
Trace level: 2  
Type: MQCFIN

#### *AccountingToken*

Description: AccountingToken. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_ACCOUNTING\_TOKEN  
PCF Parameter: MQBACF\_ACCOUNTING\_TOKEN  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_ACCOUNTING\_TOKEN\_LENGTH.

### **MQPUT1:**

Application has started the MQPUT1 MQI function

#### *ObjectType*

Description: The object type passed in MQOT.ObjectType  
PCF Parameter: MQIACF\_OBJECT\_TYPE  
Trace level: 1  
Type: MQCFIN

#### *ObjectName*

Description: The name of the object passed to the MQI call before any queue name resolution is attempted.  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 1  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ObjectQMgrName*

Description: The name of the object queue manager passed to the MQI call before any queue name resolution is attempted.  
PCF Parameter: MQCACF\_OBJECT\_Q\_MGR\_NAME  
Trace level: 1  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *PutOptions*

Description: The put options from MQPMO.Options  
PCF Parameter: MQIACF\_PUT\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *AlternateUserId*

Description: Only included if MQPMO.ALTERNATE\_USER\_AUTHORITY is specified.  
PCF Parameter: MQCACF\_ALTERNATE\_USERID  
Trace level: 2  
Type: MQCFST  
Length: MQ\_USER\_ID\_LENGTH

#### *RecsPresent*

Description: The number of object name records present  
PCF Parameter: MQIACF\_RECS\_PRESENT  
Trace level: 1  
Type: MQCFIN

#### *KnownDestCount*

Description: Number of local queues opened successfully  
PCF Parameter: MQIACF\_KNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *UnknownDestCount*

Description: Number of remote queues opened successfully  
PCF Parameter: MQIACF\_UNKNOWN\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *InvalidDestCount*

Description: Number of queues that failed to open  
PCF Parameter: MQIACF\_INVALID\_DEST\_COUNT  
Trace level: 1  
Type: MQCFIN

#### *MsgBuffer*

Description: Message data.  
PCF Parameter: MQBACF\_MESSAGE\_DATA  
Trace level: 1  
Type: MQCFBS  
Length: Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.

#### *MsgLength*

Description: Length of the message.  
PCF Parameter: MQIACF\_MSG\_LENGTH  
Trace level: 1  
Type: MQCFIN

#### *HighResTime*

Description: Time of operation in microseconds since midnight, January 1st 1970 (UTC)  
**Note:** The accuracy of this timer will vary according to platform support for high a resolution timer.  
PCF Parameter: MQIAMO64\_HIGHRES\_TIME  
Trace level: 2  
Type: MQCFIN64

#### *ResolvedQName*

Description: The name of the queue after queue name resolution has been performed.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ResolvedQMgrName*

Description: The queue manager name after name resolution has been performed.  
PCF Parameter: MQCACF\_RESOLVED\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *ResolvedLocalQName<sup>26</sup>*

Description: Contains the local queue name after name resolution has been carried out  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST

#### *ResolvedLocalQMgrName<sup>26</sup>*

Description: Contains the local queue manager name after name resolution has been carried out.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_MGR\_NAME\_LENGTH

#### *AlternateSecurityId*

Description: Alternate security identifier. Only present if MQOD.Version is equal or greater than MQOD.VERSION\_3 and MQOD.AlternateSecurityId is not equal to MQSID\_NONE.  
PCF Parameter: MQBACF\_ALTERNATE\_SECURITYID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_SECURITY\_ID\_LENGTH

#### *ObjectString*

Description: Long object name. Only included if MQOD.Version is equal or greater than MQOD.VERSION\_4 and the VSLength field of MQOD.ObjectString is MQVS\_NULL\_TERMINATED or greater than zero.  
PCF Parameter: MQCACF\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResObjectString*

Description: The long object name after the queue manager resolves the name provided in the ObjectName field. Only included for topics and queue aliases that reference a topic object if MQOD.Version is equal or greater than MQOD\_VERSION\_4 and VSLength is MQVS\_NULL\_TERMINATED or greater than zero.

PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING

Trace level: 2

Type: MQCFST

Length: Length varies.

### *ResolvedType*

Description: The type of the resolved (base) object being opened. Only included if MQOD.Version is equal or greater than MQOD\_VERSION\_4. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.

PCF Parameter: MQIACF\_RESOLVED\_TYPE

Trace level: 2

Type: MQCFIN

### *ReportOptions*

Description: Message report options

PCF Parameter: MQIACF\_REPORT

Trace level: 2

Type: MQCFIN

### *MsgType*

Description: Type of message

PCF Parameter: MQIACF\_MSG\_TYPE

Trace level: 2

Type: MQCFIN

### *Expiry*

Description: Message lifetime

PCF Parameter: MQIACF\_EXPIRY

Trace level: 2

Type: MQCFIN

### *Format*

Description: Format name of message data

PCF Parameter: MQCACH\_FORMAT\_NAME

Trace level: 2

Type: MQCFST

Length: MQ\_FORMAT\_LENGTH

### *Priority*

Description: Message priority  
PCF Parameter: MQIACF\_PRIORITY  
Trace level: 2  
Type: MQCFIN

#### *Persistence*

Description: Message persistence  
PCF Parameter: MQIACF\_PERSISTENCE  
Trace level: 2  
Type: MQCFIN

#### *MsgId*

Description: Message identifier  
PCF Parameter: MQBACF\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *CorrelId*

PCF Parameter: Correlation identifier  
Description: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *ReplyToQueue*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ReplyToQMgr*

Description:  
PCF Parameter: MQCACF\_REPLY\_TO\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQCFST

#### *CodedCharSetId*

Description: Character set identifier of message data  
PCF Parameter: MQIA\_CODED\_CHAR\_SET\_ID  
Trace level: 2  
Type: MQCFIN

#### *Encoding*

Description: Numeric encoding of message data.  
PCF Parameter: MQIACF\_ENCODING  
Trace level: 2  
Type: MQCFIN

#### *PutDate*

Description:  
PCF Parameter: MQCACF\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *PutTime*

Description:  
PCF Parameter: MQCACF\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *PolicyName*

Description: The policy name that was applied to this message.  
**Note:** AMS protected messages only  
PCF Parameter: MQCA\_POLICY\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_OBJECT\_NAME\_LENGTH

#### *XmitqMsgId*

Description: The message ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *XmitqCorrelId*

Description: The correlation ID of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQBACF\_XQH\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *XmitqPutTime*

Description: The put time of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_TIME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_TIME\_LENGTH

#### *XmitqPutDate*

Description: The put date of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_PUT\_DATE  
Trace level: 2  
Type: MQCFST  
Length: MQ\_PUT\_DATE\_LENGTH

#### *XmitqRemoteQName*

Description: The remote queue destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *XmitqRemoteQMGr*

Description: The remote queue manager destination of the message in the transmission queue header.  
**Note:** Only when Format is MQFMT\_XMIT\_Q\_HEADER  
PCF Parameter: MQCACF\_XQH\_REMOTE\_Q\_MGR  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *PutMsgOptsStructure*

|                |                                                                                       |
|----------------|---------------------------------------------------------------------------------------|
| Description:   | The MQPMO structure.                                                                  |
| PCF Parameter: | MQBACF_MQPMO_STRUCT                                                                   |
| Trace level:   | 3                                                                                     |
| Type:          | MQCFBS                                                                                |
| Length:        | The length in bytes of the MQPMO structure (actual size depends on structure version) |

*MQPUT1 AppActivityDistList PCF Group Header Structure:*

If the MQPUT1 function is putting to a distribution list, then the variable parameters include one AppActivityDistList PCF group. For each of the queues in the distribution list, see “Application Activity Distribution List PCF Group Header Structure” on page 1068. The AppActivityDistList PCF group combines information from the MQOR, MQPMR, and MQRR structures to identify the objects, and the PUT parameters, and indicate the result of the PUT operation on each queue. For MQPUT1 operations the AppActivityDistList group contains some or all of the following parameters (the CompCode, Reason, ObjectName, and ObjectQMgrName is present if the reason code is MQRC\_MULTIPLE\_REASONS and the other parameters is determined by the MQPMO.PutMsgRecFields field):

*CompCode*

|                |                                                                                                                                                                                 |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:   | The completion code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT1 is MQRC_MULTIPLE_REASONS |
| PCF Parameter: | MQIACF_COMP_CODE                                                                                                                                                                |
| Trace level:   | 2                                                                                                                                                                               |
| Type:          | MQCFIN                                                                                                                                                                          |

*Reason*

|                |                                                                                                                                                                             |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description:   | The reason code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT1 is MQRC_MULTIPLE_REASONS |
| PCF Parameter: | MQIACF_REASON_CODE                                                                                                                                                          |
| Trace level:   | 2                                                                                                                                                                           |
| Type:          | MQCFIN                                                                                                                                                                      |

*ObjectName*

|                |                                                                                              |
|----------------|----------------------------------------------------------------------------------------------|
| Description:   | The name of a queue in the distribution list. Only included if MQOR structures are provided. |
| PCF Parameter: | MQCACF_OBJECT_NAME                                                                           |
| Trace level:   | 2                                                                                            |
| Type:          | MQCFST                                                                                       |
| Length:        | MQ_Q_NAME_LENGTH                                                                             |

*MsgId*

---

26. The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

Description: Message identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_MSG\_ID  
PCF Parameter: MQBACF\_MSG\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_MSG\_ID\_LENGTH

#### *CorrelId*

Description: Correlation identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_CORREL\_ID  
PCF Parameter: MQBACF\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *GroupId*

Description: Group identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_GROUP\_ID  
PCF Parameter: MQBACF\_GROUP\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_GROUP\_ID\_LENGTH

#### *Feedback*

Description: Feedback. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_FEEDBACK  
PCF Parameter: MQIACF\_FEEDBACK  
Trace level: 2  
Type: MQCFIN

#### *AccountingToken*

Description: AccountingToken. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF\_ACCOUNTING\_TOKEN  
PCF Parameter: MQBACF\_ACCOUNTING\_TOKEN  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_ACCOUNTING\_TOKEN\_LENGTH.

#### **MQSET:**

Application has started the MQSET MQI function

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *SelectorCount*

Description: The count of selectors that are supplied in the Selectors array.  
PCF Parameter: MQIACF\_SELECTOR\_COUNT  
Trace level: 2  
Type: MQCFIN

#### *Selectors*

Description: The list of attributes (integer or character) whose values are being updated by MQSET.  
PCF Parameter: MQIACF\_SELECTORS  
Trace level: 2  
Type: MQCFIL

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *IntAttrCount*

Description: The number of integer attributes to be updated by the set operation.  
PCF Parameter: MQIACF\_INTATTR\_COUNT  
Trace level: 3  
Type: MQCFIN

#### *IntAttrs*

Description: The integer attribute values  
PCF Parameter: MQIACF\_INT\_ATTRS  
Trace level: 3  
Type: MQCFIL  
Range: This parameter is only present if IntAttrCount is > 0

#### *CharAttrs*

Description: The character attributes to be updated by the set operation. The values are concatenated together.  
PCF Parameter: MQCACF\_CHAR\_ATTRS  
Trace level: 3  
Type: MQCFST  
Range: This parameter is only included if CharAttrLength is > 0

### **MQSUB:**

Application has started the MQSUB MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *SubHandle*

Description: The subscription handle  
PCF Parameter: MQIACF\_HSUB  
Trace level: 1  
Type: MQCFIN

#### *ObjectHandle*

Description: The object handle  
PCF Parameter: MQIACF\_HOBJ  
Trace level: 1  
Type: MQCFIN

#### *Options*

Description: Subscription options  
PCF Parameter: MQIACF\_SUB\_OPTIONS  
Trace level: 1  
Type: MQCFIN

#### *ObjectName*

Description: The name of the object.  
PCF Parameter: MQCACF\_OBJECT\_NAME  
Trace level: 1  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH

#### *ObjectString*

Description: Long object name.  
PCF Parameter: MQCACF\_OBJECT\_STRING  
Trace level: 1  
Type: MQCFST  
Range: Only included if the VSLength field of MQSD.ObjectString is greater than zero or MQVS\_NULL\_TERMINATED.  
Length: Length varies.

#### *AlternateUserId*

Description:  
PCF Parameter: MQCACF\_ALTERNATE\_USERID  
Trace level: 2  
Type: MQCFST  
Range: Only included if MQSO\_ALTERNATE\_USER\_AUTHORITY is specified.  
Length: MQ\_USER\_ID\_LENGTH

#### *AlternateSecurityId*

Description: Alternate security identifier.  
PCF Parameter: MQBACF\_ALTERNATE\_SECURITYID  
Trace level: 2  
Type: MQCFBS  
Range: Only present if MQSO\_ALTERNATE\_USER\_AUTHORITY is specified and MQSD.AlternateSecurityId is not equal to MQSID\_NONE.  
Length: MQ\_SECURITY\_ID\_LENGTH

#### *SubName*

Description: Subscription Name  
PCF Parameter: MQCACF\_SUB\_NAME  
Trace level: 2  
Type: MQCFST  
Range: Only included if the VSLength field of MQSD.SubName is greater than zero or MQVS\_NULL\_TERMINATED.  
Length: Length varies.

#### *SubUserData*

Description: Subscription User Data  
PCF Parameter: MQCACF\_SUB\_USER\_DATA  
Trace level: 2  
Type: MQCFST  
Range: Only included if the VSLength field of MQSD.SubName is greater than zero or MQVS\_NULL\_TERMINATED.  
Length: Length varies.

#### *SubCorrelId*

Description: Subscription Correlation identifier  
PCF Parameter: MQBACF\_SUB\_CORREL\_ID  
Trace level: 2  
Type: MQCFBS  
Length: MQ\_CORREL\_ID\_LENGTH

#### *SelectionString*

Description: Selection string.  
PCF Parameter: MQCACF\_SELECTION\_STRING  
Trace level: 2  
Type: MQCFST  
Range: Only included if the VSLength field of MQSD.SelectionString is MQVS\_NULL\_TERMINATED or greater than zero.  
Length: Length varies.

#### *ResolvedQName*

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT\_Q.  
PCF Parameter: MQCACF\_RESOLVED\_LOCAL\_Q\_NAME  
Trace level: 2  
Type: MQCFST  
Length: MQ\_Q\_NAME\_LENGTH.

#### *ResObjectString*

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT\_TOPIC.  
PCF Parameter: MQCACF\_RESOLVED\_OBJECT\_STRING  
Trace level: 2  
Type: MQCFST  
Length: Length varies.

#### *ResolvedType*

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT\_Q, MQOT\_TOPIC, or MQOT\_NONE.  
PCF Parameter: MQIACF\_RESOLVED\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *SubDescriptorStructure*

Description: The MQSD structure.  
PCF Parameter: MQBACF\_MQSD\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQSD structure.

### **MQSUBRQ:**

Application has started the MQSUBRQ MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *SubHandle*

Description: The subscription handle  
PCF Parameter: MQIACF\_HSUB  
Trace level: 1  
Type: MQCFIN

#### *SubOptions*

Description: The sub options from MQSB.Options  
PCF Parameter: MQIACF\_SUBRQ\_OPTIONS  
Trace level: 2  
Type: MQCFIN

#### *Action*

Description: The subscription request action (MQSR\_\*)  
PCF Parameter: MQIACF\_SUBRQ\_ACTION  
Trace level: 2  
Type: MQCFIN

#### *NumPubs*

Description: The number of publications sent as a result of this call (from MQSB.NumPubs)  
PCF Parameter: MQIACF\_NUM\_PUBS  
Trace level: 2  
Type: MQCFIN

### **MQSTAT:**

Application has started the MQSTAT MQI function

#### *CompCode*

Description: The completion code indicating the result of the operation  
PCF Parameter: MQIACF\_COMP\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Reason*

Description: The reason code result of the operation  
PCF Parameter: MQIACF\_REASON\_CODE  
Trace level: 1  
Type: MQCFIN

#### *Type*

Description: Type of status information being requested  
PCF Parameter: MQIACF\_STATUS\_TYPE  
Trace level: 2  
Type: MQCFIN

#### *StatusStructure*

Description: The MQSTS structure.  
PCF Parameter: MQBACF\_MQSTS\_STRUCT  
Trace level: 3  
Type: MQCFBS  
Length: The length in bytes of the MQSTS structure (actual size depends on structure version)

## **Variable Parameters for Application Activity XA Operations**

XA operations are API calls that applications can make to enable MQ to participate in a transaction. The parameters for each operation are defined in the following section.

The trace level indicates the level of trace granularity that is required for the parameters to be included in the trace. The possible trace level values are:

### 1. Low

The parameter is included when “low”, “medium” or “high” activity tracing is configured for an application. This setting means that a parameter is always included in the AppActivityData group for the operation. This set of parameters is sufficient to trace the MQI calls an application makes, and to see if they are successful.

### 2. Medium

The parameter is only included in the AppActivityData group for the operation when “medium” or “high” activity tracing is configured for an application. This set of parameters adds information about the resources, for example, queue and topic names used by the application.

### 3. High

The parameter is only included in the AppActivityData group for the operation when “high” activity tracing is configured for an application. This set of parameters includes memory dumps of the structures passed to the MQI and XA functions. For this reason, it contains more information about the parameters used in MQI and XA calls. The structure memory dumps are shallow copies of the structures. To avoid erroneous attempts to dereference pointers, the pointer values in the structures are set to NULL.

**Note:** The version of the structure that is dumped is not necessarily identical to the version used by an application. The structure can be modified by an API crossing exit, by the activity trace code, or by the queue manager. A queue manager can modify a structure to a later version, but the queue manager never changes it to an earlier version of the structure. To do so, would risk losing data.

## **AXREG:**

Application has started the AXREG AX function

*XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

#### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

### **AXUNREG:**

Application has started the AXUNREG AX function

#### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XACLOSE:**

Application has started the XACLOSE AX function

### *Xa\_info*

Description: Information used to initialize the resource manager.  
PCF Parameter: MQCACF\_XA\_INFO  
Trace level: 1  
Type: MQCFST

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XACOMMIT:**

Application has started the XACOMMIT AX function

### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

### **XACOMPLETE:**

Application has started the XACOMPLETE AX function

#### *Handle*

Description: Handle to async operation  
PCF Parameter: MQIACF\_XA\_HANDLE  
Trace level: 1  
Type: MQCFIN

#### *Retval*

Description: Return value of the asynchronous function  
PCF Parameter: MQIACF\_XA\_RETVAL  
Trace level: 1  
Type: MQCFINMQCFBS

#### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

#### **XAEND:**

Application has started the XAEND AX function

#### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

#### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XAFORGET:**

Application has started the AXREG AX function

### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XAOPEN:**

Application has started the XAOPEN AX function

### *xa\_info*

Description: Information used to initialize the resource manager.  
PCF Parameter: MQCACF\_XA\_INFO  
Trace level: 1  
Type: MQCFST

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

### **XAPREPARE:**

Application has started the XAPREPARE AX function

#### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

#### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

#### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

#### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## XARECOVER:

Application has started the XARECOVER AX function

### *Count*

Description: Count of XIDs  
PCF Parameter: MQIACF\_XA\_COUNT  
Trace level: 1  
Type: MQCFIN

### *XIDs*

Description: The XID structures  
**Note:** There are multiple instances of this PCF parameter – one for every XID structure up to Count XIDs  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XAROLLBACK:**

Application has started the XAROLLBACK AX function

### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

## **XASTART:**

Application has started the XASTART AX function

### *XID*

Description: The XID structure  
PCF Parameter: MQBACF\_XA\_XID  
Trace level: 1  
Type: MQCFBS  
Length: Sizeof(XID)

### *Rmid*

Description: Resource manager identifier  
PCF Parameter: MQIACF\_XA\_RMID  
Trace level: 1  
Type: MQCFIN

### *Flags*

Description: Flags  
PCF Parameter: MQIACF\_XA\_FLAGS  
Trace level: 1  
Type: MQCFIN

### *XARetCode*

Description: Return code  
PCF Parameter: MQIACF\_XA\_RETCODE  
Trace level: 1  
Type: MQCFIN

---

## Real-time monitoring

Real-time monitoring is a technique that allows you to determine the current state of queues and channels within a queue manager. The information returned is accurate at the moment the command was issued.

A number of commands are available that when issued return real-time information about queues and channels. Information can be returned for one or more queues or channels and can vary in quantity. Real-time monitoring can be used in the following tasks:

- Helping system administrators understand the steady state of their IBM WebSphere MQ system. This helps with problem diagnosis if a problem occurs in the system.
- Determining the condition of your queue manager at any moment, even if no specific event or problem has been detected.
- Assisting with determining the cause of a problem in your system.

With real-time monitoring, information can be returned for either queues or channels. The amount of real-time information returned is controlled by queue manager, queue, and channel attributes.

- You monitor a queue by issuing commands to ensure that the queue is being serviced properly. Before you can use some of the queue attributes, you must enable them for real-time monitoring.
- You monitor a channel by issuing commands to ensure that the channel is running properly. Before you can use some of the channel attributes, you must enable them for real-time monitoring.

Real-time monitoring for queues and channels is in addition to, and separate from, performance and channel event monitoring.

**Related concepts:**

“Attributes that control real-time monitoring”

“The Windows performance monitor” on page 1114

“Attributes that control real-time monitoring”

“The Windows performance monitor” on page 1114

**Related tasks:**

“Displaying queue and channel monitoring data” on page 1104

“Monitoring queues” on page 1106

“Monitoring channels” on page 1108

“Displaying queue and channel monitoring data” on page 1104

“Monitoring queues” on page 1106

“Monitoring channels” on page 1108

## Attributes that control real-time monitoring

Some queue and channel status attributes hold monitoring information, if real-time monitoring is enabled. If real-time monitoring is not enabled, no monitoring information is held in these monitoring attributes. Examples demonstrate how you can use these queue and channel status attributes.

You can enable or disable real-time monitoring for individual queues or channels, or for multiple queues or channels. To control individual queues or channels, set the queue attribute `MONQ` or the channel attribute `MONCHL`, to enable or disable real-time monitoring. To control many queues or channels together, enable or disable real-time monitoring at the queue manager level by using the queue manager attributes `MONQ` and `MONCHL`. For all queue and channel objects with a monitoring attribute that is specified with the default value, `QMGR`, real-time monitoring is controlled at the queue manager level.

Automatically defined cluster-sender channels are not WebSphere MQ objects, so do not have attributes in the same way as channel objects. To control automatically defined cluster-sender channels, use the queue manager attribute, `MONACLS`. This attribute determines whether automatically defined cluster-sender channels within a queue manager are enabled or disabled for channel monitoring.

For real-time monitoring of channels, you can set the `MONCHL` attribute to one of the three monitoring levels: low, medium, or high. You can set the monitoring level either at the object level or at the queue manager level. The choice of level is dependent on your system. Collecting monitoring data might require some instructions that are relatively expensive computationally, such as obtaining system time. To reduce the effect of real-time monitoring, the medium and low monitoring options measure a sample of the data at regular intervals rather than collecting data all the time. Table 113 summarizes the monitoring levels available for real-time monitoring of channels:

*Table 113. Monitoring levels*

| Level  | Description                                               | Usage                                                                                                        |
|--------|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| Low    | Measure a small sample of the data, at regular intervals. | For objects that process a high volume of messages.                                                          |
| Medium | Measure a sample of the data, at regular intervals.       | For most objects.                                                                                            |
| High   | Measure all data, at regular intervals.                   | For objects that process only a few messages per second, on which the most current information is important. |

For real-time monitoring of queues, you can set the `MONQ` attribute to one of the three monitoring levels, low, medium or high. However, there is no distinction between these values. The values all enable data collection, but do not affect the size of the sample.

## Examples

The following examples demonstrate how to set the necessary queue, channel, and queue manager attributes to control the level of monitoring. For all of the examples, when monitoring is enabled, queue and channel objects have a medium level of monitoring.

1. To enable both queue and channel monitoring for all queues and channels at the queue manager level, use the following commands:  

```
ALTER QMGR MONQ(MEDIUM) MONCHL(MEDIUM)
ALTER QL(Q1) MONQ(QMGR)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(QMGR)
```
2. To enable monitoring for all queues and channels, with the exception of local queue, Q1, and sender channel, QM1.TO.QM2, use the following commands:  

```
ALTER QMGR MONQ(MEDIUM) MONCHL(MEDIUM)
ALTER QL(Q1) MONQ(OFF)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(OFF)
```
3. To disable both queue and channel monitoring for all queues and channels, with the exception of local queue, Q1, and sender channel, QM1.TO.QM2, use the following commands:  

```
ALTER QMGR MONQ(OFF) MONCHL(OFF)
ALTER QL(Q1) MONQ(MEDIUM)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(MEDIUM)
```
4. To disable both queue and channel monitoring for all queues and channels, regardless of individual object attributes, use the following command:  

```
ALTER QMGR MONQ(NONE) MONCHL(NONE)
```
5. To control the monitoring capabilities of automatically defined cluster-sender channels use the following command:  

```
ALTER QMGR MONACLS(MEDIUM)
```
6. To specify that automatically defined cluster-sender channels are to use the queue manager setting for channel monitoring, use the following command:  

```
ALTER QMGR MONACLS(QMGR)
```

### Related concepts:

“Real-time monitoring” on page 1102

“Using IBM WebSphere MQ online monitoring” on page 1118

“Working with queue managers” on page 85

### Related tasks:

“Displaying queue and channel monitoring data”

### Related reference:

 Monitoring (MONCHL) (*WebSphere MQ V7.1 Reference*)

## Displaying queue and channel monitoring data

To display real-time monitoring information for a queue or channel, use either the WebSphere MQ Explorer or the appropriate MQSC command. Some monitoring fields display a comma-separated pair of indicator values, which help you to monitor the operation of your queue manager. Examples demonstrate how you can display monitoring data.

### About this task

Monitoring fields that display a pair of values separated by a comma provide short term and long term indicators for the time measured since monitoring was enabled for the object, or from when the queue manager was started:

- The short term indicator is the first value in the pair and is calculated in a way such that more recent measurements are given a higher weighting and will have a greater effect on this value. This gives an indication of recent trend in measurements taken.
- The long term indicator is the second value in the pair and is calculated in a way such that more recent measurements are not given such a high weighting. This gives an indication of the longer term activity on performance of a resource.

These indicator values are most useful to detect changes in the operation of your queue manager. This requires knowledge of the times these indicators show when in normal use, in order to detect increases in these times. By collecting and checking these values regularly you can detect fluctuations in the operation of your queue manager. This can indicate a change in performance.

Obtain real-time monitoring information as follows:

## Procedure

1. To display real-time monitoring information for a queue, use either the WebSphere MQ Explorer or the MQSC command `DISPLAY QSTATUS`, specifying the optional parameter `MONITOR`.
2. To display real-time monitoring information for a channel, use either the WebSphere MQ Explorer or the MQSC command `DISPLAY CHSTATUS`, specifying the optional parameter `MONITOR`.

## Example

The queue, Q1, has the attribute `MONQ` set to the default value, `QMGR`, and the queue manager that owns the queue has the attribute `MONQ` set to `MEDIUM`. To display the monitoring fields collected for this queue, use the following command:

```
DISPLAY QSTATUS(Q1) MONITOR
```

The monitoring fields and monitoring level of queue, Q1 are displayed as follows:

```
QSTATUS(Q1)
TYPE(Queue)
MONQ(MEDIUM)
QTIME(11892157,24052785)
MSGAGE(37)
LPUTDATE(2005-03-02)
LPUTTIME(09.52.13)
LGETDATE(2005-03-02)
LGETTIME(09.51.02)
```

The sender channel, QM1.T0.QM2, has the attribute `MONCHL` set to the default value, `QMGR`, and the queue manager that owns the queue has the attribute `MONCHL` set to `MEDIUM`. To display the monitoring fields collected for this sender channel, use the following command:

```
DISPLAY CHSTATUS(QM1.T0.QM2) MONITOR
```

The monitoring fields and monitoring level of sender channel, QM1.T0.QM2 are displayed as follows:

```
CHSTATUS(QM1.T0.QM2)
XMITQ(Q1)
CONNNAME(127.0.0.1)
CURRENT
CHLTYPE(SDR)
STATUS(RUNNING)
SUBSTATE(MQGET)
MONCHL(MEDIUM)
XQTIME(755394737,755199260)
NETTIME(13372,13372)
EXITTIME(0,0)
```

XBATCSZ(50,50)  
COMPTIME(0,0)  
STOPREQ(NO)  
RQMNAME(QM2)

**Related concepts:**

“Real-time monitoring” on page 1102

**Related reference:**

 `DISPLAY QSTATUS` (*WebSphere MQ V7.1 Reference*)

## Monitoring queues

Use this page to view tasks that help you to resolve a problem with a queue and the application that services that queue. Various monitoring options are available to determine the problem

### About this task

Frequently, the first sign of a problem with a queue that is being serviced is that the number of messages on the queue (CURDEPTH) increases. If you expect an increase at certain times of day or under certain workloads, an increasing number of messages might not indicate a problem. However, if you have no explanation for the increasing number of messages, you might want to investigate the cause.

You might have an application queue where there is a problem with the application, or a transmission queue where there is a problem with the channel. Additional monitoring options are available when the application that services the queue is a channel.

The following examples investigate problems with a particular queue, called Q1, and describe the fields that you look at in the output of various commands:

**Related tasks:**

“Determining whether your application has the queue open”

“Checking that messages on the queue are available” on page 1107

“Checking whether your application is getting messages off the queue” on page 1107

“Determining whether the application can process messages fast enough” on page 1108

“Checking the queue when the current depth is not increasing” on page 1108

### Determining whether your application has the queue open

If you have a problem with a queue, check whether your application has the queue open

#### About this task

Perform the following steps to determine whether your application has the queue open:

#### Procedure

1. Ensure that the application that is running against the queue is the application that you expect. Issue the following command for the queue in question:

```
DISPLAY QSTATUS(Q1) TYPE(HANDLE) ALL
```

In the output, look at the APPLTAG field, and check that the name of your application is shown. If the name of your application is not shown, or if there is no output at all, start your application.

2. If the queue is a transmission queue, look in the output at the CHANNEL field. If the channel name is not shown in the CHANNEL field, determine whether the channel is running.
3. Ensure that the application that is running against the queue has the queue open for input. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

In the output, look at the IPPROCS field to see if any application has the queue open for input. If the value is 0 and this is a user application queue, make sure that the application opens the queue for input to get the messages off the queue.

### Checking that messages on the queue are available

If you have a large number of messages on the queue and your application is not processing any of those messages, check whether the messages on the queue are available to your application

#### About this task

Perform the following steps to investigate why your application is not processing messages from the queue:

#### Procedure

1. Ensure that your application is not asking for a specific message ID or correlation ID when it should be processing all the messages on the queue.
2. Although the current depth of the queue might show that there is an increasing number of messages on the queue, some messages on the queue might not be available to be got by an application, because they are not committed; the current depth includes the number of uncommitted MQPUTs of messages to the queue. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

In the output, look at the UNCOM field to see whether there are any uncommitted messages on the queue.

3. If your application is attempting to get any messages from the queue, check whether the putting application is committing the messages correctly. Issue the following command to find out the names of applications that are putting messages to this queue:
4. Then issue the following command, inserting in <apptag> the APPLTAG value from the output of the previous command:

```
DISPLAY QSTATUS(Q1) TYPE(HANDLE) OPENTYPE(OUTPUT)
```

```
DISPLAY CONN(*) WHERE(APPLTAG EQ <apptag>) UOWSTDA UOWSTTI
```

This shows when the unit of work was started and will help you discover whether the application is creating a long running unit of work. If the putting application is a channel, you might want to investigate why a batch is taking a long time to complete.

### Checking whether your application is getting messages off the queue

If you have a problem with a queue and the application that services that queue, check whether your application is getting messages off the queue

#### About this task

To check whether your application is getting messages off the queue, perform the following checks:

#### Procedure

1. Ensure that the application that is running against the queue is actually processing messages from the queue. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

In the output, look at the LGETDATE and LGETTIME fields which show when the last get was done from the queue.

2. If the last get from this queue was longer ago than expected, ensure that the application is processing messages correctly. If the application is a channel, check whether messages are moving through that channel

### **Determining whether the application can process messages fast enough**

If messages are building up on the queue, but your other checks have not found any processing problems, check that the application can process messages fast enough. If the application is a channel, check that the channel can process messages fast enough.

#### **About this task**

To determine whether the application is processing messages fast enough, perform the following tests:

#### **Procedure**

1. Issue the following command periodically to gather performance data about the queue:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

If the values in the QTIME indicators are high, or are increasing over the period, and you have already ruled out the possibility of long running Units of Work by checking that messages on the queue are available, the getting application might not be keeping up with the putting applications.

2. If your getting application cannot keep up with the putting applications, consider adding another getting application to process the queue. Whether you can add another getting application depends on the design of the application and whether the queue can be shared by more than one application. Features such as message grouping or getting by correlation ID might help to ensure that two applications can process a queue simultaneously.

### **Checking the queue when the current depth is not increasing**

Even if the current depth of your queue is not increasing, it might still be useful to monitor the queue to check whether your application is processing messages correctly.

#### **About this task**

To gather performance data about the queue: Issue the following command periodically:

#### **Procedure**

Issue the following command periodically:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) MSGAGE QTIME
```

In the output, if the value in MSGAGE increases over the period of time, and your application is designed to process all messages, this might indicate that some messages are not being processed at all.

### **Monitoring channels**

Use this page to view tasks that help you to resolve a problem with a transmission queue and the channel that services that queue. Various channel monitoring options are available to determine the problem.

#### **About this task**

Frequently, the first sign of a problem with a queue that is being serviced is that the number of messages on the queue (CURDEPTH) increases. If you expect an increase at certain times of day or under certain workloads, an increasing number of messages might not indicate a problem. However, if you have no explanation for the increasing number of messages, you might want to investigate the cause.

You might have a problem with the channel that services a transmission queue. Various channel monitoring options are available to help you to determine the problem.

The following examples investigate problems with a transmission queue called QM2 and a channel called QM1.TO.QM2. This channel is used to send messages from queue manager, QM1, to queue manager, QM2. The channel definition at queue manager QM1 is either a sender or server channel, and the channel definition at queue manager, QM2, is either a receiver or requester channel.

#### **Related tasks:**

“Determining whether the channel is running”

“Checking that the channel is moving messages” on page 1111

“Checking why a batch takes a long time to complete” on page 1112

“Determining whether the channel can process messages fast enough” on page 1112

“Solving problems with cluster channels” on page 1113

### **Determining whether the channel is running**

If you have a problem with a transmission queue, check whether the channel is running.

#### **About this task**

Perform the following steps to check the status of the channel that is servicing the transmission queue:

#### **Procedure**

1. Issue the following command to find out which channel you expect to process the transmission queue QM2:

```
DIS CHANNEL(*) WHERE(XMITQ EQ QM2)
```

In this example, the output of this command shows that the channel servicing the transmission queue is QM1.TO.QM2

2. Issue the following command to determine the status of the channel, QM1.TO.QM2:  

```
DIS CHSTATUS(QM1.TO.QM2) ALL
```
3. Inspect the STATUS field of the output from the **CHSTATUS** command:
  - If the value of the STATUS field is RUNNING, check that the channel is moving messages
  - If the output from the command shows no status, or the value of the STATUS field is STOPPED, RETRY, BINDING, or REQUESTING, perform the appropriate step, as follows:
4. Optional: If the value of the STATUS field shows no status, the channel is inactive, so perform the following steps:
  - a. If the channel should have been started automatically by a trigger, check that the messages on the transmission queue are available. If there are messages available on the transmission queue, check that the trigger settings on the transmission queue are correct.
  - b. Issue the following command to start the channel again manually:  

```
START CHANNEL(QM1.TO.QM2)
```
5. Optional: If the value of the STATUS field is STOPPED, perform the following steps:
  - a. Check the error logs to determine why the channel stopped. If the channel stopped owing to an error, correct the problem. Ensure also that the channel has values specified for the retry attributes: *SHORTRTY* and *LONGRTY*. In the event of transient failures such as network errors, the channel will then attempt to restart automatically.
  - b. Issue the following command to start the channel again manually:  

```
START CHANNEL(QM1.TO.QM2)
```

On WebSphere MQ for z/OS, you can detect when a user stops a channel by using command event messages.

6. Optional: If the value of the STATUS field is RETRY, perform the following steps:
- Check the error logs to identify the error, then correct the problem.
  - Issue the following command to start the channel again manually:  
START CHANNEL(QM1.TO.QM2)
- or wait for the channel to connect successfully on its next retry.
7. Optional: If the value of the STATUS field is BINDING or REQUESTING, the channel has not yet successfully connected to the partner. Perform the following steps:
- Issue the following command, at both ends of the channel, to determine the substate of the channel:  
DIS CHSTATUS(QM1.TO.QM2) ALL

**Note:**

- In some cases there might be a substate at one end of the channel only.
  - Many substates are transitory, so issue the command a few times to detect whether a channel is stuck in a particular substate.
- Check Table 114 to determine what action to take:

Table 114. Substates seen with status binding or requesting

| Initiating MCA substate <sup>1</sup>     | Responding MCA substate <sup>2</sup>     | Notes                                                                                                                                                                                                                                                                               |
|------------------------------------------|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NAMESERVER                               |                                          | The initiating MCA is waiting for a name server request to complete. Ensure that the correct host name has been specified in the channel attribute, CONNAME, and that your name servers are set up correctly.                                                                       |
| SCYEXIT                                  | SCYEXIT                                  | The MCAs are currently <i>in conversation</i> through a security exit. For more information, see "Determining whether the channel can process messages fast enough" on page 1112.                                                                                                   |
|                                          | CHADEXIT                                 | The channel autodefinition exit is currently executing. For more information, see "Determining whether the channel can process messages fast enough" on page 1112.                                                                                                                  |
| RCVEXIT<br>SENDEXIT<br>MSGEXIT<br>MREXIT | RCVEXIT<br>SENDEXIT<br>MSGEXIT<br>MREXIT | Exits are called at channel startup for MQXR_INIT. Review the processing in this part of your exit if this takes a long time. For more information, see "Determining whether the channel can process messages fast enough" on page 1112.                                            |
| SERIALIZE                                | SERIALIZE                                | This substate only applies to channels with a disposition of SHARED.                                                                                                                                                                                                                |
| NETCONNECT                               |                                          | This substate is shown if there is a delay in connecting due to incorrect network configuration.                                                                                                                                                                                    |
| SSLHANDSHAKE                             | SSLHANDSHAKE                             | An SSL handshake consists of a number of sends and receives. If network times are slow, or connection to lookup CRLs are slow, this affects the time taken to do the handshake.<br><br>On WebSphere MQ for z/OS this substate can also be indicative of not having enough SSLTASKS. |

Table 114. Substates seen with status binding or requesting (continued)

| Initiating MCA substate <sup>1</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Responding MCA substate <sup>2</sup> | Notes |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|-------|
| <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. The initiating MCA is the end of the channel which started the conversation. This can be senders, cluster-senders, fully-qualified servers and requesters. In a server-requester pair, it is the end from which you started the channel.</li> <li>2. The responding MCA is the end of the channel which responded to the request to start the conversation. This can be receivers, cluster-receivers, requesters (when the server or sender is started), servers (when the requester is started) and senders (in a requester-sender call-back pair of channels).</li> </ol> |                                      |       |

## Checking that the channel is moving messages

If you have a problem with a transmission queue, check that the channel is moving messages

### Before you begin

Issue the command `DIS CHSTATUS(QM1.TO.QM2) ALL`. If the value of the STATUS field is RUNNING, the channel has successfully connected to the partner system.

Check that there are no uncommitted messages on the transmission queue, as described in “Checking that messages on the queue are available” on page 1107.

### About this task

If there are messages available for the channel to get and send, perform the following checks:

### Procedure

1. In the output from the display channel status command, `DIS CHSTATUS(QM1.TO.QM2) ALL`, look at the following fields:

#### MSGSG

Number of messages sent or received (or, for server-connection channels, the number of MQI calls handled) during this session (since the channel was started).

#### BUFSSSENT

Number of transmission buffers sent. This includes transmissions to send control information only.

#### BYTSSSENT

Number of bytes sent during this session (since the channel was started). This includes control information sent by the message channel agent.

#### LSTMSGDA

Date when the last message was sent or MQI call was handled, see LSTMSGTI.

#### LSTMSGTI

Time when the last message was sent or MQI call was handled. For a sender or server, this is the time the last message (the last part of it if it was split) was sent. For a requester or receiver, it is the time the last message was put to its target queue. For a server-connection channel, it is the time when the last MQI call completed.

#### CURMSGSG

For a sending channel, this is the number of messages that have been sent in the current batch. For a receiving channel, it is the number of messages that have been received in the current batch. The value is reset to zero, for both sending and receiving channels, when the batch is committed.

2. Determine whether the channel has sent any messages since it started. If any have been sent, determine when the last message was sent.
3. If the channel has started a batch that has not yet completed, as indicated by a non-zero value in CURMSGs, the channel might be waiting for the other end of the channel to acknowledge the batch. Look at the SUBSTATE field in the output and refer to Table 115:

Table 115. Sender and receiver MCA substates

| Sender SUBSTATE                                                                                                                                                                      | Receiver SUBSTATE | Notes                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MQGET                                                                                                                                                                                | RECEIVE           | Normal states of a channel at rest.                                                                                                                                                                            |
| SEND                                                                                                                                                                                 | RECEIVE           | SEND is usually a transitory state. If SEND is seen it indicates that the communication protocol buffers have filled. This can indicate a network problem.                                                     |
| RECEIVE                                                                                                                                                                              |                   | If the sender is seen in RECEIVE substate for any length of time, it is waiting on a response, either to a batch completion or a heartbeat. You might want to check why a batch takes a long time to complete. |
| <p><b>Note:</b> You might also want to determine whether the channel can process messages fast enough, especially if the channel has a substate associated with exit processing.</p> |                   |                                                                                                                                                                                                                |

## Checking why a batch takes a long time to complete

Use this page to view some reasons why a batch can take a long time to complete.

### About this task

When a sender channel has sent a batch of messages it waits for confirmation of that batch from the receiver, unless the channel is pipelined. The following factors can affect how long the sender channel waits:

### Procedure

- Check whether the network is slow. A slow network can affect the time it takes to complete a batch. The measurements that result in the indicators for the NETTIME field are measured at the end of a batch. However, the first batch affected by a slowdown in the network is not indicated with a change in the NETTIME value because it is measured at the end of the batch.
- Check whether the channel is using message retry. If the receiver channel fails to put a message to a target queue, it might use message retry processing, rather than put the message to a dead-letter immediately. Retry processing can cause the batch to slow down. In between MQPUT attempts, the channel will have STATUS(PAUSED), indicating that it is waiting for the message retry interval to pass.

## Determining whether the channel can process messages fast enough

If there messages are building up on the transmission queue, but you have found no processing problems, determine whether the channel can process messages fast enough.

### Before you begin

Issue the following command repeatedly over a period of time to gather performance data about the channel:

```
DIS CHSTATUS(QM1.TO.QM2) ALL
```

### About this task

Confirm that there are no uncommitted messages on the transmission queue, as described in “Checking that messages on the queue are available” on page 1107, then check the XQTIME field in the output from

the display channel status command. When the values of the XQTIME indicators are consistently high, or increase over the measurement period, the indication is that the channel is not keeping pace with the putting applications.

Perform the following tests:

### Procedure

1. Check whether exits are processing. If exits are used on the channel that is delivering these messages, they might add to the time spent processing messages. To identify if this is the case, do the following checks:
  - a. In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `EXITTIME` field. If the time spent in exits is higher than expected, review the processing in your exits for any unnecessary loops or extra processing, especially in message, send, and receive exits. Such processing affects all messages moved across the channel.
  - b. In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `SUBSTATE` field. If the channel has one of the following substates for a significant time, review the processing in your exits:
    - SCYEXIT
    - RCVEXIT
    - SENDEXIT
    - MSGEXIT
    - MREXIT
2. Check whether the network is slow. If messages are not moving fast enough across a channel, it might be because the network is slow. To identify if this is the case, do the following checks:
  - a. In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `NETTIME` field. These indicators are measured when the sending channel asks its partner for a response. This happens at the end of each batch and, when a channel is idle during heartbeating.
  - b. If this indicator shows that round trips are taking longer than expected, use other network monitoring tools to investigate the performance of your network.
3. Check whether the channel is using compression. If the channel is using compression, this adds to the time spent processing messages. If the channel is using only one compression algorithm, do the following checks:
  - a. In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `COMPTIME` field. These indicators show the time spent during compression or decompression.
  - b. If the chosen compression is not reducing the amount of data to send by the expected amount, change the compression algorithm.
4. If the channel is using multiple compression algorithms, do the following checks:
  - a. In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `COMPTIME`, `COMPHDR`, and `COMPMSG` fields.
  - b. Change the compression algorithms specified on the channel definition, or consider writing a message exit to override the channel's choice of compression algorithm for particular messages if the rate of compression, or choice of algorithm, is not providing the required compression or performance.

### Solving problems with cluster channels

If you have a build up of messages on the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` queue, the first step in diagnosing the problem is discovering which channel, or channels, are having a problem delivering messages.

## About this task

To discover which channel, or channels, using the SYSTEM.CLUSTER.TRANSMIT.QUEUE are having a problem delivering messages. Perform the following checks:

### Procedure

1. Issue the following command:

```
DIS CHSTATUS(*) WHERE(XQMSGSA GT 1)
```

**Note:** If you have a busy cluster that has many messages moving, consider issuing this command with a higher number to eliminate the channels that have only a few messages available to deliver.

2. Look through the output for the channel, or channels, that have large values in the field XQMSGSA. Determine why the channel is not moving messages, or is not moving them fast enough. Use the tasks outlined in “Monitoring channels” on page 1108 to diagnose the problems with the channels found to be causing the build up.

## The Windows performance monitor

In WebSphere MQ Version 7.0 and earlier versions, it was possible to monitor the performance of local queues on Windows systems by using the Windows performance monitor. As of WebSphere MQ Version 7.1, this method of performance monitoring is no longer available.

You can monitor queues on all supported platforms by using methods described in “Real-time monitoring” on page 1102.

---

## Monitoring performance and resource usage

Use this topic to understand the facilities available to monitor the performance, and resource usage of your IBM WebSphere MQ for z/OS subsystems.

### Related concepts:

“Introduction to monitoring”

“Interpreting WebSphere MQ performance statistics” on page 1122

“Interpreting WebSphere MQ accounting data” on page 1137

 Planning your IBM WebSphere MQ environment on z/OS (*WebSphere MQ V7.1 Installing Guide*)

 Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*)

“Administering IBM WebSphere MQ for z/OS” on page 237

“Introduction to monitoring”

“Interpreting WebSphere MQ performance statistics” on page 1122

“Interpreting WebSphere MQ accounting data” on page 1137

 Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*)

“Administering IBM WebSphere MQ for z/OS” on page 237

## Introduction to monitoring

Use this topic as an overview of the monitoring facilities available for WebSphere MQ for z/OS. For example, obtaining snapshots, using WebSphere MQ trace, online monitoring, and events.

This topic describes how to monitor the performance and resource usage of WebSphere MQ.

- It outlines some of the information that you can retrieve and briefly describes a general approach to investigating performance problems. (You can find information about dealing with performance problems in the “Problem determination on z/OS” on page 1265.)

- It describes how you can collect statistics about the performance of WebSphere MQ by using SMF records.
- It describes how to gather accounting data to enable you to charge your customers for their use of your WebSphere MQ systems.
- It describes how to use WebSphere MQ events (alerts) to monitor your systems.

Here are some of the tools you might use to monitor WebSphere MQ; they are described in the sections that follow:

- Tools provided by WebSphere MQ:
  - Using DISPLAY commands
  - “Using CICS adapter statistics” on page 1116
  - “Using IBM WebSphere MQ events” on page 1118
- z/OS service aids:
  - “Using System Management Facility” on page 1119
- Other IBM licensed programs:
  - Using the Resource Measurement Facility
  - Using Tivoli Decision Support for z/OS
  - Using the CICS monitoring facility

Information about interpreting the data gathered by the performance statistics trace is given in “Interpreting WebSphere MQ performance statistics” on page 1122.

Information about interpreting the data gathered by the accounting trace is given in “Interpreting WebSphere MQ accounting data” on page 1137.

**Related concepts:**

- “Getting snapshots of WebSphere MQ using the DISPLAY commands”
- “Using CICS adapter statistics” on page 1116
- “Using IBM WebSphere MQ trace” on page 1116
- “Using IBM WebSphere MQ online monitoring” on page 1118
- “Using IBM WebSphere MQ events” on page 1118
- “Using System Management Facility” on page 1119
- “Using other products with WebSphere MQ” on page 1120
- “Investigating performance problems” on page 1121

**Getting snapshots of WebSphere MQ using the DISPLAY commands**

WebSphere MQ provides the MQSC facility which can give a snapshot of the performance, and resource usage using the DISPLAY commands.

You can get an idea of the current state of WebSphere MQ by using the DISPLAY commands and, for the CICS adapter, the CICS adapter panels.

**Using DISPLAY commands**

You can use the WebSphere MQ MQSC DISPLAY or PCF Inquire commands to obtain information about the current state of WebSphere MQ. They provide information about the status of the command server, process definitions, queues, the queue manager, and its associated components. These commands are:

| MQSC command     | PCF command                               |
|------------------|-------------------------------------------|
| DISPLAY ARCHIVE  | Inquire Archive                           |
| DISPLAY AUTHINFO | Inquire Authentication Information Object |
| DISPLAY CFSTATUS | Inquire CF Structure Status               |
| DISPLAY CFSTRUCT | Inquire CF Structure                      |
| DISPLAY CHANNEL  | Inquire Channel                           |
| DISPLAY CHINIT   | Inquire Channel Initiator                 |
| DISPLAY CHSTATUS | Inquire Channel Status                    |
| DISPLAY CMDSERV  |                                           |
| DISPLAY CLUSQMGR | Inquire Cluster Queue Manager             |
| DISPLAY CONN     | Inquire Connection                        |
| DISPLAY GROUP    | Inquire Group                             |
| DISPLAY LOG      | Inquire Log                               |
| DISPLAY PROCESS  | Inquire Process                           |
| DISPLAY QMGR     | Inquire Queue Manager                     |
| DISPLAY QSTATUS  | Inquire Queue Status                      |
| DISPLAY QUEUE    | Inquire Queue                             |
| DISPLAY SECURITY | Inquire Security                          |
| DISPLAY STGCLASS | Inquire Storage Class                     |
| DISPLAY SYSTEM   | Inquire System                            |
| DISPLAY TRACE    |                                           |
| DISPLAY USAGE    | Inquire Usage                             |

For the detailed syntax of each command, see  [WebSphere MQ Script \(MQSC\) Command Reference \(WebSphere MQ V7.1 Reference\)](#) or  [WebSphere MQ Programmable Command Formats and Administration Interface \(WebSphere MQ V7.1 Reference\) manual](#). All of the functions of these commands (except DISPLAY CMDSERV and DISPLAY TRACE) are also available through the operations and control panels.

These commands provide a snapshot of the system *only* at the moment the command was processed. If you want to examine trends in the system, you must start a WebSphere MQ trace and analyze the results over a period of time.

## Using CICS adapter statistics

If you are an authorized CICS user, you can use the CICS adapter control panels to display CICS adapter statistics dynamically.

These statistics provide a snapshot of information related to CICS thread usage and situations when all threads are busy. The display connection panel can be refreshed by pressing the Enter key. For more information, see “The CICS-WebSphere MQ Adapter” section in the CICS Transaction Server for z/OS Version 4.1 product documentation at:  [CICS Transaction Server for z/OS Version 4.1, The CICS-WebSphere MQ adapter](#).

## Using IBM WebSphere MQ trace

You can record performance statistics and accounting data for IBM WebSphere MQ by using the IBM WebSphere MQ trace facility. Use this topic to understand how to control IBM WebSphere MQ trace.

The data generated by IBM WebSphere MQ is sent to:

- The System Management Facility (SMF), specifically as SMF record type 115, subtypes 1 and 2 for the performance statistics trace
- The SMF, specifically as SMF record type 116, subtypes zero, 1, and 2 for the accounting trace.

If you prefer, the data generated by the IBM WebSphere MQ accounting trace can also be sent to the generalized trace facility (GTF).

### Starting IBM WebSphere MQ trace

You can start the IBM WebSphere MQ trace facility at any time by issuing the IBM WebSphere MQ START TRACE command.

Accounting data can be lost if the accounting trace is started or stopped while applications are running. To collect accounting data successfully, the following conditions must apply:

- The accounting trace must be active when an application starts, and it must still be active when the application finishes.
- If the accounting trace is stopped, any accounting data collection that was active stops.

You can also start collecting some trace information automatically if you specify YES on the SMFSTAT (SMF STATISTICS) and SMFACCT (SMF ACCOUNTING) parameters of the CSQ6SYSP macro. These parameters are described in  Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*).

You cannot use this method to start collecting class 3 accounting information (thread-level and queue-level accounting). You must use the START TRACE command to collect such information. However, you can include the command in your CSQINP2 input data set so that the trace is started automatically when you start your queue manager.

Before starting a IBM WebSphere MQ trace, read “Using System Management Facility” on page 1119.

### Controlling IBM WebSphere MQ trace

To control the IBM WebSphere MQ trace data collection at start-up, specify values for the parameters in the CSQ6SYSP macro when you customize IBM WebSphere MQ. See  Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*) for details.

You can control IBM WebSphere MQ tracing when the queue manager is running with these commands:

- START TRACE
- ALTER TRACE
- STOP TRACE

You can choose the destination to which trace data is sent. Possible destinations are:

**SMF** System Management Facility

**GTF** Generalized Trace Facility (accounting trace only)

**SRV** Serviceability routine for diagnostic use by IBM service personnel

For daily monitoring, information is sent to SMF (the default destination). SMF data sets typically contain information from other systems; this information is not available for reporting until the SMF data set is dumped.

You can also send accounting trace information to the GTF. This information has an event identifier of 5EE. The “The MQI call and user parameter, and z/OS generalized trace facility (GTF)” on page 1245 describes how to deal with IBM WebSphere MQ trace information sent to the GTF.

For information about IBM WebSphere MQ commands, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) manual.

## Effect of trace on IBM WebSphere MQ performance

Using the IBM WebSphere MQ trace facility can have a significant effect on IBM WebSphere MQ and transaction performance. For example, if you start a global trace for class 1 or for all classes, it is likely to increase processor usage and transaction response times by approximately 50%. However, if you start a global trace for classes 2 - 4 alone, the increase in processor usage and transaction response times is likely to be less than 1% additional processor cost to the cost of IBM WebSphere MQ calls. The same applies for a statistics or accounting trace.

## Using IBM WebSphere MQ online monitoring

You can collect monitoring data for queues and channels (including automatically defined cluster-server channels) by setting the MONQ, MONCHL, and MONACLS attributes.

Table 116 summarizes the commands to set these attributes at different levels and to display the monitoring information.

*Table 116. Setting and displaying attributes to control online monitoring*

| Attribute | Applicable at this level | Set using command                                              | Display monitoring information using command |
|-----------|--------------------------|----------------------------------------------------------------|----------------------------------------------|
| MONQ      | Queue                    | DEFINE QLOCAL<br>DEFINE QMODEL<br>ALTER QLOCAL<br>ALTER QMODEL | DISPLAY QSTATUS                              |
|           | Queue manager            | ALTER QMGR                                                     |                                              |
| MONCHL    | Channel                  | DEFINE CHANNEL<br>ALTER CHANNEL                                | DISPLAY CHSTATUS                             |
|           | Queue manager            | ALTER QMGR                                                     |                                              |
| MONACLS   | Queue manager            | ALTER QMGR                                                     |                                              |

For full details of these commands, see  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*). For more information about online monitoring, see “Monitoring and performance” on page 829.

## Using IBM WebSphere MQ events

IBM WebSphere MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. You can monitor the operation of all your queue managers by incorporating these events into your own system management application.

IBM WebSphere MQ instrumentation events fall into the following categories:

### Queue manager events

These events are related to the definitions of resources within queue managers. For example, an application attempts to put a message to a queue that does not exist.

### Performance events

These events are notifications that a threshold condition has been reached by a resource. For example, a queue depth limit has been reached, or the queue was not serviced within a predefined time limit.

### Channel events

These events are reported by channels as a result of conditions detected during their operation. For example, a channel instance is stopped.

### Configuration events

These events are notifications that an object has been created, changed, or deleted.

When an event occurs, the queue manager puts an *event message* on the appropriate *event queue*, if defined. The event message contains information about the event that can be retrieved by a suitable IBM WebSphere MQ application.

IBM WebSphere MQ events can be enabled using the IBM WebSphere MQ commands or the operations and control panels.

See “Event types” on page 832 for information about the IBM WebSphere MQ events that generate messages, and for information about the format of these messages. See  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for information about enabling the events.

## Using System Management Facility

You can use SMF to collect statistics and accounting information. To use SMF, certain parameters must be set in z/OS and in WebSphere MQ.

System management facility (SMF) is a z/OS service aid used to collect information from various z/OS subsystems. This information is dumped and reported periodically, for example, hourly. You can use SMF with the WebSphere MQ trace facility to collect data from WebSphere MQ. In this way you can monitor *trends*, for example, in system utilization and performance, and collect accounting information about each user ID using WebSphere MQ.

To record performance statistics (record type 115) to SMF specify the following in the SMFPRMxx member of SYS1.PARMLIB or with the SETSMF z/OS operator command.  
SYS(TYPE(115))

To record accounting information (record type 116) to SMF specify the following in the SMFPRMxx member of SYS1.PARMLIB or with the SETSMF z/OS operator command.  
SYS(TYPE(116))

You can turn on or off the recording of accounting information at the queue or queue manager level using the ACCTQ parameter of the DEFINE QLOCAL, DEFINE QMODEL, ALTER QLOCAL, ALTER QMODEL, DEFINE QMGR, or ALTER QMGR commands. See  The MQSC commands (*WebSphere MQ V7.1 Reference*) for details of these commands.

To use the z/OS command SETSMF, either PROMPT(ALL) or PROMPT(LIST) must be specified in the SMFPRMxx member. See the z/OS *MVS Initialization and Tuning Reference* and the z/OS *MVS System Commands* manuals for more information.

You must also set the SMFSTAT and SMFACCT parameters to YES; this is described in  Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*).

You can specify the interval at which WebSphere MQ collects statistics and accounting data in one of two ways:

- You can specify a value for STATIME in your system parameters (described in  Using CSQ6SYSP (*WebSphere MQ V7.1 Installing Guide*)).
- You can specify zero for STATIME and use the SMF global accounting interval (described in the *z/OS MVS Initialization and Tuning Reference*).

SMF must be running before you can send data to it. For more information about SMF, see the *MVS System Management Facilities (SMF)* manual.

For the statistics and accounting data to be reset, at least one MQI call must be issued during the accounting interval.

### **Allocating additional SMF buffers**

When you start a trace, you must ensure that you allocate adequate SMF buffers. Specify SMF buffering on the VSAM BUFSP parameter of the access method services DEFINE CLUSTER statement. Specify CISZ(4096) and BUFSP(81920) on the DEFINE CLUSTER statement for each SMF VSAM data set.

If an SMF buffer shortage occurs, SMF rejects any trace records sent to it. WebSphere MQ sends a CSQW133I message to the z/OS console when this occurs. WebSphere MQ treats the error as temporary and remains active even though SMF data can be lost. When the shortage has been alleviated and trace recording has resumed, WebSphere MQ sends a CSQW123I message to the z/OS console.

### **Reporting data in SMF**

You can use the SMF program IFASMFDP to dump SMF records to a sequential data set so that they can be processed.

There are several ways to report on this data, for example:

- Write an application program to read and report information from the SMF data set. You can then tailor the report to fit your exact needs.
- Use Performance Reporter to process the records. For more information, see “Using other products with WebSphere MQ.”

### **Using other products with WebSphere MQ**

You can use other products to help you to improve the presentation of, or to augment statistics related to, performance and accounting. For example, Resource Measurement Facility, Tivoli Decision Support, and CICS monitoring.

#### **Using Resource Measurement Facility**

Resource Measurement Facility (RMF™) is an IBM licensed program (program number 5685-029) that provides system-wide information about processor utilization, I/O activity, storage, and paging. You can use RMF to monitor the utilization of physical resources across the whole system dynamically. For more information, see the *MVS Resource Measurement Facility User's Guide*.

#### **Using Tivoli Decision Support for z/OS**

You can use Tivoli Decision Support for z/OS to interpret RMF and SMF records.

Tivoli Decision Support for z/OS is an IBM licensed program (program number 5698-A07) that enables you to manage the performance of your system by collecting performance data in a Db2 database and presenting the data in various formats for use in systems management. Tivoli Decision Support for can generate graphic and tabular reports using systems management data it stores in its Db2 database. It includes an administration dialog, a reporting dialog, and a log collector, all of which interact with a standard Db2 database.

This is described in the *Tivoli Decision Support Administrator's Guide*.

## Using the CICS monitoring facility

The CICS monitoring facility provides performance information about each CICS transaction running. It can be used to investigate the resources used and the time spent processing transactions. For background information, see the *CICS Performance Guide* and the *CICS Customization Guide*.

## Investigating performance problems

Performance problems can arise from various factors. For example, incorrect resource allocation, poor application design, and I/O restraints. Use this topic to investigate some of the possible causes of performance problems.

Performance can be adversely affected by:

- Buffer pools that are an incorrect size
- Lack of real storage
- I/O contention for page sets or logs
- Log buffer thresholds that are set incorrectly
- Incorrect setting of the number of log buffers
- Large messages
- Units of recovery that last a long time, incorporating many messages for each sync point
- Messages that remain on a queue for a long time
- RACF auditing
- Unnecessary security checks
- Inefficient program design

When you analyze performance data, always start by looking at the overall system before you decide that you have a specific IBM WebSphere MQ problem. Remember that almost all symptoms of reduced performance are magnified when there is contention. For example, if there is contention for DASD, transaction response times can increase. Also, the more transactions there are in the system, the greater the processor usage and greater the demand for both virtual and real storage.

In such situations, the system shows heavy use of *all* its resources. However, the system is actually experiencing normal system stress, and this stress might be hiding the cause of a performance reduction. To find the cause of such a loss of performance, you must consider all items that might be affecting your active tasks.

## Investigating the overall system

Within IBM WebSphere MQ, the performance problem is either increased response time or an unexpected and unexplained heavy use of resources. First check factors such as total processor usage, DASD activity, and paging. An IBM tool for checking total processor usage is resource management facility (RMF). In general, you must look at the system in some detail to see why tasks are progressing slowly, or why a specific resource is being heavily used.

Start by looking at general task activity, then focus on particular activities, such as specific tasks or a specific time interval.

Another possibility is that the system has limited real storage; therefore, because of paging interrupts, the tasks progress more slowly than expected.

## Investigating individual tasks

You can use the accounting trace to gather information about IBM WebSphere MQ tasks. These trace records tell you a great deal about the activity that the task has performed, and about how much time the

task spent suspended, waiting for latches. The trace record also includes information about how much Db2 and coupling facility activity were performed by the task.

Interpreting IBM WebSphere MQ accounting data is described in “Interpreting WebSphere MQ accounting data” on page 1137.

Long running units of work can be identified by the presence of message CSQR026I in the job log. This message indicates that a task has existed for more than three queue manager checkpoints and its log records have been shunted. For a description of log record shunting, see  The log files (*WebSphere MQ V7.1 Product Overview Guide*).

## Interpreting WebSphere MQ performance statistics

Use this topic as an index to the different SMF records created by WebSphere MQ for z/OS.

WebSphere MQ performance statistics are written as SMF type 115 records. Statistics records are produced periodically at a time interval specified by the STATIME parameter of the CSQ6SYSP system parameter module, or at the SMF global accounting interval if you specify zero for STATIME. The information provided in the SMF records comes from the following components of WebSphere MQ:

|                                       |                                                                                                                                                                                             |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Buffer manager</b>                 | Manages the buffer pools in virtual storage and the writing of pages to page sets as the buffer pools become full. Also manages the reading of pages from page sets.                        |
| <b>Coupling facility manager</b>      | Manages the interface with the coupling facility.                                                                                                                                           |
| <b>Data manager</b>                   | Manages the links between messages and queues. It calls the buffer manager to process the pages with messages on them.                                                                      |
| <b>Db2 manager</b>                    | Manages the interface with the Db2 database that is used as the shared repository.                                                                                                          |
| <b>Lock manager</b>                   | Manages locks for WebSphere MQ for z/OS.                                                                                                                                                    |
| <b>Log manager</b>                    | Manages the writing of log records, which are essential for maintaining the integrity of the system if there is a back out request, or for recovery, if there is a system or media failure. |
| <b>Message manager</b>                | Processes all WebSphere MQ API requests.                                                                                                                                                    |
| <b>Storage manager</b>                | Manages storage for WebSphere MQ for z/OS, for example, storage pool allocation, expansion, and deallocation.                                                                               |
| <b>Topic manager</b>                  | Manages the Topic and Subscription information for WebSphere MQ for z/OS.                                                                                                                   |
| <b>Coupling facility SMDS manager</b> | Manages the shared message data sets (SMDS) for large messages stored in the coupling facility.                                                                                             |

WebSphere MQ statistics can be collected for two subtypes:

- 1 System information, for example, related to the logs and storage.
- 2 Information about number of messages, buffer, and paging information. Queue-sharing group information related to the coupling facility and Db2.

The subtype is specified in the SM115STF field (shown in Table 117 on page 1123).

**Related reference:**

- “Layout of an SMF type 115 record”
- “The SMF header”
- “Self-defining sections” on page 1124
- “Examples of SMF statistics records” on page 1125
- “Processing type 115 SMF records” on page 1127
- “Storage manager data records” on page 1127
- “Log manager data records” on page 1128
- “Message manager data records” on page 1129
- “Data manager data records” on page 1129
- “Buffer manager data records” on page 1129
- “Lock manager data records” on page 1132
- “Db2 manager data records” on page 1132
- “Coupling facility manager data records” on page 1135
- “Topic manager data records” on page 1136
- “Coupling facility manager SMDS data records” on page 1137

**Layout of an SMF type 115 record**

You can use this section as a reference for the format of an SMF type 115 record.

The standard layout for SMF records involves three parts:

**SMF header**

Provides format, identification, and time and date information about the record itself.

**Self-defining section**

Defines the location and size of the individual data records within the SMF record.

**Data records**

The actual data from WebSphere MQ that you want to analyze.

For more information about SMF record formats, see the *MVS System Management Facilities (SMF)* manual.

**Related reference:**

- “The SMF header”
- “Self-defining sections” on page 1124
- “Examples of SMF statistics records” on page 1125

**The SMF header**

Use this topic as a reference for the format of the SMF header.

Table 117 shows the format of SMF record header (SM115).

*Table 117. SMF record 115 header description*

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description        | Example |
|----------------|----------------|-----------|-----|----------|--------------------|---------|
| 0              | 0              | Structure | 28  | SM115    | SMF record header. |         |
| 0              | 0              | Integer   | 2   | SM115LEN | SMF record length. | 14A0    |
| 2              | 2              |           | 2   |          | Reserved.          |         |
| 4              | 4              | Integer   | 1   | SM115FLG | System indicator.  | 5E      |

Table 117. SMF record 115 header description (continued)

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description                                                                                       | Example         |
|----------------|----------------|-----------|-----|----------|---------------------------------------------------------------------------------------------------|-----------------|
| 5              | 5              | Integer   | 1   | SM115RTY | Record type. The SMF record type, for WebSphere MQ statistics records this is always 115 (X'73'). | 73              |
| 6              | 6              | Integer   | 4   | SM115TME | Time when SMF moved record.                                                                       | 00355575        |
| 10             | A              | Integer   | 4   | SM115DTE | Date when SMF moved record.                                                                       | 0100223F        |
| 14             | E              | Character | 4   | SM115SID | z/OS subsystem ID. Defines the z/OS subsystem on which the records were collected.                | D4E5F4F1 (MV41) |
| 18             | 12             | Character | 4   | SM115SSI | WebSphere MQ subsystem ID.                                                                        | D4D8F0F7 (MQ07) |
| 22             | 16             | Integer   | 2   | SM115STF | Record subtype.                                                                                   | 0002            |
| 24             | 18             | Character | 3   | SM115REL | WebSphere MQ version.                                                                             | F6F0F0 (600)    |
| 27             | 1B             |           | 1   |          | Reserved                                                                                          |                 |
| 28             | 1C             | Character | 0   | SM115END | End of SMF header and start of self-defining section.                                             |                 |

**Note:** The (hexadecimal) values in the "Example" column relate to Figure 128 on page 1125.

## Self-defining sections

Use this topic as a reference for format of the self-defining sections of the SMF record.

A self-defining section of a type 115 SMF record tells you where to find a statistics record, how long it is, and how many times that type of record is repeated (with different values). The self-defining sections follow the header, at fixed offsets from the start of the SMF record. Each statistics record can be identified by an eye-catcher string.

The following types of self-defining section are available to users for type 115 records. Each self-defining section points to statistics data related to one of the WebSphere MQ components. Table 118 summarizes the sources of the statistics, the eye-catcher strings, and the offsets of the self-defining sections from the start of the SMF record header.

Table 118. Offsets to self-defining sections. Offsets are from the start of the SMF record and are fixed for each type of statistics source.

| Source of statistics                      | Record subtype<br>(SM115STF) | Offset of self-defining section |       | Eye-catcher of data |
|-------------------------------------------|------------------------------|---------------------------------|-------|---------------------|
|                                           |                              | Dec                             | Hex   |                     |
| Storage manager                           | 1                            | 100                             | X'64' | QSST                |
| Log manager                               | 1                            | 116                             | X'74' | QJST                |
| Message manager                           | 2                            | 36                              | X'24' | QMST                |
| Data manager                              | 2                            | 44                              | X'2C' | QIST                |
| Buffer manager - one for each buffer pool | 2                            | 52                              | X'34' | QPST                |
| Lock manager                              | 2                            | 60                              | X'3C' | QLST                |
| Db2 manager                               | 2                            | 68                              | X'44' | Q5ST                |
| Coupling Facility manager                 | 2                            | 76                              | X'4C' | QEST                |
| Topic manager                             | 2                            | 84                              | X'54' | QTST                |
| SMDS usage                                | 2                            | 92                              | X'5C' | QESD                |

**Note:** Other self-defining sections refer to data for IBM use only.

Each self-defining section is two fullwords long and has this format:

```
sssssss1111nnnn
```

where:

**sssssss** Fullword containing the offset from the start of the SMF record.  
**1111** Halfword giving the length of this data record.  
**nnnn** Halfword giving the number of data records in this SMF record.

Figure 128 shows an example of part of an SMF type 115 subtype 2 record. The numbers in the first column represent the offset, in hexadecimal, from the start of the record. Each line corresponds to 16 bytes of data, where each byte is two hexadecimal characters, for example 0C. The characters in the last column represent the printable characters for each byte. Non-printable characters are shown by a period (.) character.

In this example, alternate fields in the SMF header are underlined to help you to see them; refer to Table 117 on page 1123 to identify them. The self-defining sections for the message manager statistics data records (at the offset given in Table 118 on page 1124) and buffer manager statistics are shown in **bold**.

|        |                 |                 |                 |          |                   |
|--------|-----------------|-----------------|-----------------|----------|-------------------|
| 000000 | 14A00000        | 5E730035        | 55750100        | 223FD4E5 | *....;.....MV*    |
| 000010 | <u>F4F1D4D8</u> | <u>F0F70002</u> | <u>F6F0F000</u> | 0000147C | *41MQ07..600...@* |
| 000020 | 00240001        | <b>00000054</b> | <b>00300001</b> | 00000084 | *.....*           |
| 000030 | 00500001        | <b>000000D4</b> | <b>00680004</b> | 00000274 | *.&.....M.....*   |
| 000040 | 00200001        | 00000294        | 01E00001        | 00000474 | *.....*           |
| 000050 | 10080001        | D40F0030        | D8D4E2E3        | 00000000 | *...M...QMST...*  |

Figure 128. Part of an SMF record 115 showing the header and self-defining sections

The self-defining section for message manager statistics is located at offset X'20' from the start of the SMF record and contains this information:

- The offset of the message manager statistics is located X'00000054' bytes from the start of the SMF record.
- The message manager record is X'0030' bytes long.
- There is one record (X'0001').

Similarly, the buffer manager self-defining section at X'30' specifies that the offset to the buffer manager statistics is X'000000D4', is of length X'0068', and occurs X'0004' times.

**Note:** Always use offsets in the self-defining sections to locate the statistics records.

## Examples of SMF statistics records

Use this topic to understand some example SMF records.

Figure 129 on page 1126 shows an example of part of the SMF record for subtype 1. Subtype 1 includes the storage manager and log manager statistics records. The SMF record header is shown underlined.

The self-defining section at offset X'64' refers to storage manager statistics and the self-defining section at offset X'74' refers to log manager statistics, both shown in **bold**.

The storage manager statistics record is located at offset X'0000011C' from the start of the header and is X'48' bytes long. There is one set of storage manager statistics, identified by the eye-catcher string QSST. The start of this statistics record is also shown in the example.

The log manager statistics record is located at offset X'00000164' from the start of the header and is X'78' bytes long. There is one set of log manager statistics, identified by the eye-catcher string QJST.

|        |                 |                 |                 |          |                    |
|--------|-----------------|-----------------|-----------------|----------|--------------------|
| 000000 | 02000000        | 5E730035        | 55750100        | 223FD4E5 | *....;.....MV*     |
| 000010 | <u>F4F1D4D8</u> | <u>F0F70001</u> | <u>F6F0F000</u> | 000001DC | *41MQ07..600.....* |
| 000020 | 00240001        | 00000000        | 00000000        | 00000000 | *.....*            |
| 000030 | 00000000        | 00000000        | 00000000        | 0000007C | *.....@*           |
| 000040 | 00400001        | 000000BC        | 00600001        | 00000000 | *.....-.....*      |
| 000050 | 00000000        | 00000000        | 00000000        | 00000000 | *.....*            |
| 000060 | 00000000        | <b>0000011C</b> | <b>00480001</b> | 00000000 | *.....*            |
| 000070 | 00000000        | <b>00000164</b> | <b>00780001</b> | 00000000 | *.....*            |
| 000080 | 00000000        | 00000000        | 00000000        | 00000000 | *.....*            |
| .      |                 |                 |                 |          |                    |
| .      |                 |                 |                 |          |                    |
| 000110 | 00000000        | 00000000        | 00000000        | 003C0048 | *.....*            |
| 000120 | D8E2E2E3        | 0000004F        | 00000003        | 00000002 | *QSST... .....*    |

Figure 129. SMF record 115, subtype 1

Figure 130 on page 1127 shows an example of part of the SMF record for subtype 2. Subtype 2 includes the statistics records for the message, data, buffer, lock, coupling facility, and Db2 managers. The SMF record header is shown underlined; the self-defining sections are shown alternately **bold** and *italic*.

- The self-defining section at offset X'24' refers to message manager statistics. The message manager statistics record is located at offset X'00000054' from the start of the header and is X'30' bytes long. There is one set of these statistics, identified by the eye-catcher string QMST.
- The self-defining section at offset X'2C' refers to data manager statistics. The data manager statistics record is located at offset X'00000084' from the start of the header and is X'50' bytes long. There is one set of these statistics, identified by the eye-catcher string QIST.
- The self-defining section at offset X'34' refers to buffer manager statistics. The buffer manager statistics record is located at offset X'000000D4' from the start of the header and is X'68' bytes long. There are four sets of these statistics, identified by the eye-catcher string QPST.
- The self-defining section at offset X'3C' refers to lock manager statistics. The lock manager statistics record is located at offset X'00000274' from the start of the header and is X'20' bytes long. There is one set of these statistics, identified by the eye-catcher string QLST.
- The self-defining section at offset X'44' refers to Db2 manager statistics. The Db2 manager statistics record is located at offset X'00000294' from the start of the header and is X'1E0' bytes long. There is one set of these statistics, identified by the eye-catcher string Q5ST.
- The self-defining section at offset X'4C' refers to coupling facility manager statistics. The coupling facility manager statistics record is located at offset X'00000474' from the start of the header and is X'1008' bytes long. There is one set of these statistics, identified by the eye-catcher string QEST.

```

000000 14A00000 5E730035 55750100 223FD4E5 *....;.....MV*
000010 F4F1D4D8 F0F70002 F6F0F000 0000147C *41MQ07..600...@*
000020 00240001 00000054 00300001 00000084 *......*
000030 00500001 000000D4 00680004 00000274 *.&....M.....*
000040 00200001 00000294 01E00001 00000474 *......*
000050 10080001 D40F0030 D8D4E2E3 00000000 *...M...QMST...*
000060 00000000 00000000 00000000 00000000 *......*
000070 00000000 00000000 00000000 00000000 *......*
000080 00000000 C90F0050 D8C9E2E3 00000000 *...I..&QIST...*
000090 00000001 00000000 00000025 00000003 *......*
0000A0 00000000 0000002C 00000007 00000000 *......*
0000B0 00000000 00000000 00000000 00000012 *......*
0000C0 00000000 00000000 00000000 00000000 *......*
0000D0 00000000 D70F0068 D8D7E2E3 00000000 *...P...QPST...*
0000E0 000007D0 000007BD 000007BD 00000037 *...}.....*
0000F0 00000000 0000001B 0000003B 00000000 *......*
000100 00000000 00000000 00000000 00000000 *......*
000110 0000001B 00000000 00000000 00000000 *......*
000120 00000000 00000000 00000000 00000000 *......*
000130 00000000 00000000 00000000 D70F0068 *...P...*
000140 D8D7E2E3 00000001 000007D0 000007CD *QPST.....}*
.
.

```

Figure 130. SMF record 115, subtype 2

## Processing type 115 SMF records

Use this topic as a reference for processing type 115 SMF records.

You must process any data you collect from SMF to extract useful information. When you process the data, verify that the records are from WebSphere MQ and that they are the records you are expecting.

Validate the values of the following fields:

- SM115RTY, the SMF record number, must be X'73' (115)
- SM115STF, the record subtype, must be 0001 or 0002

Reading from the active SMF data sets is not supported. You must use the SMF program IFASMFDP to dump SMF record to a sequential data set so that they can be processed. For more information see “Using System Management Facility” on page 1119.

There is a C sample program called CSQ4SMFD which prints the contents of SMF type 115 and 116 records. The program is provided as source in thlqual.SCSQC37S and in executable format in thlqual.SCSQLOAD. Sample JCL is provided in thlqual.SCSQPROC(CSQ4SMFJ).

## Storage manager data records

Use this topic as a reference for storage manager data records.

The format of the storage manager statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQSST).

The data contains information about the number of fixed and variable storage pools that the queue manager has allocated, expanded, contracted, and deleted during the statistics interval, plus the number of GETMAIN, FREEMAIN, and STORAGE requests to z/OS, including a count of those requests that were unsuccessful. Additional information includes a count of the number of times the short-on-storage condition was detected and a count of the number of abends that occurred as a result of that condition.

## Log manager data records

Use this topic as a reference for format of log manager data records.

The format of the log manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQJST)`.

In the statistics, these counts are important:

1. The total number of log write requests:

$$N_{\text{logwrite}} = \text{QJSTWRW} + \text{QJSTWRNW} + \text{QJSTWRF}$$

2. The total number of log read requests:

$$N_{\text{logread}} = \text{QJSTRBUF} + \text{QJSTRACT} + \text{QJSTRARH}$$

The problem symptoms that can be examined using log manager statistics are described in the following table.

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Symptom 1</b><br/>QJSTWTB is nonzero.</p> <p><b>Reason</b><br/>Tasks are being suspended while the in-storage buffer is being written to the active log.<br/>There might be problems writing to the active log.<br/>The OUTBUFF parameter within CSQ6LOGP is too small.</p> <p><b>Action</b><br/>Investigate the problems writing to the active log.<br/>Increase the value of the OUTBUFF parameter within CSQ6LOGP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p><b>Symptom 2</b><br/>The ratio: <math>\text{QJSTWTL}/N_{\text{logread}}</math> is greater than 1%.</p> <p><b>Reason</b><br/>Log reads were initiated that had to read from an archive log, but WebSphere MQ could not allocate a data set because MAXRTU data sets were already allocated.</p> <p><b>Action</b><br/>Increase MAXRTU.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <p><b>Symptom 3</b><br/>The ratio: <math>\text{QJSTRARH}/N_{\text{logread}}</math> is larger than normal.</p> <p><b>Reason</b><br/>Most log read requests should come from the output buffer or the active log. To satisfy requests for back out, unit-of-recovery records are read from the in-storage buffer, the active log, and the archived logs.<br/><br/>A long-running unit of recovery, extending over a period of many minutes, might have log records spread across many different logs. This degrades performance because extra work has to be done to recover the log records.</p> <p><b>Action</b><br/>Change the application to reduce the length of a unit of recovery. Also, consider increasing the size of the active log to reduce the possibility of a single unit of recovery being spread out over more than one log.</p> <p><b>Other pointers</b><br/>The ratio <math>N_{\text{logread}}/N_{\text{logwrite}}</math> gives an indication of how much work has to be backed out.</p> |

**Symptom 4**

QJSTLLCP is more than 10 an hour.

**Reason**

On a busy system, you would expect to see typically 10 checkpoints an hour. If the QJSTLLCP value is larger than this, it indicates a problem in the setup of the queue manager.

The most likely reason for this is that the LOGLOAD parameter in CSQ6SYSP is too small. The other event that causes a checkpoint is when an active log fills up and switches to the next active log data set. If your logs are too small, this can cause frequent checkpoints. The QJSTLLCP counter is not incremented for log switch induced checkpoints; you must look in the JES logs for the queue managers to determine if the rate log files are switched.

**Action**

Increase the LOGLOAD parameter, or increase the size of your log data sets as required.

**Symptom 5**

QJSTCmpFail > 0 or QJSTCmpComp not much less than QJSTCmpUncmp

**Reason**

The queue manager is unable to significantly compress log records.

QJSTCmpFail is the number of times the queue manager was unable to achieve any reduction in record length. You should compare the number to QJSTCmpReq (number of compression requests) to see if the number of failures is significant.

QJSTCmpComp is the total of compressed bytes written to the log and QJSTCmpUncmp is the total bytes before compression. Neither total contains bytes written for log records that were not eligible for compression. If the numbers are similar then compression has achieved little benefit.

**Action**

Turn off log compression. Issue the SET LOG COMPLOG(NONE) command. See the  SET LOG (*WebSphere MQ V7.1 Reference*) command for details.

**Note:** In the first set of statistics produced after system startup, there might be significant log activity due to the resolution of in-flight units of recovery.

**Message manager data records**

Use this topic as a reference for message manager data records.

The format of the message manager statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQMST).

The data gives you counts of different WebSphere MQ API requests.

**Data manager data records**

Use this topic as a reference for the format of the Data Manager data records.

The format of the data manager statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQIST).

The data gives you counts of different object requests.

**Buffer manager data records**

Use this topic as a reference for the format of buffer manager data records.

The format of the buffer manager statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQPST).

**Note:** If you have defined a buffer pool, but not used it, no values are set so the buffer manager statistics record does not contain any data.

For information about efficiently managing your buffer pools, see “Managing your buffer pools” on page 1131.

When interpreting the statistics, you are recommended to consider the following factors because the values of these fields can be used to improve the performance of your system:

1. If QPSTSOS, QPSTSTLA, or QPSTDMC is greater than zero, you should either increase the size of the buffer pool or reallocate the page sets to different buffer pools.
  - QPSTSOS is the number of times that there were no buffers available for page get requests. If QPSTSOS ever becomes nonzero, it shows that WebSphere MQ is under severe stress. The buffer pool size should be increased. If increasing the buffer pool size does not make the value of QPSTSOS zero, there might be I/O contention on the DASD page sets.
  - QPSTDMC is the number of updates that were performed synchronously because there was either more than 95% of the pages in the buffer pool waiting for write I/O, or there was less than 5% of the buffer pool available for read requests. If this number is not zero, the buffer pool might be too small and should be enlarged. If increasing the buffer pool size does not reduce QPSTDMC to zero, there might be I/O contention on the DASD page sets.
  - QPSTIMW is a count of the number of times pages were written out synchronously. If QPSTDMC is zero, QPSTIMW is the number of times pages were found on the queue waiting for write I/O that had been there for at least two checkpoints.
2. For buffer pool zero and buffer pools that contain short-lived messages:
  - QPSTDWT should be zero, and the percentage  $QPSTCBSL/QPSTNBUF$  should be greater than 15%. QPSTDWT is the number of times the asynchronous write processor was started because there was either more than 85% of the pages in the buffer pool waiting for write I/O, or there was less than 15% of the buffer pool available for read requests. Increasing the buffer pool size should reduce this value. If it does not, the pattern of access is one of long delays between puts and gets.
  - QPSTTPW might be greater than zero due to checkpointing activity.
  - QPSTRIO should be zero unless messages are being read from a page set after the queue manager is restarted.

The ratio of QPSTRIO to QPSTGETP shows the efficiency of page retrieval within the buffer pool. Increasing the buffer pool size should decrease this ratio and, therefore, increase the page retrieval efficiency. If this does not happen, it indicates that pages are not being frequently reaccessed. This implies a transaction pattern where there is a long delay between messages being put and then later retrieved.

The ratio of QPSTGETN to QPSTGETP indicates the number of times an empty page, as opposed to a non-empty page, has been requested. This ratio is more an indication of transaction pattern, than a value that can be used to tune the system.
  - If QPSTSTL has a value greater than zero, this indicates that pages that have not been used before are now being used. This might be caused by an increased message rate, messages not being processed as fast as they were previously (leading to a buildup of messages), or larger messages being used.

QPSTSTL is a count of the number of times a page access request did not find the page already in the buffer pool. Again, the lower the ratio of QPSTSTL to  $(QPSTGETP + QPSTGETN)$  is, the higher the page retrieval efficiency. Increasing the buffer pool size should decrease this ratio but, if it does not, it is an indication that there are long delays between puts and gets.
  - You are recommended to have sufficient buffers to handle your peak message rate.
3. For buffer pools with long-lived messages, where there are more messages than can fit into the buffer pool:

- $(QPSTRIO+QPSTWIO)/\text{Statistics interval}$  is the I/O rate to page sets. If this value is high, you should consider using multiple page sets on different volumes to allow I/O to be carried out in parallel.

The higher the ratio of QPSTSTW to QPSTWIO, the better the efficiency of the asynchronous write processor. You can increase this ratio, and therefore the efficiency of the asynchronous write processor, by increasing the buffer pool size.

- Over the period of time that the messages are processed (for example, if messages are written to a queue during the day and processed overnight) the number of read I/Os (QPSTRIO) should be approximately the total number of pages written (QPSTTPW). This shows that one page is read for every page written.

If QPSTRIO is much larger than QPSTTPW, this shows that pages are being read in multiple times. This might be a result of the application using **MQGET** by *MsgId* or *CorrelId* when the queue is not indexed, or browsing messages on the queue using get next.

The following actions might relieve this problem:

- a. Increase the size of the buffer pool so that there are enough pages to hold the queue, in addition to any changed pages.
- b. Use the INDXTYPE queue attribute, which allows a queue to be indexed by *MsgId* or *CorrelId* and eliminates the need for a sequential scan of the queue.
- c. Change the design of the application to eliminate the use of **MQGET** with *MsgId* or *CorrelId*, or the get next with browse option.

**Note:** Applications using long-lived messages typically process the first available message and do not use **MQGET** with *MsgId* or *CorrelId*, and they might browse only the first available message.

- d. Move page sets to a different buffer pool to reduce contention between messages from different applications.

#### **Related concepts:**

“Managing your buffer pools”

#### **Managing your buffer pools:**

To manage your buffer pools efficiently, you must consider the factors that affect the buffer pool I/O operations and also the statistics associated with the buffer pools.

The following factors affect buffer pool I/O operations.

- If a page containing the required data is not found in the buffer pool, it is read synchronously to an available buffer from its DASD page set.
- Whenever a page is updated, it is put on an internal queue of pages to be (potentially) written out to DASD. This means that the buffer used by that page is unavailable for use by any other page until the buffer has been written to DASD.
- If the number of pages queued to be written to DASD exceeds 85% of the total number of buffers in the pool, an asynchronous write processor is started to put the buffers to DASD.

Similarly, if the number of buffers available for page get requests become less than 15% of the total number of buffers in the pool, the asynchronous write processor is started to perform the write I/O operations.

The write processor stops when the number of pages queued to be written to DASD has fallen to 75% of the total number of buffers in the pool.

- If the number of pages queued for writing to DASD exceeds 95% of the total number of buffers in the pool, all updates result in a synchronous write of the page to DASD.

Similarly, if the number of buffers available for page get requests becomes less than 5% of the total number of buffers in the pool, all updates result in a synchronous write of the page to DASD.

- If the number of buffers available for page get requests ever reaches zero, a transaction that encounters this condition is suspended until the asynchronous write processor has finished.
- If a page is frequently updated, the page spends most of its time on the queue of pages waiting to be written to DASD. Because this queue is in least recently used order, it is possible that a frequently updated page placed on this least recently used queue is never written out to DASD. For this reason, at the time of update, if the page is found to have been waiting on the write operation to DASD queue for at least 2 checkpoints, it is synchronously written to DASD. Updating occurs at checkpoint time.

The aim of this algorithm is to maximize the time pages spend in buffer pool memory while allowing the system to function if the system load puts the buffer pool usage under stress.

## Lock manager data records

Use this topic as a reference to the format of the lock manager data records.

The format of the lock manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQLST)`.

The records contain data about the following information:

- The number of lock get requests and lock release requests.
- The number of times a lock get request determined that the requested lock was already held.

## Db2 manager data records

Use this topic as a reference to the format of the Db2 manager data records.

The format of the Db2 manager statistics record is described in the following table and in assembler macro `thlqual.SCSQMACS(CSQDQ5ST)` and C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `q5st`, `q5stid`.

If the queue manager was not started as a member of a queue-sharing group, no data is recorded in this record.

*Table 119. Db2 statistics record (Q5ST)*

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description                   |
|----------------|----------------|-----------|-----|----------|-------------------------------|
| 0              | 0              | Structure | 668 | Q5ST     | Db2 manager statistics        |
| 0              | 0              | Bitstring | 2   | Q5STID   | Control block identifier      |
| 2              | 2              | Integer   | 2   | Q5STLL   | Control block length          |
| 4              | 4              | Character | 4   | Q5STEYEC | Control block eye catcher     |
| 8              | 8              | Character | 660 | Q5STZERO | QMST part cleared on occasion |
| 8              | 8              | Integer   | 4   | NUMTASK  | Number of server tasks        |
| 12             | C              | Integer   | 4   | ACTTASK  | Number of active server tasks |
| 16             | 10             | Integer   | 4   | CONNCNT  | Number of connect requests    |
| 20             | 14             | Integer   | 4   | DISCCNT  | Number of disconnect requests |
| 24             | 18             | Integer   | 4   | DHIGMAX  | Max. request queue depth      |
| 28             | 1C             | Integer   | 4   | ABNDCNT  | Number of DB2SRV task abends  |
| 32             | 20             | Integer   | 4   | REQUCNT  | Number of requests queued     |
| 36             | 24             | Integer   | 4   | DEADCNT  | Number of deadlock timeouts   |
| 40             | 28             | Integer   | 4   | DELECNT  | Number of delete requests     |
| 44             | 2C             | Integer   | 4   | LISTCNT  | Number of list requests       |

Table 119. Db2 statistics record (Q5ST) (continued)

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description                                |
|----------------|----------------|-----------|-----|----------|--------------------------------------------|
| 48             | 30             | Integer   | 4   | READCNT  | Number of read requests                    |
| 52             | 34             | Integer   | 4   | UPDTCNT  | Number of update requests                  |
| 56             | 38             | Integer   | 4   | WRITCNT  | Number of write requests                   |
| 60             | 3C             | Integer   | 4   | SCSSEL   | SCST (shared-channel-status) selects       |
| 64             | 40             | Integer   | 4   | SCSINS   | SCST inserts                               |
| 68             | 44             | Integer   | 4   | SCSUPD   | SCST updates                               |
| 72             | 48             | Integer   | 4   | SCSDEL   | SCST deletes                               |
| 76             | 4C             | Integer   | 4   | SSKSEL   | SSKT (shared-sync-key) selects             |
| 80             | 50             | Integer   | 4   | SSKINS   | SSKT inserts                               |
| 84             | 54             | Integer   | 4   | SSKDEL   | SSKT deletes                               |
| 88             | 58             | Integer   | 4   | SCSBFTS  | SCST number of times buffer too small      |
| 92             | 5C             | Integer   | 4   | SCSMAXR  | SCST maximum rows on query                 |
| 96             | 60             | Integer   | 4   | * (2)    | Reserved                                   |
| 104            | 68             | Character | 8   | DELETCUW | Cumulative STCK difference - Thread delete |
| 112            | 70             | Character | 8   | DELETMXW | Maximum STCK difference - Thread delete    |
| 120            | 78             | Character | 8   | DELESCUW | Cumulative STCK difference - SQL delete    |
| 128            | 80             | Character | 8   | DELESMXW | Maximum STCK difference - SQL delete       |
| 136            | 88             | Character | 8   | LISTTCUW | Cumulative STCK difference - Thread list   |
| 144            | 90             | Character | 8   | LISTTMXW | Maximum STCK difference - Thread list      |
| 152            | 98             | Character | 8   | LISTSCUW | Cumulative STCK difference - SQL list      |
| 160            | A0             | Character | 8   | LISTSMXW | Maximum STCK difference - SQL list         |
| 168            | A8             | Character | 8   | READTCUW | Cumulative STCK difference - Thread read   |
| 176            | B0             | Character | 8   | READTMXW | Maximum STCK difference - Thread read      |
| 184            | B8             | Character | 8   | READSCUW | Cumulative STCK difference - SQL read      |
| 192            | C0             | Character | 8   | READSMXW | Maximum STCK difference - SQL read         |
| 200            | C8             | Character | 8   | UPDTTCUW | Cumulative STCK difference - Thread update |
| 208            | D0             | Character | 8   | UPDTTMXW | Maximum STCK difference - Thread update    |
| 216            | D8             | Character | 8   | UPDTSCUW | Cumulative STCK difference - SQL update    |
| 224            | E0             | Character | 8   | UPDTSMXW | Maximum STCK difference - SQL update       |
| 232            | E8             | Character | 8   | WRITTCUW | Cumulative STCK difference - Thread write  |
| 240            | F0             | Character | 8   | WRITTMXW | Maximum STCK difference - Thread write     |
| 248            | F8             | Character | 8   | WRITSCUW | Cumulative STCK difference - SQL write     |
| 256            | 100            | Character | 8   | WRITSMXW | Maximum STCK difference - SQL write        |
| 264            | 108            | Character | 8   | SCSSTCUW | Cumulative STCK difference - Thread select |
| 272            | 110            | Character | 8   | SCSSTMXW | Maximum STCK difference - Thread select    |
| 280            | 118            | Character | 8   | SCSSCUW  | Cumulative STCK difference - SQL select    |
| 288            | 120            | Character | 8   | SCSSMXW  | Maximum STCK difference - SQL select       |
| 296            | 128            | Character | 8   | SCSITCUW | Cumulative STCK difference - Thread insert |
| 304            | 130            | Character | 8   | SCSITMXW | Maximum STCK difference - Thread insert    |

Table 119. Db2 statistics record (Q5ST) (continued)

| Offset:<br>Dec | Offset:<br>Hex | Type           | Len | Name     | Description                                            |
|----------------|----------------|----------------|-----|----------|--------------------------------------------------------|
| 312            | 138            | Character      | 8   | SCSISCUW | Cumulative STCK difference - SQL insert                |
| 320            | 140            | Character      | 8   | SCSISMWX | Maximum STCK difference - SQL insert                   |
| 328            | 148            | Character      | 8   | SCSUTCWU | Cumulative STCK difference - Thread update             |
| 336            | 150            | Character      | 8   | SCSUTMXW | Maximum STCK difference - Thread update                |
| 344            | 158            | Character      | 8   | SCSUSCUW | Cumulative STCK difference - SQL update                |
| 352            | 160            | Character      | 8   | SCSUSMXW | Maximum STCK difference - SQL update                   |
| 360            | 168            | Character      | 8   | SCSDTCWU | Cumulative STCK difference - Thread delete             |
| 368            | 170            | Character      | 8   | SCSDTMXW | Maximum STCK difference - Thread delete                |
| 376            | 178            | Character      | 8   | SCSDSCWU | Cumulative STCK difference - SQL delete                |
| 384            | 180            | Character      | 8   | SCSDSMXW | Maximum STCK difference - SQL delete                   |
| 392            | 188            | Character      | 8   | SSKSTCUW | Cumulative STCK difference - Thread select             |
| 400            | 190            | Character      | 8   | SSKSTMXW | Maximum STCK difference - Thread select                |
| 408            | 198            | Character      | 8   | SSKSSCUW | Cumulative STCK difference - SQL select                |
| 416            | 1A0            | Character      | 8   | SSKSSMXW | Maximum STCK difference - SQL select                   |
| 424            | 1A8            | Character      | 8   | SSKITCUW | Cumulative STCK difference - Thread insert             |
| 432            | 1B0            | Character      | 8   | SSKITMXW | Maximum STCK difference - Thread insert                |
| 440            | 1B8            | Character      | 8   | SSKISCUW | Cumulative STCK difference - SQL insert                |
| 448            | 1C0            | Character      | 8   | SSKISMWX | Maximum STCK difference - SQL insert                   |
| 456            | 1C8            | Character      | 8   | SSKDTCWU | Cumulative STCK difference - Thread delete             |
| 464            | 1D0            | Character      | 8   | SSKDTMXW | Maximum STCK difference - Thread delete                |
| 472            | 1D8            | Character      | 8   | SSKDSCWU | Cumulative STCK difference - SQL delete                |
| 480            | 1E0            | Character      | 8   | SSKDSMXW | Maximum STCK difference - SQL delete                   |
| 488            | 1E8            | Integer        | 4   | LMSSEL   | Number of Db2 BLOB read requests                       |
| 492            | 1EC            | Integer        | 4   | LMSINS   | Number of Db2 BLOB insert requests                     |
| 496            | 1F0            | Integer        | 4   | LMSUPD   | Number of Db2 BLOB update requests                     |
| 500            | 1F4            | Integer        | 4   | LMSDEL   | Number of Db2 BLOB delete requests                     |
| 504            | 1F8            | Integer        | 4   | LMSLIS   | Number of Db2 BLOB list requests                       |
| 508            | 1FC            | 64 bit integer | 8   | LMSSTCUW | Total elapsed time for all thread read BLOB requests   |
| 516            | 204            | 64 bit integer | 8   | LMSSTMXW | Maximum elapsed time for a thread read BLOB request    |
| 524            | 20C            | 64 bit integer | 8   | LMSSSCWU | Total elapsed time for all SQL read BLOB requests      |
| 532            | 214            | 64 bit integer | 8   | LMSSSMXW | Maximum elapsed time for an SQL read BLOB request      |
| 540            | 21C            | 64 bit integer | 8   | LMSITCUW | Total elapsed time for all thread insert BLOB requests |
| 548            | 224            | 64 bit integer | 8   | LMSITMXW | Maximum elapsed time for a thread insert BLOB request  |
| 556            | 22C            | 64 bit integer | 8   | LMSISCUW | Total elapsed time for all SQL insert BLOB requests    |

Table 119. Db2 statistics record (Q5ST) (continued)

| Offset:<br>Dec | Offset:<br>Hex | Type           | Len | Name     | Description                                            |
|----------------|----------------|----------------|-----|----------|--------------------------------------------------------|
| 564            | 234            | 64 bit integer | 8   | LMSISMXW | Maximum elapsed time for an SQL insert BLOB request    |
| 572            | 23C            | 64 bit integer | 8   | LMSUTCW  | Total elapsed time for all thread update BLOB requests |
| 580            | 244            | 64 bit integer | 8   | LMSUTMXW | Maximum elapsed time for a thread update BLOB request  |
| 588            | 24C            | 64 bit integer | 8   | LMSUSCUW | Total elapsed time for all SQL update BLOB requests    |
| 596            | 254            | 64 bit integer | 8   | LMSUSMXW | Maximum elapsed time for an SQL update BLOB request    |
| 604            | 25C            | 64 bit integer | 8   | LMSDTCW  | Total elapsed time for all thread delete BLOB requests |
| 612            | 264            | 64 bit integer | 8   | LMSDTMXW | Maximum elapsed time for a thread delete BLOB request  |
| 620            | 26C            | 64 bit integer | 8   | LMSDSCW  | Total elapsed time for all SQL delete BLOB requests    |
| 628            | 274            | 64 bit integer | 8   | LMSDSMXW | Maximum elapsed time for an SQL delete BLOB request    |
| 636            | 27C            | 64 bit integer | 8   | LMSLTCW  | Total elapsed time for all thread list BLOB requests   |
| 644            | 284            | 64 bit integer | 8   | LMSLTMXW | Maximum elapsed time for a thread list BLOB request    |
| 652            | 28C            | 64 bit integer | 8   | LMSLSCW  | Total elapsed time for all SQL list BLOB requests      |
| 660            | 294            | 64 bit integer | 8   | LMSLSMXW | Maximum elapsed time for an SQL list BLOB request      |

The data contains counts for each request type that the Db2 resource manager supports. For these request types, maximum and cumulative elapse times are kept for the following:

- The time spent in the Db2 resource manager as a whole (called the thread time).
- The time that was spent performing the RRSF and SQL parts of the request (a subset of the thread time called the SQL time).

Information is also provided for:

- The number of server tasks attached.
- The maximum overall request depth against any of the server tasks.
- The number of times any of the server task requests terminated abnormally.

If the abnormal termination count is not zero, a requeue count is provided indicating the number of queued requests that were requeued to other server tasks as a result of the abnormal termination.

If the average thread time is significantly greater than the average SQL time, this might indicate that thread requests are spending an excessive amount of time waiting for a server task to process the SQL part of the request. If this is the case, examine the DHIGMAX field and, if the value is greater than one, consider increasing the number of Db2 server tasks specified in the QSGDATA parameter of the CSQ6SYSP system parameter macro.

### Coupling facility manager data records

Use this topic as a reference to the format of the coupling facility manager data records.

The format of the coupling facility manager statistics record is described in the following table and in assembler macro thlqual.SCSQMACS(CSQDQEST) and C header file thlqual.SCSQC370(CSQDSMFC). The field names in C are all in lowercase, for example qest, qestid.

If the queue manager was not started as a member of a queue-sharing group, no data is recorded in this record.

Table 120. Coupling facility statistics record (QEST)

| Offset: Dec | Offset: Hex | Type      | Len  | Name            | Description                       |
|-------------|-------------|-----------|------|-----------------|-----------------------------------|
| 0           | 0           | Structure | 4104 | QEST            | CF manager statistics             |
| 0           | 0           | Bitstring | 2    | QESTID          | Control block identifier          |
| 2           | 2           | Integer   | 2    | QESTLL          | Control block length              |
| 4           | 4           | Character | 4    | QESTEYEC        | Control block eye catcher         |
| 8           | 8           | Character | 4096 | QESTZERO        | QEST part cleared on occasion     |
| 8           | 8           | Character | 64   | QESTSTUC (0:63) | Array (one entry per structure)   |
| 8           | 8           | Character | 12   | QESTSTR         | Structure name                    |
| 20          | 14          | Integer   | 4    | QESTSTRN        | Structure number                  |
| 24          | 18          | Integer   | 4    | QESTCSEC        | Number of IXLLSTE calls           |
| 28          | 1C          | Integer   | 4    | QESTCMEC        | Number of IXLLSTM calls           |
| 32          | 20          | Character | 8    | QESTSSTC        | Time spent doing IXLLSTE calls    |
| 40          | 28          | Character | 8    | QESTMSTC        | Time spent doing IXLLSTM calls    |
| 48          | 30          | Integer   | 4    | QESTRSEC        | Number of IXLLSTE redrives        |
| 52          | 34          | Integer   | 4    | QESTRMEC        | Number of IXLLSTM redrives        |
| 56          | 38          | Integer   | 4    | QESTSFUL        | Number of structure fulls         |
| 60          | 3C          | Integer   | 4    | QESTMNUS        | Maximum number of entries in use  |
| 64          | 40          | Integer   | 4    | QESTMLUS        | Maximum number of elements in use |
| 68          | 44          | Character | 4    | *               | Reserved                          |
| 4104        | 1008        | Character | 0    | *               | End of control block              |

The data contains information for each coupling facility list structure, including the CSQ\_ADMIN structure, that the queue manager could connect to during the statistics interval. The information for each structure includes the following:

- The number of and cumulative elapsed times for IXLLSTE and IXLLSTM requests.
- The number of times a request had to be retried because of a timeout.
- The number of times a 'structure full' condition occurred.

## Topic manager data records

Use this topic as a reference to the format of the topic manager data records.

The format of the Topic manager statistics record is described in the following table and in assembler macro thlqual.SCSQMACS(CSQDQTST) and C header file thlqual.SCSQC370(CSQDSMFC). The field names in C are all in lowercase, for example qtst, qtstid.

Table 121. Topic manager statistics record (QTST)

| Offset: Dec | Offset: Hex | Type      | Len | Name           | Description                                            |
|-------------|-------------|-----------|-----|----------------|--------------------------------------------------------|
| 0           | 0           | Structure | 96  | QTST           | Topic manager statistics                               |
| 0           | 0           | Bitstring | 2   | QTSTID         | Control block identifier                               |
| 2           | 2           | Integer   | 2   | QTSTLL         | Control block length                                   |
| 4           | 4           | Character | 4   | TESTEYEC       | Control block eye catcher                              |
| 8           | 8           | Character | 88  | QTSTZERO       | QTST part cleared on occasion                          |
| 8           | 8           | Integer   | 4   | QTSTSTOT       | Total subscription requests                            |
| 12          | 0C          | Integer   | 4   | QTSTSDUR       | Durable subscription requests                          |
| 16          | 10          | Integer   | 4   | QTSTSHIG (1:3) | Subscription high water mark array (API, ADMIN, PROXY) |
| 28          | 1C          | Integer   | 4   | QTSTSLOW (1:3) | Subscription low water mark array (API, ADMIN, PROXY)  |
| 40          | 28          | Integer   | 4   | QTSTSEXP       | Subscriptions expired                                  |
| 44          | 2C          | Integer   | 4   | QTSTTMSG       | Total messages put to Sub queue                        |
| 48          | 30          | Integer   | 4   | QTSTSPHW       | Single publish subscriber high water mark              |
| 52          | 34          | Integer   | 4   | QTSTPTOT (1:3) | Total Publication requests (API, ADMIN, PROXY)         |
| 64          | 40          | Integer   | 4   | QTSTPTHI       | Total publish high water mark                          |
| 68          | 44          | Integer   | 4   | QTSTPTLO       | Total publish low water mark                           |
| 72          | 48          | Integer   | 4   | QTSTPNOS       | Count of publishes to no subscriber                    |
| 76          | 4C          | Integer   | 4   | *              | Reserved                                               |
| 80          | 50          | Bitstring | 8   | QTSTETHW       | Elapse time HW on publish                              |
| 88          | 58          | Bitstring | 8   | QTSTETTO       | Elapse time total on publish                           |

## Coupling facility manager SMDS data records

Use this topic as a reference to the format of the coupling facility manager shared message data set (SMDS) data records.

The format of the coupling facility manager shared message data set (SMDS) statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQESD), C header file thlqual.SCSQC370(CSQDSMFC) and in WebSphere MQ SupportPac [MP1B: WebSphere MQ for z/OS V7.0 - Interpreting accounting and statistics data](#)

The statistics provide information about the utilization of the owned shared message data set, I/O activity for the group of shared message data sets, and SMDS buffer utilization.

If the queue manager was not started as a member of a queue-sharing group, no data is recorded in this record.

## Interpreting WebSphere MQ accounting data

WebSphere MQ accounting data is written as SMF type 116 records. Use this topic as a reference to the different types of accounting data records.

WebSphere MQ accounting information can be collected for three subtypes:

- 0 Message manager accounting records (how much processor time was spent processing WebSphere MQ API calls and the number of **MQPUT** and **MQGET** calls). This information is produced when a named task disconnects from WebSphere MQ, and so the information contained within the record might cover many hours.
- 1 Accounting data for each task, at thread and queue level.
- 2 Additional queue-level accounting data (if the task used more queues than could fit in the subtype 1 record).

Subtype 0 is produced with trace class 1; subtypes 1 and 2 are produced with trace class 3.

**Related concepts:**

“Message manager data records” on page 1145

“Thread-level and queue-level data records” on page 1146

**Related reference:**

“Layout of an SMF type 116 record”

“Processing type 116 SMF records” on page 1143

“Common WebSphere MQ SMF header” on page 1143

“Thread cross reference data” on page 1145

“Sample subtype zero accounting record” on page 1145

**Layout of an SMF type 116 record**

Use this topic as a reference to the format of an SMF type record.

The standard layout for SMF records involves three parts:

**SMF header**

Provides format, identification, and time and date information about the record itself.

**Self-defining section**

Defines the location and size of the individual data records within the SMF record.

**Data records**

The actual data from WebSphere MQ that you want to analyze.

For more information about SMF record formats, see the *MVS System Management Facilities (SMF)* manual.

**The SMF header**

Table 122 shows the format of SMF record header (SM116).

*Table 122. SMF record header description*

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description                                                                                       | Example  |
|----------------|----------------|-----------|-----|----------|---------------------------------------------------------------------------------------------------|----------|
| 0              | 0              | Structure | 28  | SM116    | SMF record header.                                                                                |          |
| 0              | 0              | Integer   | 2   | SM116LEN | SMF record length.                                                                                | 01A4     |
| 2              | 2              |           | 2   |          | Reserved.                                                                                         |          |
| 4              | 4              | Integer   | 1   | SM116FLG | System indicator.                                                                                 | 5E       |
| 5              | 5              | Integer   | 1   | SM116RTY | Record type. The SMF record type, for WebSphere MQ accounting records this is always 116 (X'74'). | 74       |
| 6              | 6              | Integer   | 4   | SM116TME | Time when SMF moved record.                                                                       | 00356124 |
| 10             | A              | Integer   | 4   | SM116DTE | Date when SMF moved record.                                                                       | 0100223F |

Table 122. SMF record header description (continued)

| Offset:<br>Dec                                                                                           | Offset:<br>Hex | Type      | Len | Name     | Description                                                                        | Example         |
|----------------------------------------------------------------------------------------------------------|----------------|-----------|-----|----------|------------------------------------------------------------------------------------|-----------------|
| 14                                                                                                       | E              | Character | 4   | SM116SID | z/OS subsystem ID. Defines the z/OS subsystem on which the records were collected. | D4E5F4F1 (MV41) |
| 18                                                                                                       | 12             | Character | 4   | SM116SSI | WebSphere MQ subsystem ID.                                                         | D4D8F0F7 (MQ07) |
| 22                                                                                                       | 16             | Integer   | 2   | SM116STF | Record subtype.                                                                    | 0000            |
| 24                                                                                                       | 18             | Character | 3   | SM116REL | WebSphere MQ version.                                                              | F6F0F0 (600)    |
| 27                                                                                                       | 1B             |           | 1   |          | Reserved.                                                                          |                 |
| 28                                                                                                       | 1C             | Character | 0   | SM116END | End of SMF header and start of self-defining section.                              |                 |
| <p><b>Note:</b> The (hexadecimal) values in the right-hand column relate to Figure 132 on page 1143.</p> |                |           |     |          |                                                                                    |                 |

## Self-defining sections

A self-defining section of an SMF record tells you where to find an accounting record, how long it is, and how many times that type of record is repeated (with different values). The self-defining sections follow the header, at a fixed offset from the start of the SMF record.

Each self-defining section points to accounting related data. Table 123 summarizes the offsets from the start of the SMF record header.

Table 123. Offsets to self-defining sections. Offsets are from the start of the SMF record and are fixed for each type of accounting source.

| Record subtype<br>(SMF116STF) | Source of accounting data    | Offset of self-defining section |       | See...                                                                                                                                                                     |
|-------------------------------|------------------------------|---------------------------------|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                               |                              | Dec                             | Hex   |                                                                                                                                                                            |
| All                           | Common header                | 28                              | X'1C' | "Common WebSphere MQ SMF header" on page 1143                                                                                                                              |
| 0                             | Message manager              | 44                              | X'2C' | "Message manager data records" on page 1145                                                                                                                                |
| 1                             | Thread identification record | 36                              | X'24' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |
| 1                             | Thread-level accounting      | 44                              | X'2C' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |
| 1                             | Queue-level accounting       | 52                              | X'34' | "Thread-level and queue-level data records" on page 1146. This section is present only if the WTASWQCT field in the task-related information (WTAS) structure is non-zero. |
| 2                             | Thread identification record | 36                              | X'24' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |

Table 123. Offsets to self-defining sections (continued). Offsets are from the start of the SMF record and are fixed for each type of accounting source.

| Record subtype (SMF116STF) | Source of accounting data | Offset of self-defining section |       | See...                                                   |
|----------------------------|---------------------------|---------------------------------|-------|----------------------------------------------------------|
|                            |                           | Dec                             | Hex   |                                                          |
| 2                          | Queue-level accounting    | 44                              | X'2C' | "Thread-level and queue-level data records" on page 1146 |

**Note:** Other self-defining sections refer to data for IBM use only.

Each self-defining section is two fullwords long and has this format:

sssssss1111nnnn

where:

**sssssss**

Fullword containing the offset from start of the SMF record.

**1111** Halfword giving the length of this data record.

**nnnn** Halfword giving the number of data records in this SMF record.

Figure 131 shows an example of part of an SMF type 116 record. The numbers in the left-hand column represent the offset, in hexadecimal, from the start of the record. Each line corresponds to sixteen bytes of data, where each byte is two hexadecimal characters, for example 0C. The characters in the right-hand column represent the printable characters for each byte. Non-printable characters are shown by a period (.) character.

In this example, alternate fields in the SMF header are underlined to help you to see them; refer to Table 122 on page 1138 to identify them. The self defining section for one of the message manager accounting data records (at the offset given in Table 125 on page 1141) is shown in **bold**.

|        |                 |          |          |                 |                    |
|--------|-----------------|----------|----------|-----------------|--------------------|
| 000000 | 01A40000        | 5E740035 | 61240100 | 223FD4E5        | *....;.../.....MV* |
| 000000 | F4F1D4D8        | F0F70000 | F6F0F000 | 00000134        | *41MQ07..600.....* |
| 000000 | 00700001        | 00000054 | 00B00001 | <b>00000104</b> | *.....*            |
| 000000 | <b>00300001</b> | 00000000 | 00000000 | 00000000        | *.....*            |
| 000000 | 00000000        | 00000000 | 00000000 | 00000000        | *.....*            |

Figure 131. Part of an SMF record 116 showing the header and self-defining sections

The self-defining section for the type of message manager accounting data is located at offset X'2C' from the start of the SMF record and contains this information:

- The offset of the message manager accounting data is located X'00000104' bytes from the start of the SMF record.
- This message manager record is X'0030' bytes long.
- There is one record (X'0001').

**Note:** Always use offsets in the self-defining sections to locate the accounting records.

**Related reference:**

“The SMF header”

“Self-defining sections”

**The SMF header:**

Table 124 shows the format of SMF record header (SM116).

*Table 124. SMF record header description*

| Offset:<br>Dec | Offset:<br>Hex | Type      | Len | Name     | Description                                                                                       | Example         |
|----------------|----------------|-----------|-----|----------|---------------------------------------------------------------------------------------------------|-----------------|
| 0              | 0              | Structure | 28  | SM116    | SMF record header.                                                                                |                 |
| 0              | 0              | Integer   | 2   | SM116LEN | SMF record length.                                                                                | 01A4            |
| 2              | 2              |           | 2   |          | Reserved.                                                                                         |                 |
| 4              | 4              | Integer   | 1   | SM116FLG | System indicator.                                                                                 | 5E              |
| 5              | 5              | Integer   | 1   | SM116RTY | Record type. The SMF record type, for WebSphere MQ accounting records this is always 116 (X'74'). | 74              |
| 6              | 6              | Integer   | 4   | SM116TME | Time when SMF moved record.                                                                       | 00356124        |
| 10             | A              | Integer   | 4   | SM116DTE | Date when SMF moved record.                                                                       | 0100223F        |
| 14             | E              | Character | 4   | SM116SID | z/OS subsystem ID. Defines the z/OS subsystem on which the records were collected.                | D4E5F4F1 (MV41) |
| 18             | 12             | Character | 4   | SM116SSI | WebSphere MQ subsystem ID.                                                                        | D4D8F0F7 (MQ07) |
| 22             | 16             | Integer   | 2   | SM116STF | Record subtype.                                                                                   | 0000            |
| 24             | 18             | Character | 3   | SM116REL | WebSphere MQ version.                                                                             | F6F0F0 (600)    |
| 27             | 1B             |           | 1   |          | Reserved.                                                                                         |                 |
| 28             | 1C             | Character | 0   | SM116END | End of SMF header and start of self-defining section.                                             |                 |

**Note:** The (hexadecimal) values in the right-hand column relate to Figure 132 on page 1143.

**Self-defining sections:**

A self-defining section of an SMF record tells you where to find an accounting record, how long it is, and how many times that type of record is repeated (with different values). The self-defining sections follow the header, at a fixed offset from the start of the SMF record.

Each self-defining section points to accounting related data. Table 125 summarizes the offsets from the start of the SMF record header.

*Table 125. Offsets to self-defining sections.* Offsets are from the start of the SMF record and are fixed for each type of accounting source.

| Record subtype<br>(SMF116STF) | Source of accounting data | Offset of self-defining section |       | See...                                        |
|-------------------------------|---------------------------|---------------------------------|-------|-----------------------------------------------|
|                               |                           | Dec                             | Hex   |                                               |
| All                           | Common header             | 28                              | X'1C' | “Common WebSphere MQ SMF header” on page 1143 |
| 0                             | Message manager           | 44                              | X'2C' | “Message manager data records” on page 1145   |

Table 125. Offsets to self-defining sections (continued). Offsets are from the start of the SMF record and are fixed for each type of accounting source.

| Record subtype (SMF116STF) | Source of accounting data    | Offset of self-defining section |       | See...                                                                                                                                                                     |
|----------------------------|------------------------------|---------------------------------|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                            |                              | Dec                             | Hex   |                                                                                                                                                                            |
| 1                          | Thread identification record | 36                              | X'24' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |
| 1                          | Thread-level accounting      | 44                              | X'2C' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |
| 1                          | Queue-level accounting       | 52                              | X'34' | "Thread-level and queue-level data records" on page 1146. This section is present only if the WTASWQCT field in the task-related information (WTAS) structure is non-zero. |
| 2                          | Thread identification record | 36                              | X'24' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |
| 2                          | Queue-level accounting       | 44                              | X'2C' | "Thread-level and queue-level data records" on page 1146                                                                                                                   |

**Note:** Other self-defining sections refer to data for IBM use only.

Each self-defining section is two fullwords long and has this format:

```
sssssssl1111nnnn
```

where:

**ssssssss**

Fullword containing the offset from start of the SMF record.

**l111** Halfword giving the length of this data record.

**nnnn** Halfword giving the number of data records in this SMF record.

Figure 132 on page 1143 shows an example of part of an SMF type 116 record. The numbers in the left-hand column represent the offset, in hexadecimal, from the start of the record. Each line corresponds to sixteen bytes of data, where each byte is two hexadecimal characters, for example 0C. The characters in the right-hand column represent the printable characters for each byte. Non-printable characters are shown by a period (.) character.

In this example, alternate fields in the SMF header are underlined to help you to see them; refer to Table 124 on page 1141 to identify them. The self defining section for one of the message manager accounting data records (at the offset given in Table 125 on page 1141) is shown in **bold**.

```

000000 01A40000 5E740035 61240100 223FD4E5 *....;.../.....MV*
000000 F4F1D4D8 F0F70000 F6F0F000 00000134 *41MQ07..600....*
000000 00700001 00000054 00B00001 00000104 *.....*
000000 00300001 00000000 00000000 00000000 *.....*
000000 00000000 00000000 00000000 00000000 *.....*

```

Figure 132. Part of an SMF record 116 showing the header and self-defining sections

The self-defining section for the type of message manager accounting data is located at offset X'2C' from the start of the SMF record and contains this information:

- The offset of the message manager accounting data is located X'00000104' bytes from the start of the SMF record.
- This message manager record is X'0030' bytes long.
- There is one record (X'0001').

**Note:** Always use offsets in the self-defining sections to locate the accounting records.

## Processing type 116 SMF records

Use this topic as a reference to the format of the processing type accounting record.

Any accounting data you collect from SMF must be processed to extract useful information. When you process the data, verify that the records are from WebSphere MQ and that they are the records you are expecting.

Validate the value of the following fields:

- SM116RTY, the SMF record number = X'74' (116)
- SM116STF, the record subtype, must be 0000, 0001, or 0002

Reading from the active SMF data sets is not supported. You must use the SMF program IFASMFDP to dump SMF records to a sequential data set so that they can be processed. For more information see "Using System Management Facility" on page 1119.

There is a C sample program called CSQ4SMFD which prints the contents of SMF type 115 and 116 records. The program is provided as source in thlqual.SCSQC37S and in executable format in thlqual.SCSQLOAD. Sample JCL is provided in thlqual.SCSQPROC.

You need to update the SMFIN DD card with the name of the SMF dataset. Use the z/OS command '/D SMF' to show the name of the dataset, and you need to update the DUMPDD DD card with the name for the output dataset.

You also need to specify the START and END times that you require.

The following sample JCL extracts SMF records:

```

//SMFDUMP EXEC PGM=IFASMFDP,REGION=0M
//SYSPRINT DD SYSOUT=
//SMFIN DD DSN=xxxxxx.MANA,DISP=SHR
//SMFOUT DD DSN=xxxxxx.SMFOUT,SPACE=(CYL,(1,1)),DISP=(NEW,CATLG)
//SYSIN DD *
 INDD(SMFIN,OPTIONS(DUMP))
 OUTDD(SMFOUT,TYPE(116))
 OUTDD(SMFOUT,TYPE(115))
 START(1159) END(1210)
/*

```

## Common WebSphere MQ SMF header

Use this topic as a reference to the common WebSphere MQ SMF header type accounting record.

The format of this record is described in Table 126 and in assembler macros `thlqual.SCSQMACS(CSQDQWHS)` and `thlqual.SCSQMACS(CSQDQWHC)`, and C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `qwhs`, `qwhsnsda`.

Details of the structures and fields can be found in WebSphere MQ supportpac [MP1B: WebSphere MQ for z/OS V7.1 - Interpreting accounting and statistics data](#).

The QWHS data includes the subsystem name. For subtype 1 records, it also shows whether there are queue-level accounting records present. If the QWHNSDA field is 3 or less, there are not, and the corresponding self-defining section (at offset X'34') is not set.

The QWHC data gives you information about the user (for example, the user ID (QWHCAID) and the type of application (QWHCATYP)). The QWHC section is completed only for subtype 0 records. The equivalent information is present in the thread identification record for subtype 1 and 2 records.

*Table 126. Structure of the common WebSphere MQ SMF header record QWHS*

| Offset:<br>Dec | Offset:<br>Hex | Type      | Length | Name     | Description                                                                                                                                    |
|----------------|----------------|-----------|--------|----------|------------------------------------------------------------------------------------------------------------------------------------------------|
| 0              | 0              | Structure | 128    | QWHS     |                                                                                                                                                |
| 0              | 0              |           | 6      |          | Reserved                                                                                                                                       |
| 6              | 6              | Character | 1      | QWHNSDA  | Number of self defining sections in the SMF records                                                                                            |
| 7              | 7              |           | 5      |          | Reserved                                                                                                                                       |
| 12             | C              | Character | 4      | QWHSSID  | Subsystem name                                                                                                                                 |
| 16             | 10             |           | 24     |          | Reserved                                                                                                                                       |
| 40             | 28             | Character | 8      | QWHCAID  | User ID associated with the z/OS job                                                                                                           |
| 48             | 30             | Character | 12     | QWHCCV   | Thread cross reference                                                                                                                         |
| 60             | 3C             | Character | 8      | QWHCCN   | Connection name                                                                                                                                |
| 68             | 44             |           | 8      |          | Reserved                                                                                                                                       |
| 76             | 4C             | Character | 8      | QWHCOPID | User ID associated with the transaction                                                                                                        |
| 84             | 54             | Integer   | 4      | QWHCATYP | Type of connecting system (1=CICS, 2=Batch or TSO, 3=IMS control region, 4=IMS MPP or BMP, 5=Command server, 6=Channel initiator, 7=RRS Batch) |
| 88             | 58             | Character | 22     | QWHCTOKN | Accounting token set to the z/OS accounting information for the user                                                                           |
| 110            | 6E             | Character | 16     | QWHCNID  | Network identifier                                                                                                                             |
| 126            | 7E             |           | 2      |          | Reserved                                                                                                                                       |

## Combining CICS and WebSphere MQ performance data

Use this topic as a reference to the combination of WebSphere MQ and CICS performance data.

The common WebSphere MQ SMF header type accounting record section, QWHCTOKN, is used to correlate CICS type 110 SMF records with WebSphere MQ type 116 SMF records.

CICS generates an LU6.2 unit-of-work token, for each CICS task. The token is used to generate an accounting token that is written to QWHCTOKN in the correlation header of subtype zero records.

Details are also written to the WTIDACCT section in subtype 1 and 2 records. The accounting token enables correlation between CICS and WebSphere MQ performance data for a transaction.

## Thread cross reference data

Use this topic as a reference to the format of the thread cross reference type accounting record.

The interpretation of the data in the thread cross reference (QWHCCV) field varies. This depends on what the data relates to:

- CICS connections (QWHCATYP=1) – see Table 127
- IMS connections (QWHCATYP=3 or 4) – see Table 128
- Batch connections (QWHCATYP=2 or 7) – this field consists of binary zeros
- Others – no meaningful data

Table 127. Structure of the thread cross reference for a CICS system

| Offset: Dec | Offset: Hex | Type      | Length | Description            |
|-------------|-------------|-----------|--------|------------------------|
| 48          | 30          | Character | 4      | CICS thread number.    |
| 52          | 34          | Character | 4      | CICS transaction name. |
| 56          | 38          | Integer   | 4      | CICS task number.      |

Some entries contain blank characters. These apply to the task, rather than to a specific transaction.

Table 128. Structure of the thread cross reference for an IMS system

| Offset: Dec | Offset: Hex | Type      | Length | Description                                                |
|-------------|-------------|-----------|--------|------------------------------------------------------------|
| 48          | 30          | Character | 4      | IMS partition specification table (PST) region identifier. |
| 52          | 34          | Character | 8      | IMS program specification block (PSB) name.                |

## Message manager data records

Use this topic as a reference to the format of the message manager accounting records.

The message manager is the component of WebSphere MQ that processes all API requests. The format of the message manager accounting records is described in assembler macro `thlqual.SCSQMACS(CSQDQMAC)`.

The QMAC data gives you information about the processor time spent processing WebSphere MQ calls, and counts of the number of **MQPUT** and **MQGET** requests for messages of different sizes.

**Note:** A single IMS application might write two SMF records. In this case, add the figures from both records to provide the correct totals for the IMS application.

## Records containing zero processor time

Records are sometimes produced that contain zero processor time in the QMACCPUT field. These records occur when long running tasks identified to WebSphere MQ either terminate or are prompted to output accounting records by accounting trace being stopped. Such tasks exist in the CICS adapter and in the channel initiator (for distributed queuing). The number of these tasks with zero processor time depends upon how much activity there has been in the system:

- For the CICS adapter, this can result in up to nine records with zero processor time.
- For the channel initiator, the number of records with zero processor time can be up to the sum of Adapters + Dispatchers + 6, as defined in the queue manager attributes.

These records reflect the amount of work done under the task, and can be ignored.

## Sample subtype zero accounting record

Use this topic as a reference to the format of the subtype zero accounting records.

Figure 133 shows a type 116, subtype zero SMF record. In this figure, the SMF record header and the QMAC accounting data record are underlined. The self-defining sections are in bold.

|        |                 |                 |                 |                 |                    |
|--------|-----------------|-----------------|-----------------|-----------------|--------------------|
| 000000 | 01A40000        | 5E740035        | 61240100        | 223FD4E5        | *....;.../.....MV* |
| 000010 | <u>F4F1D4D8</u> | <u>F0F70000</u> | <u>F6F0F000</u> | <u>00000134</u> | *41MQ07..600....*  |
| 000020 | <b>00700001</b> | 00000054        | 00B00001        | <b>00000104</b> | *.....*            |
| 000030 | <b>00300001</b> | 00000000        | 00000000        | 00000000        | *.....*            |
| 000040 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 000050 | 00000000        | B478AB43        | 9C6C2280        | B478AB47        | *.....%.....*      |
| 000060 | 9DB47E02        | 00000000        | 04C0F631        | 00000001        | *..=.....{6.....*  |
| 000070 | 9880E72D        | 00000000        | 014D9540        | 00000000        | *..X.....(. ....*  |
| 000080 | 08480C80        | 00000010        | 40404040        | 40404040        | *.....*            |
| 000090 | 00000000        | 00000000        | 00000051        | 00000000        | *.....*            |
| 0000A0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 0000B0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 0000C0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 0000D0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 0000E0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 0000F0 | 00000000        | 00000000        | 00000000        | 00000000        | *.....*            |
| 000100 | 00000000        | D4140030        | D8D4C1C3        | 00000000        | *...M...QMAC...*   |
| 000110 | <u>689C738D</u> | <u>00000050</u> | <u>00000000</u> | <u>00000050</u> | *.....&.....&*     |
| 000120 | <u>0000000A</u> | <u>00000000</u> | <u>00000000</u> | <u>00000000</u> | *.....*            |
| 000130 | <u>00000000</u> | <u>0024011A</u> | <u>00030710</u> | <u>02DAACF0</u> | *.....0*           |

Figure 133. Example SMF type 116, subtype zero record

## Thread-level and queue-level data records

Use this topic as a reference to the format of the thread-level and queue-level accounting records.

Thread level accounting records are collected for each task using WebSphere MQ. For each task, a thread-level accounting data record is written to the SMF when the task finishes. For a long running task, data is also written at the statistics interval set by the STATIME parameter of the CSQ6SYSP system parameter macro (or by the system SMF statistics broadcast), provided that the task was running the previous time statistics were gathered. In addition, accounting information is gathered about each queue that the task opens. A queue-level accounting record is written for each queue that the task has used since the thread-level accounting record was last written.

Thread-level and queue-level accounting records are produced if you specify class 3 when you start the accounting trace.

The thread level accounting information is written to an SMF type 116, subtype 1 record, and is followed by queue-level records. If the task opened many queues, further queue information is written to one or more SMF type 116 subtype 2 records. A thread identification control block is included in each subtype 1 and 2 record to enable you to relate each record to the correct task. Typically, the maximum number of queue-level records in each SMF record is about 45.

The format of the thread-level accounting record is described in assembler macro thlqual.SCSQMACS(CSQDWTAS). The format of the queue-level accounting record is described in assembler macro thlqual.SCSQMACS(CSQDWQ). The format of the thread identification record is described in assembler macro thlqual.SCSQMACS(CSQDWTID). All these records are also described in C header file thlqual.SCSQC370(CSQDSMFC). The field names in C are all in lowercase, for example wtas, wtasshex.

**Related concepts:**

“Meaning of the channel names”

**Related reference:**

“Sample subtype 1 and subtype 2 records”

**Meaning of the channel names:**

Use this topic as a reference to the meaning of channel names..

The channel name in the WTID is constructed as shown in the following example. In this example a sender channel exists from queue manager QM1 to queue manager QM2.

Table 129. Meaning of channel names

| Field name                                                               | Meaning                                       | Example        |
|--------------------------------------------------------------------------|-----------------------------------------------|----------------|
| For queue manager QM1 the sender channel has the following fields set:   |                                               |                |
| WTIDCCN                                                                  | The job name                                  | QM1CHIN        |
| WTIDCHL                                                                  | The channel name                              | QM1.QM2        |
| WTIDCHLC                                                                 | This is defined in the CONNAME of the channel | WINMVS2B(2162) |
| For queue manager QM2 the receiver channel has the following fields set: |                                               |                |
| WTIDCCN                                                                  | The job name                                  | QM2CHIN        |
| WTIDCHL                                                                  | The channel name                              | QM1.QM2        |
| WTIDCHLC                                                                 | Where the channel came from                   | 9.20.101.14    |

**Sample subtype 1 and subtype 2 records:**

Use this topic as a reference to the format of the subtype 1 and subtype 2 accounting records.

Figure 134 and Figure 135 on page 1148 show examples of SMF type 116, subtype 1 and subtype 2 records. These two accounting records were created for a batch job that opened 80 queues. Because many queues were opened, a subtype 2 record was required to contain all the information produced.

```

000000 703C0000 5E74002D 983B0100 229FD4E5 *....;.....MV*
000010 F4F1D4D8 F0F70001 F6F0F000 00006FCC *41MQ07..600...?.*
000020 00700001 0000003C 00D00001 0000010C *.....}.....*
000030 02C00001 000003CC 02400030 F70000D0 *.{..... ..7..}*
000040 E6E3C9C4 00000000 00000000 00000040 *WTID..... *
.
.
.
000100 00000000 00000000 7F4A4BB8 F70102C0 *....."...7..{*
000110 E6E3C1E2 B4802373 0BF07885 7F4AE718 *WTAS.....0.."X.*

```

Figure 134. Example SMF type 116, subtype 1 record. This record contains a CSQDWTID control block, the CSQDWTAS control block, and the first set of CSQDWQST control blocks.

The first self-defining section starts at X'24' and is **bold** in the example; X'0000003C' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'0000010C' is the offset to the WTAS data record, X'02C0' is the length of the WTAS record, and X'0001' is the number of WTAS records.

The third self-defining section starts at X'34' and is **bold** in the example; X'000003CC' is the offset to the first WQST data record, X'0240' is the length of the WQST record, and X'0030' is the number of WQST records.

Figure 135 shows an example of an SMF type 116, subtype 2 record.

```

000000 49740000 5E74002D 983B0100 229FD4E5 *....;.....MV*
000010 F4F1D4D8 F0F70002 F6F0F000 00004904 *41MQ07..600.....*
000020 00700001 00000034 00D00001 00000104 *.....}.....*
000030 02400020 F70000D0 E6E3C9C4 00000002 *..7..}WTID....*
.
.
.
000100 7F4A4BB8 F7020240 E6D8E2E3 00000001 *"...7.. WQST....*
```

Figure 135. Example SMF type 116, subtype 2 record. This record contains a CSQDWTID control block and the remaining CSQDWQST control blocks.

The first self-defining section starts at X'24' and is **bold** in the example; X'00000034' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'00000104' is the offset to the first WQST data record, X'0240' is the length of the WQST record, and X'0020' is the number of WQST records.

Figure 136 shows an example of an SMF type 116, subtype 1 record where no queues have been opened and there are consequently no self-defining sections for WQST records..

```

000000 5E740039 4E9B0104 344FD4E5 *|MV*
000010 F4F1D4D8 F0F70001 F6F0F000 000003DC *41MQ07..600.....*
000020 00800001 00000034 00D00001 00000104 *.....*
000030 02D80001 F70000D0 E6E3C9C4 00000002 *.Q..7...WTID....*
000040 C1F8C5C1 C4C5D740 C1F8C5C1 C4C54040 *A8EADEP A8EADE *
000050 40404040 40404040 00000000 00000000 **
000060 40404040 40404040 4040 * *
```

Figure 136. Example SMF type 116, subtype 1 record with no WQST data records

The first self-defining section starts at X'24' and is **bold** in the example; X'00000034' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'0000010C' is the offset to the WTAS data record, X'02D8' is the length of the WTAS record, and X'0001' is the number of WTAS records.

There is no self-defining section describing a WQST data record, equivalent to the third self-defining section in Figure 134 on page 1147.

---

## Troubleshooting and support

If you are having problems with your queue manager network or IBM WebSphere MQ applications, use the techniques described to help you diagnose and solve the problems.

For an introduction to troubleshooting and support, see “Troubleshooting overview.”

There are some initial checks that you can make for your platform to help determine the causes of some common problems. See the appropriate topic for your platform:

- “Making initial checks on Windows, UNIX and Linux systems” on page 1151
- “Making initial checks on IBM i” on page 1161
- “Making initial checks on z/OS” on page 1169

For information about solving problems, see “Dealing with problems” on page 1186.

For information about solving problems for IBM WebSphere MQ Telemetry, see “Troubleshooting for IBM WebSphere MQ Telemetry” on page 1206.

For information about solving problems when you are using channel authentication records, see “Troubleshooting channel authentication records” on page 1222.

Information that is produced by IBM WebSphere MQ can help you to find and resolve problems. For more information, see the following topics:

- “Using logs” on page 1226
- “Using trace” on page 1234
- “Problem determination on z/OS” on page 1265
- “First Failure Support Technology (FFST)” on page 1315

For information about recovering after a problem, see “Recovering after failure” on page 1355.

You can also read the general troubleshooting guidance in the following topics:

- “IBM Support Assistant (ISA)” on page 1324
- “Searching knowledge bases” on page 1326
- “Contacting IBM Software Support” on page 1346
- “Getting product fixes” on page 1355

If a IBM WebSphere MQ component or command has returned an error, and you want further information about a message written to the screen or the log, you can browse for details of the message, see “Reason codes” on page 1379.

### Related concepts:



Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

---

## Troubleshooting overview

Troubleshooting is the process of finding and eliminating the cause of a problem. Whenever you have a problem with your IBM software, the troubleshooting process begins as soon as you ask yourself “what happened?”

A basic troubleshooting strategy at a high level involves:

1. "Recording the symptoms of the problem"
2. "Re-creating the problem"
3. "Eliminating possible causes"

## Recording the symptoms of the problem

Depending on the type of problem that you have, whether it be with your application, your server, or your tools, you might receive a message that indicates that something is wrong. Always record the error message that you see. As simple as this sounds, error messages sometimes contain codes that might make more sense as you investigate your problem further. You might also receive multiple error messages that look similar but have subtle differences. By recording the details of each one, you can learn more about where your problem exists.

Sources of error messages:

- Problems view
- Local error log
- Eclipse log
- User trace
- Service trace
- Error dialog boxes

## Re-creating the problem

Think back to what steps you were doing that led to the problem. Try those steps again to see if you can easily re-create the problem. If you have a consistently repeatable test case, it is easier to determine what solutions are necessary.

- How did you first notice the problem?
- Did you do anything different that made you notice the problem?
- Is the process that is causing the problem a new procedure, or has it worked successfully before?
- If this process worked before, what has changed? (The change can refer to any type of change that is made to the system, ranging from adding new hardware or software, to reconfiguring existing software.)
- What was the first symptom of the problem that you witnessed? Were there other symptoms occurring around the same time?
- Does the same problem occur elsewhere? Is only one machine experiencing the problem or are multiple machines experiencing the same problem?
- What messages are being generated that might indicate what the problem is?

You can find more information about these types of question in "Making initial checks on Windows, UNIX and Linux systems" on page 1151.

## Eliminating possible causes

Narrow the scope of your problem by eliminating components that are not causing the problem. By using a process of elimination, you can simplify your problem and avoid wasting time in areas that are not responsible. Consult the information in this product and other available resources to help you with your elimination process.

- Has anyone else experienced this problem? See: "Searching knowledge bases" on page 1326.
- Is there a fix you can download? See: "Getting product fixes" on page 1355.

---

## Making initial checks on Windows, UNIX and Linux systems

Before you start problem determination in detail, consider whether there is an obvious cause of the problem, or an area of investigation that is likely to give useful results. This approach to diagnosis can often save a lot of work by highlighting a simple error, or by narrowing down the range of possibilities.

The cause of your problem could be in:

- IBM WebSphere MQ
- The network
- The application
- Other applications that you have configured to work with IBM WebSphere MQ

This section contains a list of questions to consider. As you go through the list, make a note of anything that might be relevant to the problem. Even if your observations do not suggest a cause straight away, they might be useful later if you have to carry out a systematic problem determination exercise.

- “Has IBM WebSphere MQ run successfully before?” on page 1152
- “Have any changes been made since the last successful run?” on page 1153
- “Are there any error messages or return codes to explain the problem?” on page 1153
- “Can you reproduce the problem?” on page 1154
- “Are you receiving an error code when creating or starting a queue manager? (Windows only)” on page 1154
- “Does the problem affect only remote queues?” on page 1154
- “Have you obtained incorrect output?” on page 1155
- “Are some of your queues failing?” on page 1157
- “Have you failed to receive a response from a PCF command?” on page 1157
- “Has the application run successfully before?” on page 1158
- “Is your application or system running slowly?” on page 1160
- “Does the problem affect specific parts of the network?” on page 1160
- “Does the problem occur at specific times of the day?” on page 1160
- “Is the problem intermittent?” on page 1160

See the following sections for some additional tips for problem determination for system administrators and application developers.

### Tips for system administrators

- Check the error logs for messages for your operating system:
  - “Error logs on Windows, UNIX and Linux systems” on page 1227
  - “Error logs on IBM i” on page 1231
  - “Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272
- Check the contents of `qm.ini` for any configuration changes or errors. For more information on changing configuration information, see:
  -  UNIX and Linux (*WebSphere MQ V7.1 Installing Guide*)
  -  IBM i (*WebSphere MQ V7.1 Installing Guide*)
  -  z/OS (*WebSphere MQ V7.1 Installing Guide*)
- If your application development teams are reporting something unexpected, you use trace to investigate the problems. For information about using trace, see “Using trace” on page 1234.

## Tips for application developers

- Check the return codes from the MQI calls in your applications. For a list of reason codes, see “API reason codes” on page 1380. Use the information provided in the return code to determine the cause of the problem. Follow the steps in the Programmer response sections of the reason code to resolve the problem.
- If you are unsure whether your application is working as expected, for example, you are not unsure of the parameters being passed into the MQI or out of the MQI, you can use trace to collect information about all the inputs and outputs of your MQI calls. For more information about using trace, see “Using trace” on page 1234.
- For more information about handling errors in MQI applications, see  Handling program errors (*WebSphere MQ V7.1 Programming Guide*).

### Related concepts:

“Troubleshooting and support” on page 1149  
“Making initial checks on z/OS” on page 1169  
“Making initial checks on IBM i” on page 1161  
“Dealing with problems” on page 1186  
“IBM Support Assistant (ISA)” on page 1324  
“Reason codes” on page 1379

 Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

### Related tasks:

“Searching knowledge bases” on page 1326  
“Contacting IBM Software Support” on page 1346

### Related reference:

“PCF reason codes” on page 1590

## Has IBM WebSphere MQ run successfully before?

If IBM WebSphere MQ has not run successfully before, it is likely that you have not yet set it up correctly. See  Installing IBM WebSphere MQ (*WebSphere MQ V7.1 Installing Guide*) and select the platform, or platforms, that your enterprise uses to check that you have installed the product correctly.

To run the verification procedure, see:

-  Verifying a server installation (*WebSphere MQ V7.1 Installing Guide*)
-  Verifying a client installation (*WebSphere MQ V7.1 Installing Guide*)

Also look at  Configuring (*WebSphere MQ V7.1 Installing Guide*) for information about post-installation configuration of IBM WebSphere MQ.

## Have any changes been made since the last successful run?

Changes that have been made to your IBM WebSphere MQ configuration, maintenance updates, or changes to other programs that interact with IBM WebSphere MQ could be the cause of your problem.

When you are considering changes that might recently have been made, think about the WebSphere MQ system, and also about the other programs it interfaces with, the hardware, and any new applications. Consider also the possibility that a new application that you are not aware of might have been run on the system.

- Have you changed, added, or deleted any queue definitions?
- Have you changed or added any channel definitions? Changes might have been made to either WebSphere MQ channel definitions or any underlying communications definitions required by your application.
- Do your applications deal with return codes that they might get as a result of any changes you have made?
- Have you changed any component of the operating system that could affect the operation of WebSphere MQ? For example, have you modified the Windows Registry.

## Have you applied any maintenance updates?

If you have applied a maintenance update to WebSphere MQ, check that the update action completed successfully and that no error message was produced.

- Did the update have any special instructions?
- Was any test run to verify that the update was applied correctly and completely?
- Does the problem still exist if WebSphere MQ is restored to the previous maintenance level?
- If the installation was successful, check with the IBM Support Center for any maintenance package errors.
- If a maintenance package has been applied to any other program, consider the effect it might have on the way WebSphere MQ interfaces with it.

## Are there any error messages or return codes to explain the problem?

You might find error messages or return codes that help you to determine the location and cause of your problem.

IBM WebSphere MQ uses error logs to capture messages concerning its own operation, any queue managers that you start, and error data coming from the channels that are in use. Check the error logs to see if any messages have been recorded that are associated with your problem.

WebSphere MQ also logs errors in the Windows Application Event Log. On Windows, check if the Windows Application Event Log shows any WebSphere MQ errors. To open the log, from the Computer Management panel, expand **Event Viewer** and select **Application**.

For information about the locations and contents of the error logs, see “Error logs on Windows, UNIX and Linux systems” on page 1227

For each WebSphere MQ Message Queue Interface (MQI) and WebSphere MQ Administration Interface (MQAI) call, a completion code and a reason code are returned by the queue manager or by an exit routine, to indicate the success or failure of the call. If your application gets a return code indicating that a Message Queue Interface (MQI) call has failed, check the reason code to find out more about the problem.

For a list of reason codes, see “API completion and reason codes” on page 1380.

Detailed information on return codes is contained within the description of each MQI call.

## Related concepts:

 IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

 Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

## Related reference:

 Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)

“PCF reason codes” on page 1590

“Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes” on page 1672

“WCF custom channel exceptions” on page 1676

## Can you reproduce the problem?

If you can reproduce the problem, consider the conditions under which it is reproduced:

- Is it caused by a command or an equivalent administration request?  
Does the operation work if it is entered by another method? If the command works if it is entered on the command line, but not otherwise, check that the command server has not stopped, and that the queue definition of the SYSTEM.ADMIN.COMMAND.QUEUE has not been changed.
- Is it caused by a program? Does it fail on all WebSphere MQ systems and all queue managers, or only on some?
- Can you identify any application that always seems to be running in the system when the problem occurs? If so, examine the application to see if it is in error.

## Are you receiving an error code when creating or starting a queue manager? (Windows only)

If the WebSphere MQ Explorer, or the `amqmdain` command, fails to create or start a queue manager, indicating an authority problem, it might be because the user under which the IBM WebSphere MQ Windows service is running has insufficient rights.

Ensure that the user with which the IBM WebSphere MQ Windows service is configured has the rights described in “User rights required for a IBM WebSphere MQ Windows Service” on page 69. By default this service is configured to run as the `MUSR_MQADMIN` user. For subsequent installations, the Prepare IBM WebSphere MQ Wizard creates a user account named `MUSR_MQADMINx`, where `x` is the next available number representing a user ID that does not exist.

## Does the problem affect only remote queues?

Things to check if the problem affects only remote queues.

If the problem affects only remote queues, perform the following checks:

- Check that required channels have started, can be triggered, and any required initiators are running.
- Check that the programs that should be putting messages to the remote queues have not reported problems.
- If you use triggering to start the distributed queuing process, check that the transmission queue has triggering set on. Also, check that the trigger monitor is running.
- Check the error logs for messages indicating channel errors or problems.
- If necessary, start the channel manually.

## Have you obtained incorrect output?

In this section, *incorrect output* refers to your application: not receiving a message that you were expecting it to receive; receiving a message containing unexpected or corrupted information; receiving a message that you were not expecting it to receive, for example, one that was destined for a different application.

### Messages that do not arrive on the queue

If messages do not arrive when you are expecting them, check for the following:

- Has the message been put on the queue successfully?
  - Has the queue been defined correctly? For example, is MAXMSGL sufficiently large?
  - Is the queue enabled for putting?
  - Is the queue already full?
  - Has another application got exclusive access to the queue?
- Are you able to get any messages from the queue?
  - Do you need to take a sync point?

If messages are being put or retrieved within sync point, they are not available to other tasks until the unit of recovery has been committed.
  - Is your wait interval long enough?

You can set the wait interval as an option for the MQGET call. Ensure that you are waiting long enough for a response.
  - Are you waiting for a specific message that is identified by a message or correlation identifier (*MsgId* or *CorrelId*)?

Check that you are waiting for a message with the correct *MsgId* or *CorrelId*. A successful MQGET call sets both these values to that of the message retrieved, so you might need to reset these values in order to get another message successfully.

Also, check whether you can get other messages from the queue.
  - Can other applications get messages from the queue?
  - Was the message you are expecting defined as persistent?

If not, and IBM WebSphere MQ has been restarted, the message has been lost.
  - Has another application got exclusive access to the queue?

If you cannot find anything wrong with the queue, and IBM WebSphere MQ is running, check the process that you expected to put the message onto the queue for the following:

- Did the application start?

If it should have been triggered, check that the correct trigger options were specified.
- Did the application stop?
- Is a trigger monitor running?
- Was the trigger process defined correctly?
- Did the application complete correctly?

Look for evidence of an abnormal end in the job log.
- Did the application commit its changes, or were they backed out?

If multiple transactions are serving the queue, they can conflict with one another. For example, suppose one transaction issues an MQGET call with a buffer length of zero to find out the length of the message, and then issues a specific MQGET call specifying the *MsgId* of that message. However, in the meantime, another transaction issues a successful MQGET call for that message, so the first application receives a reason code of MQRC\_NO\_MSG\_AVAILABLE. Applications that are expected to run in a multiple server environment must be designed to cope with this situation.

Consider that the message could have been received, but that your application failed to process it in some way. For example, did an error in the expected format of the message cause your program to reject it? If so, refer to the subsequent information in this topic.

## Messages that contain unexpected or corrupted information

If the information contained in the message is not what your application was expecting, or has been corrupted in some way, consider the following:

- Has your application, or the application that put the message onto the queue, changed?  
Ensure that all changes are simultaneously reflected on all systems that need to be aware of the change.  
For example, the format of the message data might have been changed, in which case, both applications must be recompiled to pick up the changes. If one application has not been recompiled, the data will appear corrupted to the other.
- Is an application sending messages to the wrong queue?  
Check that the messages your application is receiving are not intended for an application servicing a different queue. If necessary, change your security definitions to prevent unauthorized applications from putting messages on to the wrong queues.  
If your application uses an alias queue, check that the alias points to the correct queue.
- Has the trigger information been specified correctly for this queue?  
Check that your application should have started; or should a different application have started?

If these checks do not enable you to solve the problem, check your application logic, both for the program sending the message, and for the program receiving it.

## Problems with incorrect output when using distributed queues

If your application uses distributed queues, consider the following points:

- Has IBM WebSphere MQ been correctly installed on both the sending and receiving systems, and correctly configured for distributed queuing?
- Are the links available between the two systems?  
Check that both systems are available, and connected to IBM WebSphere MQ. Check that the connection between the two systems is active.  
You can use the MQSC command PING against either the queue manager (PING QMGR) or the channel (PING CHANNEL) to verify that the link is operable.
- Is triggering set on in the sending system?
- Is the message for which you are waiting a reply message from a remote system?  
Check that triggering is activated in the remote system.
- Is the queue already full?  
If so, check if the message has been put onto the dead-letter queue.  
The dead-letter queue header contains a reason or feedback code explaining why the message could not be put onto the target queue. See  Using the dead-letter (undelivered message) queue (*WebSphere MQ V7.1 Programming Guide*) and  MQDLH – Dead-letter header (*WebSphere MQ V7.1 Reference*) for information about the dead-letter queue header structure.
- Is there a mismatch between the sending and receiving queue managers?  
For example, the message length could be longer than the receiving queue manager can handle.
- Are the channel definitions of the sending and receiving channels compatible?

For example, a mismatch in sequence number wrap can stop the distributed queuing component. See

 Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*) for more information about distributed queuing.

- Is data conversion involved? If the data formats between the sending and receiving applications differ, data conversion is necessary. Automatic conversion occurs when the MQGET call is issued if the format is recognized as one of the built-in formats.

If the data format is not recognized for conversion, the data conversion exit is taken to allow you to perform the translation with your own routines.

Refer to  Data conversion (*WebSphere MQ V7.1 Reference*) for further information about data conversion.

## Are some of your queues failing?

If you suspect that the problem occurs with only a subset of queues, check the local queues that you think are having problems.

Perform the following checks:

1. Display the information about each queue. You can use the MQSC command DISPLAY QUEUE to display the information.
2. Use the data displayed to do the following checks:
  - If CURDEPTH is at MAXDEPTH, the queue is not being processed. Check that all applications are running normally.
  - If CURDEPTH is not at MAXDEPTH, check the following queue attributes to ensure that they are correct:
    - If triggering is being used:
      - Is the trigger monitor running?
      - Is the trigger depth too great? That is, does it generate a trigger event often enough?
      - Is the process name correct?
      - Is the process available and operational?
    - Can the queue be shared? If not, another application could already have it open for input.
    - Is the queue enabled appropriately for GET and PUT?
  - If there are no application processes getting messages from the queue, determine why this is so. It could be because the applications need to be started, a connection has been disrupted, or the MQOPEN call has failed for some reason.

Check the queue attributes IPPROCS and OPPROCS. These attributes indicate whether the queue has been opened for input and output. If a value is zero, it indicates that no operations of that type can occur. The values might have changed; the queue might have been open but is now closed.

You need to check the status at the time you expect to put or get a message.

If you are unable to solve the problem, contact your IBM Support Center for help.

## Have you failed to receive a response from a PCF command?

Considerations if you have issued a command but have not received a response.

If you have issued a command but have not received a response, consider the following checks:

- Is the command server running?

Work with the **dspmqcsv** command to check the status of the command server.

- If the response to this command indicates that the command server is not running, use the **strmqcsv** command to start it.

- If the response to the command indicates that the `SYSTEM.ADMIN.COMMAND.QUEUE` is not enabled for `MQGET` requests, enable the queue for `MQGET` requests.
- Has a reply been sent to the dead-letter queue?

The dead-letter queue header structure contains a reason or feedback code describing the problem. See

 [MQDLH – Dead-letter header \(WebSphere MQ V7.1 Reference\)](#) and  [Using the dead-letter \(undelivered message\) queue \(WebSphere MQ V7.1 Programming Guide\)](#) for information about the dead-letter queue header structure (MQDLH).

If the dead-letter queue contains messages, you can use the provided browse sample application (`amqsbcg`) to browse the messages using the `MQGET` call. The sample application steps through all the messages on a named queue for a named queue manager, displaying both the message descriptor and the message context fields for all the messages on the named queue.

- Has a message been sent to the error log?  
See “Error log directories” on page 1229 for further information.
- Are the queues enabled for put and get operations?
- Is the *WaitInterval* long enough?  
If your `MQGET` call has timed out, a completion code of `MQCC_FAILED` and a reason code of `MQRC_NO_MSG_AVAILABLE` are returned. (See  [WaitInterval \(MQLONG\) \(WebSphere MQ V7.1 Reference\)](#) for information about the *WaitInterval* field, and completion and reason codes from `MQGET`.)
- If you are using your own application program to put commands onto the `SYSTEM.ADMIN.COMMAND.QUEUE`, do you need to take a sync point?  
Unless you have excluded your request message from sync point, you need to take a sync point before receiving reply messages.
- Are the `MAXDEPTH` and `MAXMSGL` attributes of your queues set sufficiently high?
- Are you using the *CorrelId* and *MsgId* fields correctly?  
Set the values of *MsgId* and *CorrelId* in your application to ensure that you receive all messages from the queue.

Try stopping the command server and then restarting it, responding to any error messages that are produced.

If the system still does not respond, the problem could be with either a queue manager or the whole of the IBM WebSphere MQ system. First, try stopping individual queue managers to isolate a failing queue manager. If this step does not reveal the problem, try stopping and restarting IBM WebSphere MQ, responding to any messages that are produced in the error log.

If the problem still occurs after restart, contact your IBM Support Center for help.

## Has the application run successfully before?

Use the information in this topic to help diagnose common problems with applications.

If the problem appears to involve one particular application, consider whether the application has run successfully before.

Before you answer **Yes** to this question, consider the following:

- Have any changes been made to the application since it last ran successfully?  
If so, it is likely that the error lies somewhere in the new or modified part of the application. Take a look at the changes and see if you can find an obvious reason for the problem. Is it possible to retry using a back level of the application?
- Have all the functions of the application been fully exercised before?

Could it be that the problem occurred when part of the application that had never been invoked before was used for the first time? If so, it is likely that the error lies in that part of the application. Try to find out what the application was doing when it failed, and check the source code in that part of the program for errors.

If a program has been run successfully on many previous occasions, check the current queue status and the files that were being processed when the error occurred. It is possible that they contain some unusual data value that invokes a rarely-used path in the program.

- Does the application check all return codes?

Has your WebSphere MQ system been changed, perhaps in a minor way, such that your application does not check the return codes it receives as a result of the change. For example, does your application assume that the queues it accesses can be shared? If a queue has been redefined as exclusive, can your application deal with return codes indicating that it can no longer access that queue?

- Does the application run on other WebSphere MQ systems?

Could it be that there is something different about the way that this WebSphere MQ system is set up that is causing the problem? For example, have the queues been defined with the same message length or priority?

Before you look at the code, and depending upon which programming language the code is written in, examine the output from the translator, or the compiler and linkage editor, to see if any errors have been reported.

If your application fails to translate, compile, or link-edit into the load library, it will also fail to run if you attempt to invoke it. See  Developing applications (*WebSphere MQ V7.1 Programming Guide*) for information about building your application.

If the documentation shows that each of these steps was accomplished without error, consider the coding logic of the application. Do the symptoms of the problem indicate the function that is failing and, therefore, the piece of code in error? See the following section for some examples of common errors that cause problems with WebSphere MQ applications.

## Common programming errors

The errors in the following list illustrate the most common causes of problems encountered while running WebSphere MQ programs. Consider the possibility that the problem with your WebSphere MQ system could be caused by one or more of these errors:

- Assuming that queues can be shared, when they are in fact exclusive.
- Passing incorrect parameters in an MQI call.
- Passing insufficient parameters in an MQI call. This might mean that WebSphere MQ cannot set up completion and reason codes for your application to process.
- Failing to check return codes from MQI requests.
- Passing variables with incorrect lengths specified.
- Passing parameters in the wrong order.
- Failing to initialize *MsgId* and *CorrelId* correctly.
- Failing to initialize *Encoding* and *CodedCharSetId* following MQRC\_TRUNCATED\_MSG\_ACCEPTED.

## Is your application or system running slowly?

If your application is running slowly, it might be in a loop, or waiting for a resource that is not available, or there might be a performance problem.

Perhaps your system is operating near the limits of its capacity. This type of problem is probably worst at peak system load times, typically at mid-morning and mid-afternoon. (If your network extends across more than one time zone, peak system load might seem to occur at some other time.)

A performance problem might be caused by a limitation of your hardware.

If you find that performance degradation is not dependent on system loading, but happens sometimes when the system is lightly loaded, a poorly-designed application program is probably to blame. This could appear to be a problem that only occurs when certain queues are accessed.

If the performance issue persists, the problem might lie with IBM WebSphere MQ itself. If you suspect this, contact your IBM Support Center for help.

A common cause of slow application performance, or the build up of messages on a queue (usually a transmission queue) is one or more applications that write persistent messages outside a unit of work; for more information, see  Message persistence (*WebSphere MQ V7.1 Programming Guide*).

## Does the problem affect specific parts of the network?

You might be able to identify specific parts of the network that are affected by the problem (remote queues, for example). If the link to a remote message queue manager is not working, the messages cannot flow to a remote queue.

Check that the connection between the two systems is available, and that the intercommunication component of WebSphere MQ has started.

Check that messages are reaching the transmission queue, and check the local queue definition of the transmission queue and any remote queues.

Have you made any network-related changes, or changed any WebSphere MQ definitions, that might account for the problem?

## Does the problem occur at specific times of the day?

If the problem occurs at specific times of day, it could be that it depends on system loading. Typically, peak system loading is at mid-morning and mid-afternoon, so these are the times when load-dependent problems are most likely to occur. (If your WebSphere MQ network extends across more than one time zone, peak system loading might seem to occur at some other time of day.)

## Is the problem intermittent?

An intermittent problem could be caused by the way that processes can run independently of each other. For example, a program might issue an MQGET call without specifying a wait option before an earlier process has completed. An intermittent problem might also be seen if your application tries to get a message from a queue before the call that put the message has been committed.

---

## Making initial checks on IBM i

Before you start problem determination in detail, consider whether there is an obvious cause of the problem, or an area of investigation that is likely to give useful results. This approach to diagnosis can often save a lot of work by highlighting a simple error, or by narrowing down the range of possibilities.

The cause of your problem could be in any of the following:

- Hardware
- Operating system
- Related software, for example, a language compiler
- The network
- The IBM WebSphere MQ product
- Your IBM WebSphere MQ application
- Other applications
- Site operating procedures

This section contains a list of questions to consider. As you go through the list, make a note of anything that might be relevant to the problem. Even if your observations do not suggest a cause straight away, they might be useful later if you have to carry out a systematic problem determination exercise.

The following steps are intended to help you isolate the problem and are taken from the viewpoint of a IBM WebSphere MQ application. Check all the suggestions at each stage.

1. Has IBM WebSphere MQ for IBM i run successfully before?

**Yes** Proceed to Step 2.

**No** It is likely that you have not installed or set up IBM WebSphere MQ correctly.

2. Has the IBM WebSphere MQ application run successfully before?

**Yes** Proceed to Step 3.

**No** Consider the following:

a. The application might have failed to compile or link, and fails if you attempt to invoke it. Check the output from the compiler or linker.

Refer to the appropriate programming language reference information, or see



Developing applications (*WebSphere MQ V7.1 Programming Guide*), for information on how to build your application.

b. Consider the logic of the application. For example, do the symptoms of the problem indicate that a function is failing and, therefore, that a piece of code is in error.

Check the following common programming errors:

- Assuming that queues can be shared, when they are in fact exclusive.
- Trying to access queues and data without the correct security authorization.
- Passing incorrect parameters in an MQI call; if the wrong number of parameters is passed, no attempt can be made to complete the completion code and reason code fields, and the task is ended abnormally.
- Failing to check return codes from MQI requests.
- Using incorrect addresses.
- Passing variables with incorrect lengths specified.
- Passing parameters in the wrong order.
- Failing to initialize *MsgId* and *CorrelId* correctly.

3. Has the IBM WebSphere MQ application changed since the last successful run?

- Yes** It is likely that the error lies in the new or modified part of the application. Check all the changes and see if you can find an obvious reason for the problem.
- a. Have all the functions of the application been fully exercised before?
 

Could it be that the problem occurred when part of the application that had never been invoked before was used for the first time? If so, it is likely that the error lies in that part of the application. Try to find out what the application was doing when it failed, and check the source code in that part of the program for errors.
  - b. If the program has run successfully before, check the current queue status and files that were being processed when the error occurred. It is possible that they contain some unusual data value that causes a rarely used path in the program to be invoked.
  - c. The application received an unexpected MQI return code. For example:
    - Does your application assume that the queues it accesses are shareable? If a queue has been redefined as exclusive, can your application deal with return codes indicating that it can no longer access that queue?
    - Have any queue definition or security profiles been changed? An MQOPEN call could fail because of a security violation; can your application recover from the resulting return code?

See  Developing applications reference (*WebSphere MQ V7.1 Reference*) for your programming language for a description of each return code.
  - d. If you have applied any PTF to IBM WebSphere MQ for IBM i, check that you received no error messages when you installed the PTF.

**No** Ensure that you have eliminated all the preceding suggestions and proceed to Step 4.

4. Has the server system remained unchanged since the last successful run?

**Yes** Proceed to “Problem characteristics” on page 1163.

**No** Consider all aspects of the system and review the appropriate documentation on how the change might have affected the IBM WebSphere MQ application. For example :

- Interfaces with other applications
- Installation of new operating system or hardware
- Application of PTFs
- Changes in operating procedures

See the following sections for some additional tips for problem determination for system administrators and application developers.

### Tips for system administrators

- Check the error logs for messages for your operating system:
  - “Error logs on Windows, UNIX and Linux systems” on page 1227
  - “Error logs on IBM i” on page 1231
  - “Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272
- Check the contents of qm.ini for any configuration changes or errors. For more information on changing configuration information, see:
  -  UNIX and Linux (*WebSphere MQ V7.1 Installing Guide*)
  -  IBM i (*WebSphere MQ V7.1 Installing Guide*)
  -  z/OS (*WebSphere MQ V7.1 Installing Guide*)
- If your application development teams are reporting something unexpected, you use trace to investigate the problems. For information about using trace, see “Using trace” on page 1234.

## Tips for application developers

- Check the return codes from the MQI calls in your applications. For a list of reason codes, see “API reason codes” on page 1380. Use the information provided in the return code to determine the cause of the problem. Follow the steps in the Programmer response sections of the reason code to resolve the problem.
- If you are unsure whether your application is working as expected, for example, you are not unsure of the parameters being passed into the MQI or out of the MQI, you can use trace to collect information about all the inputs and outputs of your MQI calls. For more information about using trace, see “Using trace” on page 1234.
- For more information about handling errors in MQI applications, see  Handling program errors (*WebSphere MQ V7.1 Programming Guide*).

### Related concepts:

“Troubleshooting and support” on page 1149

“Required authority” on page 1166

“Making initial checks on Windows, UNIX and Linux systems” on page 1151

“Making initial checks on z/OS” on page 1169

“Dealing with problems” on page 1186

“IBM Support Assistant (ISA)” on page 1324

“Reason codes” on page 1379

 Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

### Related tasks:

“Searching knowledge bases” on page 1326

“Contacting IBM Software Support” on page 1346

### Related reference:

“Determining problems with IBM WebSphere MQ for IBM i applications” on page 1167

“PCF reason codes” on page 1590

## Problem characteristics

Perhaps by using the preliminary checks, you have found the cause of the problem. If so, you can probably now resolve it, possibly with the help of other sections in the IBM WebSphere MQ product documentation, and in the libraries of other licensed programs.

If you have not yet found the cause, start to look at the problem in greater detail. Use the following questions as pointers to the problem. Answering the appropriate question, or questions, should lead you to the cause of the problem:

- “Can you reproduce the problem?”
- “Is the problem intermittent?” on page 1164
- “Problems with commands” on page 1164
- “Does the problem affect all users of the IBM WebSphere MQ for IBM i application?” on page 1164
- “Does the problem affect specific parts of the network?” on page 1164
- “Does the problem occur only on WebSphere MQ?” on page 1165
- “Does the problem occur at specific times of the day?” on page 1165
- “Have you failed to receive a response from a command?” on page 1165

## Can you reproduce the problem?

If you can reproduce the problem, consider the conditions under which you do so:

- Is it caused by a command?

Does the operation work if it is entered by another method? If the command works if it is entered on the command line, but not otherwise, check that the command server has not stopped. You must also check that the queue definition of the `SYSTEM.ADMIN.COMMAND.QUEUE` has not been changed.

- Is it caused by a program? If so, does it fail in batch? Does it fail on all IBM WebSphere MQ for IBM i systems, or only on some?
- Can you identify any application that always seems to be running in the system when the problem occurs? If so, examine the application to see if it is in error.
- Does the problem occur with any queue manager, or when connected to one specific queue manager?
- Does the problem occur with the same type of object on any queue manager, or only one particular object? What happens after this object has been cleared or redefined?
- Is the problem independent of any message persistence settings?
- Does the problem occur only when sync points are used?
- Does the problem occur only when one or more queue-manager events are enabled?

### **Is the problem intermittent?**

An intermittent problem might be caused by failing to take into account the fact that processes can run independently of each other. For example, a program might issue an `MQGET` call, without specifying a wait option, before an earlier process has completed. You might also encounter this problem if your application tries to get a message from a queue while the call that put the message is in-doubt (that is, before it has been committed or backed out).

### **Problems with commands**

Use this information to avoid potential problems with special characters. Be careful when including special characters, for example backslash (`\`) and quotation marks (`"`) characters, in descriptive text for some commands. If you use either of these characters in descriptive text, precede them with a backslash (`\`) character, for example:

- Enter `\\` if you need a backslash (`\`) character in your text.
- Enter `\"` if you need quotation marks (`"`) characters in your text.

Queue managers and their associated object names are case-sensitive. By default, IBM i uses uppercase characters, unless you surround the name in apostrophe (`'`) characters.

For example, `MYQUEUE` and `myqueue` translate to `MYQUEUE`, whereas `'myqueue'` translates to `myqueue`.

### **Does the problem affect all users of the IBM WebSphere MQ for IBM i application?**

If the problem affects only some users, look for differences in how the users configure their systems and queue manager settings.

Check the library lists and user profiles. Can the problem be circumvented by having `*ALLOBJ` authority?

### **Does the problem affect specific parts of the network?**

You might be able to identify specific parts of the network that are affected by the problem (remote queues, for example). If the link to a remote message queue manager is not working, the messages cannot flow to a remote queue.

Check these points:

- Is the connection between the two systems available, and has the intercommunication component of IBM WebSphere MQ for IBM i started?

Check that messages are reaching the transmission queue, the local queue definition of the transmission queue, and any remote queues.

- Have you made any network-related changes that might account for the problem or changed any IBM WebSphere MQ for IBM i definitions?
- Can you distinguish between a channel definition problem and a channel message problem?  
For example, redefine the channel to use an empty transmission queue. If the channel starts correctly, the definition is correctly configured.

## Does the problem occur only on WebSphere MQ?

If the problem occurs only on this version of IBM WebSphere MQ, check the appropriate database on RETAIN, or the  [https://www.ibm.com/support/home/product/P439881V74305Y86/IBM\\_MQ](https://www.ibm.com/support/home/product/P439881V74305Y86/IBM_MQ), to ensure that you have applied all the relevant PTFs.

## Does the problem occur at specific times of the day?

If the problem occurs at specific times of day, it might be that it is dependent on system loading. Typically, peak system loading is at mid-morning and mid-afternoon, and so these times are when load-dependent problems are most likely to occur. (If your IBM WebSphere MQ for IBM i network extends across more than one time zone, peak system loading might seem to occur at some other time of day.)

## Have you failed to receive a response from a command?

If you have issued a command but you have not received a response, consider the following questions:

- Is the command server running?

Work with the **DSPMQMCSVR** command to check the status of the command server.

- If the response to this command indicates that the command server is not running, use the **STRMQMCSVR** command to start it.
- If the response to the command indicates that the **SYSTEM.ADMIN.COMMAND.QUEUE** is not enabled for MQGET requests, enable the queue for MQGET requests.

- Has a reply been sent to the dead-letter queue?

The dead-letter queue header structure contains a reason or feedback code describing the problem. See

 **MQDLH** – Dead-letter header (*WebSphere MQ V7.1 Reference*) for information about the dead-letter queue header structure (MQDLH).

If the dead-letter queue contains messages, you can use the provided browse sample application (amqsbcg) to browse the messages using the MQGET call. The sample application steps through all the messages on a named queue for a named queue manager, displaying both the message descriptor and the message context fields for all the messages on the named queue.

- Has a message been sent to the error log?

See “Error logs on IBM i” on page 1231 for further information.

- Are the queues enabled for put and get operations?
- Is the *WaitInterval* long enough?

If your MQGET call has timed out, a completion code of MQCC\_FAILED and a reason code of

MQRC\_NO\_MSG\_AVAILABLE are returned. (See  Getting messages from a queue using the MQGET call (*WebSphere MQ V7.1 Programming Guide*) for more information about the *WaitInterval* field, and completion and reason codes from MQGET.)

- If you are using your own application program to put commands onto the **SYSTEM.ADMIN.COMMAND.QUEUE**, do you need to take a sync point?

Unless you have excluded your request message from sync point, you must take a sync point before attempting to receive reply messages.

- Are the MAXDEPTH and MAXMSGL attributes of your queues set sufficiently high?
- Are you using the *CorrelId* and *MsgId* fields correctly?

Set the values of *MsgId* and *CorrelId* in your application to ensure that you receive all messages from the queue.

**Related concepts:**

“Troubleshooting and support” on page 1149

“Required authority”

**Related reference:**

“Determining problems with IBM WebSphere MQ for IBM i applications” on page 1167

## Required authority

Some IBM WebSphere MQ commands rely on using IBM i system commands for creating and managing objects, files, and libraries.

For example, CRTMQM (create queue manager) and DLTMQM (delete queue manager). Similarly some IBM WebSphere MQ program code, for example a queue manager, relies on using IBM i system programs.

To enable this reliance, the commands and programs listed must either have \*PUBLIC \*USE authority, or explicit \*USE authority to the IBM WebSphere MQ user profiles QMQM and QMQADM.

Such authority is applied automatically as part of the install process, and you do not need to apply it yourself.

However, if you encounter problems, here is how to do it manually:

1. Set the authorities for commands using GRTOBJAUT with an OBJTYPE(\*CMD) parameter, for example:

```
GRTOBJAUT OBJ(QSYS/ADDLIB) OBJTYPE(*CMD) USER(QMQADM) AUT(*USE)
```

- QSYS/ADDLIB
- QSYS/ADDPFM
- QSYS/CALL
- QSYS/CHGCURLIB
- QSYS/CHGJOB
- QSYS/CRTJRN
- QSYS/CRTJRNRVC
- QSYS/CRTJOBQ
- QSYS/CRTJOB
- QSYS/CRTL
- QSYS/CRTMSGQ
- QSYS/CRTPF
- QSYS/CRTPGM
- QSYS/CRTSRCPF
- QSYS/DLTJRN
- QSYS/DLTJRNRVC
- QSYS/DLTL
- QSYS/DLTMSGQ
- QSYS/OVRPRTF
- QSYS/RCLACTGRP
- QSYS/RTVJRNE
- QSYS/RCVJRNE
- QSYS/SBMJOB

2. Set the authorities for programs using GRTOBJAUT with an OBJTYPE(\*PGM) parameter, for example:

GRTOBJAUT OBJ(QSYS/QWTSETP) OBJTYPE(\*PGM) USER(QMQMADM) AUT(\*USE)

- QSYS/QWTSETP(\*PGM)
- QSYS/QSYRLSPH(\*PGM)
- QSYS/QSYGETPH(\*PGM)

## Determining problems with IBM WebSphere MQ for IBM i applications

If you have not yet found the cause, start to look at the problem in greater detail.

This section contains information about problems you might encounter with IBM WebSphere MQ applications, commands, and messages. Use the following questions as pointers to the problem. Answering the appropriate question, or questions, should lead you to the cause of the problem.

### Are some of your queues working?

If you suspect that the problem occurs with only a subset of queues, select the name of a local queue that you think is having problems.

1. Display the information about this queue, using WRKMQMSTS or DSPMQMQ.
2. Use the data displayed to do the following checks:
  - If CURDEPTH is at MAXDEPTH, the queue is not being processed. Check that all applications are running normally.
  - If CURDEPTH is not at MAXDEPTH, check the following queue attributes to ensure that they are correct:
    - If triggering is being used:
      - Is the trigger monitor running?
      - Is the trigger depth too large?
      - Is the process name correct?
    - Can the queue be shared? If not, another application might already have it open for input.
    - Is the queue enabled appropriately for GET and PUT?
  - If there are no application processes getting messages from the queue, determine why (for example, because the applications must be started, a connection has been disrupted, or because the MQOPEN call has failed for some reason).

If you cannot solve the problem, contact your IBM support center for help.

### Does the problem affect only remote queues?

If the problem affects only remote queues, check the subsequent points:

1. Check that the programs that should be putting messages to the remote queues have run successfully.
2. If you use triggering to start the distributed queuing process, check that the transmission queue has triggering set on. Also, check that the trigger monitor is running.
3. If necessary, start the channel manually. See  Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*) for information about how to do so.
4. Check the channel with a PING command.

### Are messages failing to arrive on the queue?

If messages do not arrive when you are expecting them, check for the following:

- Have you selected the correct queue manager, that is, the default queue manager or a named queue manager?
- Has the message been put on the queue successfully?

- Has the queue been defined correctly, for example, is MAXMSGLEN sufficiently large?
- Can applications put messages on the queue (is the queue enabled for putting)?
- If the queue is already full, it might mean that an application was unable to put the required message on the queue.
- Can you get the message from the queue?
  - Must you take a sync point?
 

If messages are being put or retrieved within sync point, they are not available to other tasks until the unit of recovery has been committed.
  - Is your timeout interval long enough?
  - Are you waiting for a specific message that is identified by a message identifier or correlation identifier (*MsgId* or *CorrelId*)?
 

Check that you are waiting for a message with the correct *MsgId* or *CorrelId*. A successful MQGET call sets both these values to that of the message retrieved, so you might need to reset these values in order to get another message successfully.

Also check if you can get other messages from the queue.
  - Can other applications get messages from the queue?
  - Was the message you are expecting defined as persistent?
 

If not, and IBM WebSphere MQ for IBM i has been restarted, the message has been lost.

If you cannot find anything wrong with the queue, and the queue manager itself is running, make the following checks on the process that you expected to put the message on to the queue:

- Did the application start?
 

If it should have been triggered, check that the correct trigger options were specified.
- Is a trigger monitor running?
- Was the trigger process defined correctly?
- Did it complete correctly?
 

Look for evidence of an abnormal end in the job log.
- Did the application commit its changes, or were they backed out?

If multiple transactions are serving the queue, they might occasionally conflict with one another. For example, one transaction might issue an MQGET call with a buffer length of zero to find out the length of the message, and then issue a specific MQGET call specifying the *MsgId* of that message. However, in the meantime, another transaction might have issued a successful MQGET call for that message, so the first application receives a completion code of MQRC\_NO\_MSG\_AVAILABLE. Applications that are expected to run in a multi-server environment must be designed to cope with this situation.

Consider that the message might have been received, but that your application failed to process it in some way. For example, did an error in the expected format of the message cause your program to reject it? If so, see “Are unexpected messages received when using distributed queues?” on page 1169.

## Do messages contain unexpected or corrupted information?

If the information contained in the message is not what your application was expecting, or has been corrupted in some way, consider the following points:

- Has your application, or the application that put the message on to the queue, changed?
 

Ensure that all changes are simultaneously reflected on all systems that need to be aware of the change.

For example, a copyfile formatting the message might have been changed, in which case, recompile both applications to pick up the changes. If one application has not been recompiled, the data appears corrupted to the other.

- Is an application sending messages to the wrong queue?  
Check that the messages your application is receiving are not intended for an application servicing a different queue. If necessary, change your security definitions to prevent unauthorized applications from putting messages on to the wrong queues.  
If your application has used an alias queue, check that the alias points to the correct queue.
- Has the trigger information been specified correctly for this queue?  
Check that your application should have been started, or should a different application have been started?
- Has the CCSID been set correctly, or is the message format incorrect because of data conversion.

If these checks do not enable you to solve the problem, check your application logic, both for the program sending the message, and for the program receiving it.

## Are unexpected messages received when using distributed queues?

If your application uses distributed queues, consider the following points:

- Has distributed queuing been correctly installed on both the sending and receiving systems?
- Are the links available between the two systems?

Check that both systems are available, and connected to IBM WebSphere MQ for IBM i. Check that the connection between the two systems is active.

- Is triggering set on in the sending system?
- Is the message you are waiting for a reply message from a remote system?

Check that triggering is activated in the remote system.

- Is the queue already full?

If so, it might mean that an application was unable to put the required message on to the queue. Check that the message has been put onto the undelivered-message queue.

The dead-letter queue message header (dead-letter header structure) contains a reason or feedback code explaining why the message could not be put on to the target queue. See  MQDLH –

Dead-letter header (*WebSphere MQ V7.1 Reference*) or  IBM i Application Programming Reference (ILE/RPG) (*WebSphere MQ V7.1 Reference*), as appropriate, for information about the dead-letter header structure.

- Is there a mismatch between the sending and receiving queue managers?

For example, the message length could be longer than the receiving queue manager can handle.

- Are the channel definitions of the sending and receiving channels compatible?

For example, a mismatch in sequence number wrap stops the distributed queuing component. See

 Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*) for more information about distributed queuing.

---

## Making initial checks on z/OS

Before you start problem determination in detail, consider whether there is an obvious cause of the problem, or an area of investigation that is likely to give useful results. This approach to diagnosis can often save a lot of work by highlighting a simple error, or by narrowing down the range of possibilities.

The cause of your problem could be in:

- IBM WebSphere MQ
- The network
- The application
- Other applications that you have configured to work with IBM WebSphere MQ

This section contains a list of questions to consider. As you go through the list, make a note of anything that might be relevant to the problem. Even if your observations do not suggest a cause straight away, they might be useful later if you have to carry out a systematic problem determination exercise.

- “Has IBM WebSphere MQ for z/OS run successfully before?” on page 1171
- “Have you applied any APARs or PTFs?” on page 1172
- “Are there any error messages, return codes or other error conditions?” on page 1172
- “Has your application or WebSphere MQ for z/OS stopped processing work?” on page 1174
- “Is there a problem with the WebSphere MQ queues?” on page 1175
- “Are some of your queues working?” on page 1175
- “Are the correct queues defined?” on page 1176
- “Does the problem affect only remote or cluster queues?” on page 1176
- “Does the problem affect only shared queues?” on page 1177
- “Does the problem affect specific parts of the network?” on page 1178
- “Problems that occur at specific times of the day or affect specific users” on page 1178
- “Is the problem intermittent or does the problem occur with all z/OS, CICS, or IMS systems?” on page 1178
- “Has the application run successfully before?” on page 1179
- “Have any changes been made since the last successful run?” on page 1180
- “Do you have a program error?” on page 1181
- “Has there been an abend?” on page 1182
- “Have you obtained incorrect output?” on page 1183
- “Can you reproduce the problem?” on page 1183
- “Have you failed to receive a response from an MQSC command?” on page 1184
- “Is your application or IBM WebSphere MQ for z/OS running slowly?” on page 1185

See the following sections for some additional tips for problem determination for system administrators and application developers.

### Tips for system administrators

- Check the error logs for messages for your operating system:
  - “Error logs on Windows, UNIX and Linux systems” on page 1227
  - “Error logs on IBM i” on page 1231
  - “Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272
- Check the contents of `qm.ini` for any configuration changes or errors. For more information on changing configuration information, see:
  -  UNIX and Linux (*WebSphere MQ V7.1 Installing Guide*)
  -  IBM i (*WebSphere MQ V7.1 Installing Guide*)
  -  z/OS (*WebSphere MQ V7.1 Installing Guide*)
- If your application development teams are reporting something unexpected, you use trace to investigate the problems. For information about using trace, see “Using trace” on page 1234.

### Tips for application developers

- Check the return codes from the MQI calls in your applications. For a list of reason codes, see “API reason codes” on page 1380. Use the information provided in the return code to determine the cause of the problem. Follow the steps in the Programmer response sections of the reason code to resolve the problem.

- If you are unsure whether your application is working as expected, for example, you are not unsure of the parameters being passed into the MQI or out of the MQI, you can use trace to collect information about all the inputs and outputs of your MQI calls. For more information about using trace, see “Using trace” on page 1234.
- For more information about handling errors in MQI applications, see  Handling program errors (*WebSphere MQ V7.1 Programming Guide*).

**Related concepts:**

- “Troubleshooting and support” on page 1149
- “Making initial checks on Windows, UNIX and Linux systems” on page 1151
- “Making initial checks on IBM i” on page 1161
- “Dealing with problems” on page 1186
- “IBM Support Assistant (ISA)” on page 1324
- “Reason codes” on page 1379

 Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

**Related tasks:**

- “Searching knowledge bases” on page 1326
- “Contacting IBM Software Support” on page 1346

**Related reference:**

- “PCF reason codes” on page 1590

## Has IBM WebSphere MQ for z/OS run successfully before?

Whether IBM WebSphere MQ for z/OS has successfully run before can help with problem determination and there are checks you can perform to help you.

If the answer to this question is **No**, consider the following:

- Check your setup
  - If IBM WebSphere MQ has not run successfully on z/OS before, it is likely that you have not yet set it up correctly. See the information about installing and customizing the queue manager in  Installing the WebSphere MQ for z/OS product (*WebSphere MQ V7.1 Installing Guide*) for further guidance.
- Verify the installation
- Check that message CSQ9022I was issued in response to the START QMGR command (indicating normal completion).
- Ensure that z/OS displays IBM WebSphere MQ as an installed subsystem. To determine if IBM WebSphere MQ is an installed subsystem use the z/OS command D OPDATA.
- Check that the installation verification program (IVP) ran successfully.
- Issue the command DISPLAY DQM to check that the channel initiator address space is running, and that the appropriate listeners are started.

## Have you applied any APARs or PTFs?

APARs and PTFs can occasionally cause unexpected problems with WebSphere MQ. These fixes can have been applied to WebSphere MQ or to other z/OS systems.

If an APAR or PTF has been applied to WebSphere MQ for z/OS, check that no error message was produced. If the installation was successful, check with the IBM support center for any APAR or PTF error.

If an APAR or PTF has been applied to any other product, consider the effect it might have on the way WebSphere MQ interfaces with it.

Ensure that you have followed any instructions in the APAR that affect your system. (For example, you might have to redefine a resource.)

## Are there any error messages, return codes or other error conditions?

Use this topic to investigate error messages, return codes, and conditions where the queue manager or channel initiator terminated.

The problem might produce the following types of error message or return codes:

### CSQ messages and reason codes

IBM WebSphere MQ for z/OS error messages have the prefix CSQ; if you receive any messages with this prefix (for example, in the console log, or the CICS log), see  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) for an explanation.

### Other messages

For messages with a different prefix, look in the appropriate messages and codes topic for a suggested course of action. See “The message keyword” on page 1338 which lists possible message prefixes and appropriate manuals.

### Unusual messages

Be aware of unusual messages associated with the startup of IBM WebSphere MQ for z/OS, or issued while the system was running before the error occurred. Any unusual messages might indicate some system problem that prevented your application from running successfully.

### Application MQI return codes

If your application gets a return code indicating that an MQI call has failed, see  Return codes (*WebSphere MQ V7.1 Reference*) for a description of that return code.

## Have you received an unexpected error message or return code?

If your application has received an unexpected error message, consider whether the error message has originated from IBM WebSphere MQ or from another program.

### IBM WebSphere MQ error messages

IBM WebSphere MQ for z/OS error messages are prefixed with the letters CSQ.

If you get an unexpected IBM WebSphere MQ error message (for example, in the console log, or the CICS log), see  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) for an explanation.

 IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) might give you enough information to resolve the problem quickly, or it might redirect you to another manual for further guidance. If you cannot deal with the message, you might have to contact the IBM support center for help.

## Non-IBM WebSphere MQ error messages

If you get an error message from another IBM program, or from the operating system, look in the messages and codes manual from the appropriate library for an explanation of what it means.

In a queue-sharing environment, look for the following error messages:

- XES (prefixed with the letters IXL)
- Db2 (prefixed with the letters DSN)
- RRS (prefixed with the letters ATR)

See “The message keyword” on page 1338 for a list of the most common message prefixes.

## Unexpected return codes

If your application has received an unexpected return code from IBM WebSphere MQ, see

 Return codes (*WebSphere MQ V7.1 Reference*) for information about how your application can handle IBM WebSphere MQ return codes.

## Check for error messages

Issue the DISPLAY THREAD(\*) command to check if the queue manager is running. For more information about the command, see  DISPLAY THREAD (*WebSphere MQ V7.1 Reference*). If the queue manager has stopped running, look for any message that might explain the situation. Messages are displayed on the z/OS console, or on your terminal if you are using the operations and control panels. Use the DISPLAY DQM command to see if the channel initiator is working, and the listeners are active. The z/OS command

```
DISPLAY R,L
```

lists messages with outstanding replies. Check to see whether any of these replies are relevant. In some circumstances, for example, when it has used all its active logs, WebSphere MQ for z/OS waits for operator intervention.

## No error messages issued

If no error messages have been issued, perform the following procedure to determine what is causing the problem:

1. Issue the z/OS commands

```
DISPLAY A,xxxxMSTR
DISPLAY A,xxxxCHIN
```

(where xxxx is the WebSphere MQ for z/OS subsystem name). If you receive a message telling you that the queue manager or channel initiator has not been found, this message indicates that the subsystem has terminated. This condition could be caused by an abend or by operator shutdown of the system.

2. If the subsystem is running, you receive message IEE105I. This message includes the CT=nnnn field, which contains information about the processor time being used by the subsystem. Note the value of this field, and reissue the command.
  - If the CT= value has not changed, this indicates that the subsystem is not using any processor time. This could indicate that the subsystem is in a wait state (or that it has no work to do). If you can issue a command like DISPLAY DQM and you get output back, this indicates there is no work to do rather than a hang condition.
  - If the CT= value has changed dramatically, and continues to do so over repeated displays, this could indicate that the subsystem is busy or possibly in a loop.

- If the reply indicates that the subsystem is now not found, this indicates that it was in the process of terminating when the first command was issued. If a dump is being taken, the subsystem might take a while to terminate. A message is produced at the console before terminating.

To check that the channel initiator is working, issue the DISPLAY DQM command. If the response does not show the channel initiator working this could be because it is getting insufficient resources (like the processor). In this case, use the z/OS monitoring tools, such as RMF, to determine if there is a resource problem. If it is not, restart the channel initiator.

## **Has the queue manager or channel initiator terminated abnormally?**

Look for any messages saying that the queue manager or channel initiator address space has abnormally terminated. If you get a message for which the system action is to terminate WebSphere MQ, find out whether a system dump was produced, see WebSphere MQ dumps.

## **WebSphere MQ for z/OS might still be running**

Consider also that WebSphere MQ for z/OS might still be running, but only slowly. If it is running slowly, you probably have a performance problem. To confirm this, see Is your application or WebSphere MQ for z/OS running slowly. Refer to Dealing with performance problems for advice about what to do next.

## **Has your application or WebSphere MQ for z/OS stopped processing work?**

There are several reasons why your system would unexpectedly stop processing work including problems with the queue manager, the application, z/OS, and the data sets.

There are several reasons why your system would unexpectedly stop processing work. These include:

### **Queue manager problems**

The queue manager might be shutting down.

### **Application problems**

An application programming error might mean that the program branches away from its normal processing, or the application might get in a loop. There might also have been an application abend.

### **WebSphere MQ problems**

Your queues might have become disabled for MQPUT or MQGET calls, the dead-letter queue might be full, or WebSphere MQ for z/OS might be in a wait state, or a loop.

### **z/OS and other system problems**

z/OS might be in a wait state, or CICS or IMS might be in a wait state or a loop. There might be problems at the system or sysplex level that are affecting the queue manager or the channel initiator. For example, excessive paging. It might also indicate DASD problems, or higher priority tasks with high processor usage.

### **Db2 and RRS problems**

Check that Db2 and RRS are active.

In all cases, carry out the following checks to determine the cause of the problem:

## Is there a problem with the WebSphere MQ queues?

Use this topic for investigating potential problems with WebSphere MQ queues.

If you suspect that there is a problem affecting the queues on your subsystem, use the operations and control panels to display the system-command input queue.

### If the system responds

If the system responds, then at least one queue is working. In this case, follow the procedure in "Are some of your queues working?"

### If the system does not respond

The problem might be with the whole subsystem. In this instance, try stopping and restarting the queue manager, responding to any error messages that are produced.

Check for any messages on the console needing action. Resolve any that might affect WebSphere MQ, such as a request to mount a tape for an archive log. See if other subsystems or CICS regions are affected.

Use the DISPLAY QMGR COMMANDQ command to identify the name of the system command input queue.

### If the problem still occurs after restart

Contact your IBM support center for help (see "Contacting IBM Software Support" on page 1346).

### Related concepts:

"Are the correct queues defined?" on page 1176

"Does the problem affect only remote or cluster queues?" on page 1176

"Does the problem affect only shared queues?" on page 1177

## Are some of your queues working?

Use this topic to investigate when problems occur with a subset of your queues.

If you suspect that the problem occurs with only a subset of queues, select the name of a local queue that you think is having problems and perform the following procedures:

### Display queue information

Use the DISPLAY QUEUE and DISPLAY QSTATUS commands to display information about the queue.

### Is the queue being processed?

- If CURDEPTH is at MAXDEPTH, it might indicate that the queue is not being processed. Check that all applications that use the queue are running normally (for example, check that transactions in your CICS system are running or that applications started in response to Queue Depth High events are running).
- Issue DISPLAY QSTATUS(xx) IPPROCS to see if the queue is open for input. If not, start the application.
- If CURDEPTH is not at MAXDEPTH, check the following queue attributes to ensure that they are correct:
  - If triggering is being used:
    - Is the trigger monitor running?
    - Is the trigger depth too big?
    - Is the process name correct?
    - Have **all** the trigger conditions been met?

Issue DISPLAY QSTATUS(xx) IPPROCS to see if an application has the same queue open for input. In some triggering scenarios, a trigger message is not produced if the queue is open for input. Stop the application to cause the triggering processing to be invoked.

- Can the queue be shared? If not, another application (batch, IMS, or CICS) might already have it open for input.
- Is the queue enabled appropriately for GET and PUT?

#### **Do you have a long-running unit of work?**

If CURDEPTH is not zero, but when you attempt to MQGET a message the queue manager replies that there is no message available, issue either DIS QSTATUS(xx) TYPE(HANDLE) to show you information about applications that have the queue open, or issue DIS CONN(xx) to give you more information about an application that is connected to the queue.

#### **How many tasks are accessing the queues?**

Issue DISPLAY QSTATUS(xx) OPPOCS IPPROCS to see how many tasks are putting messages on to, and getting messages from the queue. In a queue-sharing environment, check OPPOCS and IPPROCS on each queue manager. Alternatively, use the CMDSCOPE attribute to check all the queue managers. If there are no application processes getting messages from the queue, determine the reason (for example, because the applications need to be started, a connection has been disrupted, or because the MQOPEN call has failed for some reason).

#### **Is this queue a shared queue? Does the problem affect only shared queues?**

Check that there is not a problem with the sysplex elements that support shared queues. For example, check that there is not a problem with the WebSphere MQ-managed Coupling Facility list structure.

Use D XCF, STRUCTURE, STRNAME=ALL to check that the Coupling Facility structures are accessible.

Use D RRS to check that RRS is active.

#### **Is this queue part of a cluster?**

Check to see if the queue is part of a cluster (from the CLUSTER or CLUSNL attribute). If it is, verify that the queue manager that hosts the queue is still active in the cluster.

#### **If you cannot solve the problem**

Contact your IBM support center for help (see “Contacting IBM Software Support” on page 1346).

#### **Are the correct queues defined?**

WebSphere MQ requires certain predefined queues. Problems can occur if these queues are not defined correctly.

Check that the system-command input queue, the system-command reply model queue, and the reply-to queue are correctly defined, and that the MQOPEN calls were successful.

If you are using the system-command reply model queue, check that it was defined correctly.

If you are using clusters, you need to define the SYSTEM.CLUSTER.COMMAND.QUEUE to use commands relating to cluster processing.

#### **Does the problem affect only remote or cluster queues?**

Use this topic for further investigation if the problem only occurs on remote or cluster queues.

If the problem affects only remote or cluster queues, check:

#### **Are the remote queues being accessed?**

Check that the programs putting messages to the remote queues have run successfully (see “Dealing with incorrect output” on page 1299).

#### **Is the system link active?**

Use APPC or TCP/IP commands as appropriate to check whether the link between the two systems is active.

Use PING or OPING for TCP/IP or D NET ID=xxxxx, E for APPC.

### Is triggering working?

If you use triggering to start the distributed queuing process, check that the transmission queue has triggering set on and that the queue is get-enabled.

### Is the channel or listener running?

If necessary, start the channel or the listener manually, or try stopping and restarting the channel.

See  Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*) for more information.

Look for error messages on the startup of the channel initiator and listener. See  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

and  Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*) to determine the cause.

### What is the channel status?

Check the channel status using the DISPLAY CHSTATUS (channel\_name) command.

### Are your process and channel definitions correct?

Check your process definitions and your channel definitions.

See  Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*) for information about how to use distributed queuing, and for information about how to define channels.

### Does the problem affect only shared queues?

Use this topic to investigate possible queue-sharing group issues which can cause problems for shared queues.

If the problem affects only queue-sharing groups, use the VERIFY QSG function of the CSQ5PQSG utility. This command verifies that the Db2 setup is consistent in terms of the bitmap allocation fields, and object definition for the Db2 queue manager, structure, and shared queue objects, and reports details of any inconsistency that is discovered.

The following is an example of a VERIFY QSG report with errors:

```
CSQU501I VERIFY QSG function requested
CSQU503I QSG=SQ02, DB2 DSG=DSN710P5, DB2 ssid=DFP5
CSQU517I XCF group CSQGSQ02 already defined
CSQU520I Summary information for XCF group CSQGSQ02
CSQU522I Member=MQ04, state=QUIESCED, system=MV4A
CSQU523I User data=D4E5F4C15AD4D8F0F4404040C4C5....
CSQU522I Member=MQ03, state=QUIESCED, system=MV4A
CSQU523I User data=D4E5F4C15AD4D8F0F3404040C4C6....
CSQU526I Connected to DB2 DF4A
CSQU572E Usage map T01_ARRAY_QMGR and DB2 table CSQ.ADMIN_B_QMGR inconsistent
CSQU573E QMGR MQ04 in table entry 1 not set in usage map
CSQU574E QMGR 27 in usage map has no entry in table
CSQU572E Usage map T01_ARRAY_STRUC and DB2 table CSQ.ADMIN_B_STRUCTURE inconsistent
CSQU575E Structure APPL2 in table entry 4 not set in usage map
CSQU576E Structure 55 in usage map has no entry in table
CSQU572E Usage map T03_LH_ARRAY and DB2 table CSQ.OBJ_B_QUEUE inconsistent
CSQU577E Queue MYSQ in table entry 13 not set in usage map for structure APPL1
CSQU576E Queue 129 in usage map for structure APPL1 has no entry in table
CSQU528I Disconnected from DB2 DF4A
CSQU148I CSQ5PQSG Utility completed, return code=12
```

## Does the problem affect specific parts of the network?

Network problems can cause related problems for MQ for z/OS. Use this topic to review possible sources of networks problems.

You might be able to identify specific parts of the network that are affected by the problem (remote queues, for example). If the link to a remote queue manager is not working, the messages cannot flow to a target queue on the target queue manager. Check that the connection between the two systems is available, and that the channel initiator and listener have been started. Use the MQSC PING CHANNEL command to check the connection.

Check that messages are reaching the transmission queue, and check the local queue definition of the transmission queue, and any remote queues. Use the MQSC BYTSENT keyword of the DISPLAY CHSTATUS command to check that data is flowing along the channel. Use DISPLAY QLOCAL (XMITQ) CURDEPTH to check whether there are messages to be sent on the transmission queue. Check for diagnostic messages at both ends of the channel informing you that messages have been sent to the dead-letter queue.

If you are using WebSphere MQ clusters, check that the clustering definitions have been set up correctly.

Have you made any network-related changes that might account for the problem?

Have you changed any WebSphere MQ definitions, or any CICS or IMS definitions? Check the triggering attributes of the transmission queue.

## Problems that occur at specific times of the day or affect specific users

Use this topic to review IBM WebSphere MQ problems that occur at specific times of the day or specific groups of users.

If the problem occurs at specific times of day, it might be that it is dependent on system loading. Typically, peak system loading is at mid-morning and mid-afternoon, and so these periods are the times when load-dependent problems are most likely to occur. (If your network extends across more than one time zone, peak system loading might seem to occur at some other time of day.)

If you think that your IBM WebSphere MQ for z/OS system has a performance problem, see “Dealing with performance problems on z/OS” on page 1293.

If the problem only affects some users, is it because some users do not have the correct security authorization? See “User IDs for security checking” on page 568 for information about user IDs checked by IBM WebSphere MQ for z/OS.

## Is the problem intermittent or does the problem occur with all z/OS, CICS, or IMS systems?

Review this topic to consider if problems are caused by application interaction or are related to other z/OS systems.

An intermittent problem could be caused by failing to take into account the fact that processes can run independently of each other. For example, a program might issue an MQGET call, without specifying WAIT, before an earlier process has completed. You might also encounter this type of problem if your application tries to get a message from a queue while it is in sync point (that is, before it has been committed).

If the problem only occurs when you access a particular z/OS, IMS, or CICS system, consider what is different about this system. Also consider whether any changes have been made to the system that might affect the way it interacts with WebSphere MQ.

## Has the application run successfully before?

Application errors can often be determined by determining if they have run successfully before or if they have produced error messages and unexpected return codes.

If the problem appears to involve one particular application, consider whether the application has run successfully before.

Before you answer Yes to this question, consider:

### Have any changes been made to the application since it last ran successfully?

If so, it is likely that the error lies somewhere in the new or modified part of the application. Investigate the changes and see if you can find an obvious reason for the problem.

### Have all the functions of the application been fully exercised before?

Did problem occur when part of the application that had never been started before was used for the first time? If so, it is likely that the error lies in that part of the application. Try to find out what the application was doing when it failed, and check the source code in that part of the program for errors.

If a program has been run successfully on many previous occasions, check the current queue status and files that were being processed when the error occurred. It is possible that they contain some unusual data value that causes a rarely used path in the program to be invoked.

### Does the application check all return codes?

Has your system has been changed, perhaps in a minor way. Check the return codes your application receives as a result of the change. For example:

- Does your application assume that the queues it accesses can be shared? If a queue has been redefined as exclusive, can your application deal with return codes indicating that it can no longer access that queue?
- Have any security profiles been altered? An `MQOPEN` call might fail because of a security violation; can your application recover from the resulting return code?

### Does the application expect particular message formats?

If a message with an unexpected message format has been put onto a queue (for example, a message from a queue manager on a different platform), it might require data conversion or another different form of processing.

### Does the application run on other IBM WebSphere MQ for z/OS systems?

Is something different about the way that this queue manager is set up that is causing the problem? For example, have the queues been defined with the same maximum message length, or default priority?

### Does the application use the MQSET call to change queue attributes?

Is the application is designed to set a queue to have no trigger, then process some work, then set the queue to have a trigger? The application might have failed before the queue had been reset to have a trigger.

### Does the application handle messages that cause an application to fail?

If an application fails because of a corrupted message, the message retrieved is rolled back. The next application might get the same message and fail in the same way. Ensure that applications use the backout count; when the backout count threshold has been reached, the message in question is put onto the backout queue.

If your application has never run successfully before, examine your application carefully to see if you can find any of the following errors:

### Translation and compilation problems

Before you look at the code, examine the output from the translator, the compiler or assembler, and the linkage editor, to see if any errors have been reported. If your application fails to

translate, compile/assemble, or link edit into the load library, it also fails to run if you attempt to invoke it. See  *Developing applications (WebSphere MQ V7.1 Programming Guide)* for information about building your application, and for examples of the job control language (JCL) statements required.

### **Batch and TSO programs**

For batch and TSO programs, check that the correct stub has been included. There is one batch stub and two RRS stubs. If you are using RRS, check that you are not using the MQCMIT and MQBACK calls with the CSQBRSTB stub. Use the CSQBRRSI stub if you want to continue using these calls with RRS.

### **CICS programs**

For CICS programs, check that the program, the IBM WebSphere MQ CICS stub, and the CICS stub have been linked in the correct order. Also, check that your program or transaction is defined to CICS.

### **IMS programs**

For IMS programs, check that the link includes the program, the IBM WebSphere MQ stub, and the IMS language interface module. Ensure that the correct entry point has been specified. A program that is loaded dynamically from an IMS program must have the stub and language interface module linked also if it is to use IBM WebSphere MQ.

### **Possible code problems**

If the documentation shows that each step was accomplished without error, consider the coding of the application. Do the symptoms of the problem indicate the function that is failing and, therefore, the piece of code in error? See "Do you have a program error?" on page 1181 for some examples of common errors that cause problems with IBM WebSphere MQ applications.

### **Do applications report errors from IBM WebSphere MQ?**

For example, a queue might not be enabled for "gets". It receives a return code specifying this condition but does not report it. Consider where your applications report any errors or problems.

## **Have any changes been made since the last successful run?**

Recent changes made since the last successful run are often the source of unexpected errors. This topic contains information about some of the changes which can be investigated as part of your problem determination.

When you are considering changes that might recently have been made, think about WebSphere MQ, and also about the other programs it interfaces with, the hardware, and any new applications. Consider also the possibility that a new application that you don't yet know about might have been run on the system.

### **Has your initialization procedure been changed?**

Consider whether that might be the cause of the problem. Have you changed any data sets, or changed a library definition? Has z/OS been initialized with different parameters? In addition, check for error messages sent to the console during initialization.

### **Have you changed any queue definitions or security profiles?**

Consider whether some of your queues have been altered so that they are members of a cluster. This change might mean that messages arrive from different sources (for example, other queue managers or applications).

### **Have you changed any definitions in your sysplex that relate to the support and implementation of shared queues?**

Consider the effect that changes to such definitions as your sysplex couple data set, or Coupling Facility resource management policy. These changes might have on the operation of shared queues. Also, consider the effect of changes to the Db2 data sharing environment.

### **Has any of the software on your z/OS system been upgraded to a later release?**

Consider whether there are any necessary post-installation or migration activities that you need to perform.

### **Has your z/OS subsystem name table been changed?**

Changes to levels of corequisite software like z/OS or LE might require additional changes to WebSphere MQ.

**Do your applications deal with return codes that they might get as a result of any changes you have made?** Ensure that your applications deal with any new return codes that you introduce.

## **Do you have a program error?**

Use this topic to investigate if a program error is causing a WebSphere MQ problem.

The examples that follow illustrate the most common causes of problems encountered while running WebSphere MQ programs. Consider the possibility that the problem with your system could be caused by one of these errors.

- Programs issue MQSET to change queue attributes and fail to reset attributes of a queue. For example, setting a queue to NOTRIGGER.
- Making incorrect assumptions about the attributes of a queue. This assumption could include assuming that queues can be opened with MQOPEN when they are MQOPEN-exclusive, and assuming that queues are not part of a cluster when they are.
- Trying to access queues and data without the correct security authorization.
- Linking a program with no stub, or with the wrong stub (for example, a TSO program with the CICS stub). This can cause either a long-running unit of work, or an X'0C4' or other abend.
- Passing incorrect or invalid parameters in an MQI call; if the wrong number of parameters are passed, no attempt can be made to complete the completion code and reason code fields, and the task is abended. (This is an X'0C4' abend.)

This problem might occur if you attempt to run an application on an earlier version of MQSeries than it was written for, where some of the MQI values are invalid.

- Failing to define the WebSphere MQ modules to z/OS correctly (this error causes an X'0C4' abend in CSQYASCP).
- Failing to check return codes from MQI requests.

This problem might occur if you attempt to run an application on a later version of WebSphere MQ than it was written for, where new return codes have been introduced that are not checked for.

- Failing to open objects with the correct options needed for later MQI calls, for example using the MQOPEN call to open a queue but not specifying the correct options to enable the queue for subsequent MQGET calls.
- Failing to initialize *MsgId* and *CorrelId* correctly.

This error is especially true for MQGET.

- Using incorrect addresses.
- Using storage before it has been initialized.
- Passing variables with incorrect lengths specified.
- Passing parameters in the wrong order.
- Failing to define the correct security profiles and classes to RACF.

This might stop the queue manager or prevent you from carrying out any productive work.

- Relying on default MQI options for a ported application.

For example, z/OS defaults to MQGET and MQPUT in sync point. The distributed-platform default is out of sync point.

- Relying on default behavior at a normal or abnormal end of a portal application.

On z/OS, a normal end does an implicit MQCMIT and an abnormal end does an implicit rollback.

## Has there been an abend?

Use this topic to investigate common causes of abends and the different types of abend that can cause problems.

If your application has stopped running, it can be caused by an abnormal termination (abend).

You are notified of an abend in one of the following places, depending on what type of application you are using:

**Batch** Your listing shows the abend.

**CICS** You see a CICS transaction abend message. If your task is a terminal task, this message is displayed on your screen. If your task is not attached to a terminal, the message is displayed on the CICS CSMT log.

**IMS** In all cases, you see a message at the IBM WebSphere MQ for IMS master terminal and in the listing of the dependent region involved. If an IMS transaction that had been entered from a terminal was being processed, an error message is also sent to that terminal.

**TSO** You might see a TSO message with a return code on your screen. (Whether this message is displayed depends on the way your system is set up, and the type of error.)

## Common causes of abends

Abends can be caused by the user ending the task being performed before it terminates normally; for example, if you purge a CICS transaction. Abends can also be caused by an error in an application program.

## Address space dumps and transaction dumps

For some abends, an address space dump is produced. For CICS transactions, a transaction dump showing the storage areas of interest to the transaction is provided.

- If an application passes some data, the address of which is no longer valid, a dump is sometimes produced in the address space of the user.

**Note:** For a batch dump, the dump is formatted and written to SYSUDUMP. For information about SYSUDUMPs, see “SYSUDUMP information” on page 1290. For CICS, a system dump is written to the SYS1.DUMP data sets, as well as a transaction dump being taken.

- If a problem with IBM WebSphere MQ for z/OS itself causes an abend, an abend code of X'5C6' or X'6C6' is returned, along with an abend reason code. This reason code uniquely describes the cause of the problem. See “WebSphere MQ for z/OS abends” on page 1269 for information about the abend codes, and see  Return codes (*WebSphere MQ V7.1 Reference*) for an explanation of the reason code.

## Abnormal program termination

If your program has terminated abnormally, see “Dealing with program abends” on page 1270.

If your system has terminated abnormally, and you want to analyze the dump produced, see “WebSphere MQ dumps” on page 1276. This section tells you how to format the dump, and how to interpret the data contained in it.

## Have you obtained incorrect output?

Use this topic to review any incorrect output you have received.

If you have obtained what you believe to be some incorrect output, consider the following:

### Classifying incorrect output

“Incorrect output” might be regarded as any output that you were not expecting. However, use this term with care in the context of problem determination because it might be a secondary effect of some other type of error. For example, looping could be occurring if you get any repetitive output, even though that output is what you expected.

### Error messages

WebSphere MQ also responds to many errors it detects by sending error messages. You might regard these messages as “incorrect output”, but they are only symptoms of another type of problem. If you have received an error message from WebSphere MQ that you were not expecting, refer to “Are there any error messages, return codes or other error conditions?” on page 1172.

### Unexpected messages

If your application has not received a message that it was expecting, has received a message containing unexpected or corrupted information, or has received a message that it was not expecting (for example, one that was destined for a different application), refer to “Dealing with incorrect output” on page 1299.

## Can you reproduce the problem?

Reproducing the problem can be used to assist problem determination for WebSphere MQ for z/OS. Use this topic to further isolate the type of problem reproduction.

If you can reproduce the problem, consider the conditions under which you can reproduce it. For example:

### Is it caused by a command?

If so, is the command issued from the z/OS console, from CSQUTIL, from a program written to put commands onto the SYSTEM.COMMAND.INPUT queue, or by using the operations and control panels?

### Does the command work if it is entered by another method?

If the command works when it is entered at the console, but not otherwise, check that the command server has not stopped, and that the queue definition of the SYSTEM.COMMAND.INPUT queue has not been changed.

### Is the command server running?

Issue the command DIS CMDSERV to check.

### Is it caused by an application?

If so, does it fail in CICS, IMS, TSO, or batch?

Does it fail on all WebSphere MQ systems, or only on some?

### Is an application causing the problem?

Can you identify any application that always seems to be running in the system when the problem occurs? If so, examine the application to see if it is in error.

## Have you failed to receive a response from an MQSC command?

Use this topic for investigating problems where you fail to receive a response from an MQSC command.

If you have issued an MQSC command from an application (and not from a z/OS console), but you have not received a response, consider the subsequent questions:

### Is the command server running?

Check that the command server is running, as follows:

1. Use the DISPLAY CMDSERV command at the z/OS console to display the status of the command server.
2. If the command server is not running, start it using the START CMDSERV command.
3. If the command server is running, issue the DISPLAY QUEUE command. Use the name of the system-command input queue and the CURDEPTH and MAXDEPTH attributes to define the data displayed.

If these values show that the queue is full, and the command server has been started, the messages are not being read from the queue.

4. Try stopping the command server and then restarting it, responding to any error messages that are produced.
5. Issue the display command again to see if it is working now.

### Has a reply been sent to the dead-letter queue?

Use the DISPLAY QMGR DEADQ command to find out the name of the system dead-letter queue (if you do not know what it is).

Use this name in the DISPLAY QUEUE command with the CURDEPTH attribute to see if there are any messages on the queue.

The dead-letter queue message header (dead-letter header structure) contains a reason or feedback code describing the problem. (See  Reason (MQLONG) (*WebSphere MQ V7.1 Reference*) for information about the dead-letter header structure.)

### Are the queues enabled for PUTs and GETs?

Use the DISPLAY QUEUE command from the console to check, for example, DISPLAY QUEUE(SYSTEM.COMMAND.INPUT) PUT GET.

### Is the WaitInterval parameter set to a sufficiently long time?

If your MQGET call has timed out, your application receives completion code of 2 and a reason code of 2033 (MQRC\_NO\_MSG\_AVAILABLE). (See  WaitInterval (MQLONG) (*WebSphere MQ*

*V7.1 Reference*) and  MQGET - Get message (*WebSphere MQ V7.1 Reference*) for information about the *WaitInterval* parameter, and completion and reason codes from MQGET.)

### Is a sync point required?

If you are using your own application program to put commands onto the system-command input queue, consider whether you must take a sync point.

You must take a sync point after putting messages to a queue, and before attempting to receive reply messages, or use MQPMO\_NO\_SYNCPOINT when putting them. Unless you have excluded your request message from sync point, you must take a sync point before attempting to receive reply messages.

### Are the MaxDepth and MaxMsgL parameters of your queues set sufficiently high?

See  CSQO016E: MQPUT to q-name unsuccessful. Reason code=mqrc. (*WebSphere MQ V7.1 Reference*) for information about defining the system-command input queue and the reply-to queue.

### **Are you using the *CorrelId* and *MsgId* parameters correctly?**

You must identify the queue and then display the CURDEPTH. Use the DISPLAY QUEUE command from the console (for example, DISPLAY QUEUE (MY.REPLY.QUEUE) CURDEPTH), to see if there are messages on the reply-to queue that you have not received.

Set the values of *MsgId* and *CorrelId* in your application to ensure that you receive all messages from the queue.

The following questions are applicable if you have issued an MQSC command from either a z/OS console (or its equivalent), or an application, but have not received a response:

#### **Is the queue manager still running, or did your command cause an abend?**

Look for error messages indicating an abend, and if one occurred, see “WebSphere MQ dumps” on page 1276.

#### **Were any error messages issued?**

Check to see if any error messages were issued that might indicate the nature of the error.

See “Issuing commands” on page 238 for information about the different methods you can use to enter MQSC commands.

### **Is your application or IBM WebSphere MQ for z/OS running slowly?**

Slow applications can be caused by the application itself or underlying software including WebSphere MQ. Use this topic for initial investigations into slow applications.

If your application is running slowly, this could indicate that it is in a loop, or waiting for a resource that is not available.

#### **Is the problem worse at peak system load times?**

This could also be caused by a performance problem. Perhaps it is because your system needs tuning, or because it is operating near the limits of its capacity. This type of problem is probably worst at peak system load times, typically at mid-morning and mid-afternoon. (If your network extends across more than one time zone, peak system load might seem to you to occur at some other time.)

#### **Does the problem occur when the system is lightly loaded?**

If you find that degrading performance is not dependent on system loading, but happens sometimes when the system is lightly loaded, a poorly designed application program is probably to blame. This could manifest itself as a problem that only occurs when specific queues are accessed.

#### **Is IBM WebSphere MQ for z/OS running slowly?**

The following symptoms might indicate that IBM WebSphere MQ for z/OS is running slowly:

- If your system is slow to respond to commands.
- If repeated displays of the queue depth indicate that the queue is being processed slowly for an application with which you would expect a large amount of queue activity.

You can find guidance on dealing with waits and loops in “Dealing with applications that are running slowly or have stopped on z/OS” on page 1294, and on dealing with performance problems in “Dealing with performance problems on z/OS” on page 1293.

---

## Dealing with problems

Learn how to resolve some of the typical problems that can occur.

There are some initial checks that you can make that may provide answers to common problems that you may have. Carry out the initial checks for your platform:

- “Making initial checks on Windows, UNIX and Linux systems” on page 1151
- “Making initial checks on z/OS” on page 1169
- “Making initial checks on IBM i” on page 1161

You can use the information acquired from the following locations to help you rectify the problem:

- Logs, see “Using logs” on page 1226
- Trace, see “Using trace” on page 1234
- Dumps, see “Problem determination on z/OS” on page 1265
- z/OS, see “Problem determination on z/OS” on page 1265

Use the following topics to help you solve specific problems:

- “Resolving problems with commands”
- “Resolving problems with queue managers” on page 1187
- “Resolving problems with channels and DQM” on page 1305
- “Resolving problems with queue manager clusters” on page 1187
- “Resolving problems with undelivered messages” on page 88
- “Resolving problems with IBM WebSphere MQ MQI clients” on page 1200

See the following topics for resolving specific problems on IBM WebSphere MQ for z/OS:

- “Dealing with performance problems on z/OS” on page 1293
- “Dealing with incorrect output” on page 1299

### Related concepts:

“Troubleshooting and support” on page 1149  
“Making initial checks on z/OS” on page 1169  
“Making initial checks on IBM i” on page 1161  
“Dealing with problems”  
“IBM Support Assistant (ISA)” on page 1324  
“Reason codes” on page 1379

 Troubleshooting and support reference (*WebSphere MQ V7.1 Reference*)

### Related tasks:

“Searching knowledge bases” on page 1326  
“Contacting IBM Software Support” on page 1346

### Related reference:

“PCF reason codes” on page 1590

## Resolving problems with commands

- **Scenario:** You receive errors when you use special characters in descriptive text for some commands.
- **Explanation:** Some characters, for example, back slash (\) and double quote (") characters have special meanings when used with commands.

- **Solution:** Precede special characters with a \, that is, enter \\ or \" if you want \ or " in your text. Not all characters are allowed to be used with commands. For more information about characters with special meanings and how to use them, see  Characters with special meanings (*WebSphere MQ V7.1 Reference*).

## Resolving problems with queue managers

Use the advice given here to help you to resolve common problems that can arise when you use queue managers.

### Queue manager unavailable error

- **Scenario:** You receive a *queue manager unavailable* error.
- **Explanation:** Configuration file errors typically prevent queue managers from being found, and result in *queue manager unavailable* errors. On Windows, problems in the `qm.ini` file can cause *queue manager unavailable* errors when a queue manager is started.
- **Solution:** Ensure that the configuration files exist, and that the IBM WebSphere MQ configuration file references the correct queue manager and log directories. On Windows, check for problems in the `qm.ini` file.

## Resolving problems with queue manager clusters

Use the advice given here to help you to resolve common problems that can arise when you use queue manager clusters.

- "A cluster-sender channel is continually trying to start" on page 1188
- "DISPLAY CLUSQMGR shows CLUSQMGR names starting SYSTEM.TEMP." on page 1189
- "Return code=2035 MQRC\_NOT\_AUTHORIZED" on page 1189
- "Return code=2085 MQRC\_UNKNOWN\_OBJECT\_NAME when trying to open a queue in the cluster" on page 1190
- "Return code=2189 MQRC\_CLUSTER\_RESOLUTION\_ERROR when trying to open a queue in the cluster" on page 1191
- "Return code=2082 MQRC\_UNKNOWN\_ALIAS\_BASE\_Q opening a queue in the cluster" on page 1191
- "Messages are not arriving on the destination queues" on page 1192
- "Messages put to a cluster alias queue go to SYSTEM.DEAD.LETTER.QUEUE" on page 1192
- "A queue manager has out of date information about queues and channels in the cluster" on page 1193
- "No changes in the cluster are being reflected in the local queue manager" on page 1193
- "DISPLAY CLUSQMGR displays a queue manager twice" on page 1194
- "A queue manager does not rejoin the cluster" on page 1194
- "Out of date information in a restored cluster" on page 1195
- "Cluster queue manager force removed from a full repository by mistake" on page 1195
- "Possible repository messages deleted" on page 1196
- "Two full repositories moved at the same time" on page 1196
- "Unknown state of a cluster" on page 1197
- "What happens when a cluster queue manager fails" on page 1198
- "What happens when a repository fails" on page 1198
- "What happens if a cluster queue is disabled for MQPUT" on page 1199

## Related concepts:

“Troubleshooting and support” on page 1149

“Making initial checks on Windows, UNIX and Linux systems” on page 1151

“Making initial checks on z/OS” on page 1169

“Making initial checks on IBM i” on page 1161

“Reason codes” on page 1379

## Related tasks:

 [Configuring a queue manager cluster \(\*WebSphere MQ V7.1 Installing Guide\*\)](#)

## A cluster-sender channel is continually trying to start

Check the queue manager and listener are running, and the cluster-sender and cluster-receiver channel definitions are correct.

### Symptom

```
1 : display chs(*)
AMQ8417: Display Channel Status details.
CHANNEL(DEMO.QM2) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
CONNNAME(computer.ibm.com(1414))
CURRENT CHLTYPE(CLUSSDR)
STATUS(RETRYING)
```

### Cause

1. The remote queue manager is not available.
2. An incorrect parameter is defined either for the local manual cluster-sender channel or the remote cluster-receiver channel.

### Solution

Check whether the problem is the availability of the remote queue manager.

1. Are there any error messages?
2. Is the queue manager active?
3. Is the listener running?
4. Is the cluster-sender channel able to start?

If the remote queue manager is available, is there a problem with a channel definition? Check the definition type of the cluster queue manager to see if the channel is continually trying to start; for example:

```
1 : dis clusqmgr(*) deftype where(channel eq DEMO.QM2)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM2) CHANNEL(DEMO.QM2) CLUSTER(DEMO)
DEFTYPE(CLUSSDRA)
```

If the definition type is CLUSSDR the channel is using the local manual cluster-sender definition. Alter any incorrect parameters in the local manual cluster-sender definition and restart the channel.

If the definition type is either CLUSSDRA or CLUSSDRB the channel is using an auto-defined cluster-sender channel. The auto-defined cluster-sender channel is based on the definition of a remote cluster receiver channel. Alter any incorrect parameters in the remote cluster receiver definition. For example, the connname parameter might be incorrect:

```
1 : alter chl(demo.qm2) chltype(clusrcvr) connname('newhost(1414)')
AMQ8016: WebSphere MQ channel changed.
```

Changes to the remote cluster-receiver definition are propagated out to any cluster queue managers that are interested. The corresponding auto-defined channels are updated accordingly. You can check that the updates have been propagated correctly by checking the changed parameter. For example:

```
1 : dis clusqmgr(qm2) conname
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM2) CHANNEL(DEMO.QM2) CLUSTER(DEMO) CONNAME(newhost(1414))
```

If the auto-defined definition is now correct, restart the channel.

## **DISPLAY CLUSQMGR shows CLUSQMGR names starting SYSTEM.TEMP.**

The queue manager has not received any information from the full repository queue manager that the manually defined CLUSSDR channel points to. Check that the cluster channels are defined correctly.

### **Symptom**

```
1 : display clusqmgr(*)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(SYSTEM.TEMPUUID.computer.hursley.ibm.com(1414))
CLUSQMGR(QM2) CHANNEL(DEMO.QM2) CLUSTER(DEMO)
```

### **Cause**

The queue manager has not received any information from the full repository queue manager that the manually defined CLUSSDR channel points to. The manually defined CLUSSDR channel must be in running state.

### **Solution**

Check that the CLUSRCVR definition is also correct, especially its CONNAME and CLUSTER parameters. Alter the channel definition, if the definition is wrong.

It might take some time for the remote queue managers to attempt a new restart, and start their channels with the corrected definition.

## **Return code=2035 MQRC\_NOT\_AUTHORIZED**

The RC2035 reason code is displayed for various reasons including an error on opening a queue or a channel, an error received when you attempt to use a user ID that has administrator authority, an error when using an IBM WebSphere MQ JMS application, and opening a queue on a cluster. MQS\_REPORT\_NOAUTH and MQSAUTHERRORS can be used to further diagnose RC2035.

### **Specific problems**

See "Specific problems generating RC2035" on page 1398 for information on:

- JMSWQM2013 invalid security authentication
- MQRC\_NOT\_AUTHORIZED on a queue or channel
- MQRC\_NOT\_AUTHORIZED (AMQ4036 on a client) as an administrator
- MQS\_REPORT\_NOAUTH and MQSAUTHERRORS environment variables

### **Opening a queue in a cluster**

The solution for this error depends on whether the queue is on z/OS or not. On z/OS use your security manager. On other platforms create a local alias to the cluster queue, or authorize all users to have access to the transmission queue.

## Symptom

Applications receive a return code of 2035 MQRC\_NOT\_AUTHORIZED when trying to open a queue in a cluster.

## Cause

Your application receives the return code of MQRC\_NOT\_AUTHORIZED when trying to open a queue in a cluster. The authorization for that queue is correct. It is likely that the application is not authorized to put to the cluster transmission queue.

## Solution

The solution depends on whether the queue is on z/OS or not. See the related information topic.

## Return code=2085 MQRC\_UNKNOWN\_OBJECT\_NAME when trying to open a queue in the cluster

### Symptom

Applications receive a return code of 2085 MQRC\_UNKNOWN\_OBJECT\_NAME when trying to open a queue in the cluster.

### Cause

The queue manager where the object exists or this queue manager might not have successfully entered the cluster.

### Solution

Make sure that they can each display all the full repositories in the cluster. Also make sure that the CLUSSDR channels to the full repositories are trying to start.

If the queue is in the cluster, check that you have used appropriate open options. You cannot get messages from a remote cluster queue, so make sure that the open options are for output only.

```
1 : display clusqmgr(*) qmtype status
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) QMTYPE(NORMAL)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM2) CLUSTER(DEMO)
CHANNEL(DEMO.QM2) QMTYPE(REPOS)
STATUS(RUNNING)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM3) CLUSTER(DEMO)
CHANNEL(DEMO.QM3) QMTYPE(REPOS)
STATUS(RUNNING)
```

## Return code=2189 MQRC\_CLUSTER\_RESOLUTION\_ERROR when trying to open a queue in the cluster

Make sure that the CLUSSDR channels to the full repositories are not continually trying to start.

### Symptom

Applications receive a return code of 2189 MQRC\_CLUSTER\_RESOLUTION\_ERROR when trying to open a queue in the cluster.

### Cause

The queue is being opened for the first time and the queue manager cannot contact any full repositories.

### Solution

Make sure that the CLUSSDR channels to the full repositories are not continually trying to start.

```
1 : display clusqmgr(*) qmtype status
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) QMTYPE(NORMAL)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM2) CLUSTER(DEMO)
CHANNEL(DEMO.QM2) QMTYPE(REPOS)
STATUS(RUNNING)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM3) CLUSTER(DEMO)
CHANNEL(DEMO.QM3) QMTYPE(REPOS)
STATUS(RUNNING)
```

## Return code=2082 MQRC\_UNKNOWN\_ALIAS\_BASE\_Q opening a queue in the cluster

Applications get rc=2082 MQRC\_UNKNOWN\_ALIAS\_BASE\_Q when trying to open a queue in the cluster.

### Problem

An MQOPEN or MQPUT1 call was issued specifying an alias queue as the target, but the *BaseQName* in the alias queue attributes is not recognized as a queue name.

This reason code can also occur when *BaseQName* is the name of a cluster queue that cannot be resolved successfully.

MQRC\_UNKNOWN\_ALIAS\_BASE\_Q might indicate that the application is specifying the **ObjectQmgrName** of the queue manager that it is connecting to, and the queue manager that is hosting the alias queue. This means that the queue manager looks for the alias target queue on the specified queue manager and fails because the alias target queue is not on the local queue manager.

### Solution

The **ObjectQmgrName** parameter is typically left blank so that the clustering decides which queue manager to route to.

## Messages are not arriving on the destination queues

Make sure that the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` is empty and also that the channel to the destination queue manager is running.

### Symptom

Messages are not arriving on the destination queues.

### Cause

The messages might be stuck at their origin queue manager.

### Solution

Make sure that the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` is empty and also that the channel to the destination queue manager is running.

```
1 : display ql(SYSTEM.CLUSTER.TRANSMIT.QUEUE) curdepth
AMQ8409: Display Queue details.
QUEUE(SYSTEM.CLUSTER.TRANSMIT.QUEUE) CURDEPTH(0)
2 : display chs(DEMO.QM2)
AMQ8417: Display Channel Status details.
CHANNEL(DEMO.QM2) XMITQ(SYSTEM.CLUSTER.TRANSMIT.QUEUE)
CONNNAME(comfrey.hursley.ibm.com(1415))
CURRENT CHLTYPE(CLUSSDR)
STATUS(RUNNING)
```

## Messages put to a cluster alias queue go to `SYSTEM.DEAD.LETTER.QUEUE`

A cluster alias queue resolves to a local queue that does not exist.

### Symptom

Messages put to an alias queue go to `SYSTEM.DEAD.LETTER.QUEUE` with reason `MQRC_UNKNOWN_ALIAS_BASE_Q`.

### Cause

A message is routed to a queue manager where a clustered alias queue is defined. A local target queue is not defined on that queue manager. Because the message was put with the `MQOO_BIND_ON_OPEN` open option, the queue manager cannot requeue the message.

When `MQOO_BIND_ON_OPEN` is used, the cluster queue alias is firmly bound. The resolved name is the name of the target queue and any queue manager on which the cluster queue alias is defined. The queue manager name is placed in the transmission queue header. If the target queue does not exist on the queue manager to which the message is sent, the message is put on the dead letter queue. The destination is not recomputed, because the transmission header contains the name of the target queue manager resolved by `MQOO_BIND_ON_OPEN`. If the alias queue had been opened with `MQOO_BIND_NOT_FIXED`, then the transmission queue header would contain a blank queue manager name, and the destination would be recomputed. In which case, if the local queue is defined elsewhere in the cluster, the message would be sent there.

### Solution

1. Change all alias queue definitions to specify `DEFBIND(NOTFIXED)`.
2. Use `MQOO_BIND_NOT_FIXED` as an open option when the queue is opened.
3. If you specify `MQOO_BIND_ON_OPEN`, ensure that a cluster alias that resolves to a local queue defined on the same queue manager as the alias.

## A queue manager has out of date information about queues and channels in the cluster

### Symptom

DISPLAY QCLUSTER and DISPLAY CLUSQMGR show objects which are out of date.

### Cause

Updates to the cluster only flow between the full repositories over manually defined CLUSSDR channels. After the cluster has formed CLUSSDR channels display as DEFTYPE(CLUSSDRB) channels because they are both manual and automatic channels. There must be enough CLUSSDR channels to form a complete network between all the full repositories.

### Solution

- Check that the queue manager where the object exists and the local queue manager are still connected to the cluster.
- Check that each queue manager can display all the full repositories in the cluster.
- Check whether the CLUSSDR channels to the full repositories are continually trying to restart.
- Check that the full repositories have enough CLUSSDR channels defined to correctly connect them together.

```
1 : dis clusqmgr(QM1) CHANNEL(*) STATUS DEFTYPE QMTYPE
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) DEFTYPE(CLUSSDRA)
QMTYPE(NORMAL) STATUS(RUNNING)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM2) CLUSTER(DEMO)
CHANNEL(DEMO.QM2) DEFTYPE(CLUSRCVR)
QMTYPE(REPOS)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM3) CLUSTER(DEMO)
CHANNEL(DEMO.QM3) DEFTYPE(CLUSSDRB)
QMTYPE(REPOS) STATUS(RUNNING)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM4) CLUSTER(DEMO)
CHANNEL(DEMO.QM4) DEFTYPE(CLUSSDRA)
QMTYPE(NORMAL) STATUS(RUNNING)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM4)
```

### No changes in the cluster are being reflected in the local queue manager

The repository manager process is not processing repository commands, possibly because of a problem with receiving or processing messages in the command queue.

### Symptom

No changes in the cluster are being reflected in the local queue manager.

### Cause

The repository manager process is not processing repository commands.

### Solution

1. Check that the SYSTEM.CLUSTER.COMMAND.QUEUE is empty.

- ```
1 : display ql(SYSTEM.CLUSTER.COMMAND.QUEUE) curdepth
AMQ8409: Display Queue details.
QUEUE(SYSTEM.CLUSTER.COMMAND.QUEUE) CURDEPTH(0)
```
2. Check that the channel initiator is running on z/OS.
 3. Check that there are no error messages in the error logs indicating the queue manager has a temporary resource shortage.

DISPLAY CLUSQMGR displays a queue manager twice

Use the RESET CLUSTER command to remove all traces of an old instance of a queue manager.

```
1 : display clusqmgr(QM1) qmid
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) QMID(QM1_2002-03-04_11.07.01)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) QMID(QM1_2002-03-04_11.04.19)
```

The cluster functions correctly with the older version of the queue manager being ignored, until it ages out of the cluster completely after about 90 days.

Cause

1. The queue manager might have been deleted, and then recreated and redefined.
2. The queue manager might have been cold-started on z/OS, without first following the procedure to remove a queue manager from a cluster; see  Removing a queue manager from a cluster (*WebSphere MQ V7.1 Installing Guide*).

Solution

To remove all trace of the queue manager immediately use the RESET CLUSTER command from a full repository queue manager. The command removes the older unwanted queue manager and its queues from the cluster.

```
2 : reset cluster(DEMO) qmid('QM1_2002-03-04_11.04.19') action(FORCEREMOVE) queues(yes)
AMQ8559: RESET CLUSTER accepted.
```

Using the RESET CLUSTER command stops auto-defined cluster sender channels for the affected queue manager. You must manually restart any cluster sender channels that are stopped, after completing the RESET CLUSTER command.

A queue manager does not rejoin the cluster

After issuing a RESET or REFRESH cluster command the channel from the queue manager to the cluster might be stopped. Check the cluster channel status and restart the channel.

Symptom

A queue manager does not rejoin a cluster after issuing the RESET CLUSTER and REFRESH CLUSTER commands.

Cause

A side effect of the RESET and REFRESH commands might be that a channel is stopped. A channel is stopped in order that the correct version of the channel runs when RESET or REFRESH command is completed.

Solution

Check that the channels between the problem queue manager and the full repositories are running and use the `START CHANNEL` command if necessary.

Related concepts:

 Clustering: Using `REFRESH CLUSTER` best practices

Out of date information in a restored cluster

After restoring a queue manager, its cluster information is out of date. Refresh the cluster information with the `REFRESH CLUSTER` command.

Problem

After an image backup of QM1, a partial repository in cluster DEMO has been restored and the cluster information it contains is out of date.

Solution

On QM1, issue the command `REFRESH CLUSTER(DEMO)`.

Note: For large clusters, use of the `REFRESH CLUSTER` command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

QM1 removes all information it has about the cluster DEMO, except that relating to the cluster queue managers which are the full repositories in the cluster. Assuming that this information is still correct, QM1 contacts the full repositories. QM1 informs the full repositories about itself and its queues. It recovers the information for queues and queue managers that exist elsewhere in the cluster as they are opened.

Cluster queue manager force removed from a full repository by mistake

Restore the queue manager to the full repository by issuing the command `REFRESH CLUSTER` on the queue manager that was removed from the repository.

Problem

The command, `RESET CLUSTER(DEMO) QMNAME(QM1) ACTION(FORCEREMOVE)` was issued on a full repository in cluster DEMO by mistake.

Solution

On QM1, issue the command `REFRESH CLUSTER(DEMO)`.

Note: For large clusters, use of the `REFRESH CLUSTER` command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

Possible repository messages deleted

Messages destined for a queue manager were removed from the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` in other queue managers. Restore the information by issuing the `REFRESH CLUSTER` command on the affected queue manager.

Problem

Messages destined for QM1 were removed from the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` in other queue managers and they might have been repository messages.

Solution

On QM1, issue the command `REFRESH CLUSTER(DEMO)`.

Note: For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

QM1 removes all information it has about the cluster DEMO, except that relating to the cluster queue managers which are the full repositories in the cluster. Assuming that this information is still correct, QM1 contacts the full repositories. QM1 informs the full repositories about itself and its queues. It recovers the information for queues and queue managers that exist elsewhere in the cluster as they are opened.

Two full repositories moved at the same time

If you move both full repositories to new network addresses at the same time, the cluster is not updated with the new addresses automatically. Follow the procedure to transfer the new network addresses. Move the repositories one at a time to avoid the problem.

Problem

Cluster DEMO contains two full repositories, QM1 and QM2. They were both moved to a new location on the network at the same time.

Solution

1. Alter the `CONNAME` in the `CLUSRCVR` and `CLUSSDR` channels to specify the new network addresses.
2. Alter one of the queue managers (QM1 or QM2) so it is no longer a full repository for any cluster.
3. On the altered queue manager, issue the command `REFRESH CLUSTER(*) REPOS(YES)`.

Note: For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

4. Alter the queue manager so it is acting as a full repository.

Recommendation

You could avoid the problem as follows:

1. Move one of the queue managers, for example QM2, to its new network address.
2. Alter the network address in the QM2 `CLUSRCVR` channel.
3. Start the QM2 `CLUSRCVR` channel.
4. Wait for the other full repository queue manager, QM1, to learn the new address of QM2.

5. Move the other full repository queue manager, QM1, to its new network address.
6. Alter the network address in the QM1 CLUSRCVR channel.
7. Start the QM1 CLUSRCVR channel.
8. Alter the manually defined CLUSSDR channels for the sake of clarity, although at this stage they are not needed for the correct operation of the cluster.

The procedure forces QM2 to reuse the information from the correct CLUSSDR channel to re-establish contact with QM1 and then rebuild its knowledge of the cluster. Additionally, having once again contacted QM1, it is given its own correct network address based on the CONNAME in QM2 CLUSRCVR definition.

Unknown state of a cluster

Restore the cluster information in all the full repositories to a known state by rebuilding the full repositories from all the partial repositories in the cluster.

Problem

Under normal conditions the full repositories exchange information about the queues and queue managers in the cluster. If one full repository is refreshed, the cluster information is recovered from the other.

The problem is how to completely reset all the systems in the cluster to restore a known state to the cluster.

Solution

To stop cluster information being updated from the unknown state of the full repositories, all the CLUSRCVR channels to full repositories are stopped. The CLUSSDR channels change to inactive.

When you refresh the full repository systems, none of them are able to communicate, so they start from the same cleared state.

When you refresh the partial repository systems, they rejoin the cluster and rebuild it to the complete set of queue managers and queues. The cluster information in the rebuilt full is restored to a known state.

Note: For large clusters, use of the **REFRESH CLUSTER** command can be disruptive to the cluster while it is in progress, and again at 27 day intervals thereafter when the cluster objects automatically send status updates to all interested queue managers. See  Refreshing in a large cluster can affect performance and availability of the cluster.

1. On all the full repository queue managers, follow these steps:
 - a. Alter queue managers that are full repositories so they are no longer full repositories.
 - b. Resolve any in doubt CLUSSDR channels.
 - c. Wait for the CLUSSDR channels to become inactive.
 - d. Stop the CLUSRCVR channels.
 - e. When all the CLUSRCVR channels on all the full repository systems are stopped, issue the command `REFRESH CLUSTER(DEMO) REPOS(YES)`.
 - f. Alter the queue managers so they are full repositories.
 - g. Start the CLUSRCVR channels to re-enable them for communication.
2. On all the partial repository queue managers, follow these steps:
 - a. Resolve any in doubt CLUSSDR channels.
 - b. Make sure all CLUSSDR channels on the queue manager are stopped or inactive.
 - c. Issue the command `REFRESH CLUSTER(DEMO) REPOS(YES)`.

What happens when a cluster queue manager fails

When a cluster queue manager fails, some undelivered messages are sent to other queue managers in the cluster. Messages that are in-flight wait until the queue manager is restarted. Use a high-availability mechanism to restart a queue manager automatically.

Problem

If a message-batch is sent to a particular queue manager and that queue manager becomes unavailable, what happens at the sending queue manager?

Explanation

Except for non-persistent messages on an NPMSPEED(FAST) channel, the undelivered batch of messages is backed out to the cluster transmission queue on the sending queue manager. On an NPMSPEED(FAST) channel, non-persistent messages are not batched, and one might be lost.

- Indoubt messages, and messages that are bound to the unavailable queue manager, wait until the queue manager becomes available again.
- Other messages are delivered to alternative queue managers selected by the workload management routine.

Solution

The unavailable cluster queue manager can be restarted automatically, either by being configured as a multi-instance queue manager, or by a platform-specific high availability mechanism.

What happens when a repository fails

How you know a repository has failed and what to do to fix it?

Problem

1. Cluster information is sent to repositories (whether full or partial) on a local queue called `SYSTEM.CLUSTER.COMMAND.QUEUE`. If this queue fills up, perhaps because the queue manager has stopped working, the cluster-information messages are routed to the dead-letter queue.
2. The repository runs out of storage.

Solution

1. Monitor the messages on your queue manager log or z/OS system console to detect if `SYSTEM.CLUSTER.COMMAND.QUEUE` is filling up. If it is, you need to run an application to retrieve the messages from the dead-letter queue and reroute them to the correct destination.
2. If errors occur on a repository queue manager, messages tell you what error has occurred and how long the queue manager waits before trying to restart.
 - a. On WebSphere MQ for z/OS, the `SYSTEM.CLUSTER.COMMAND.QUEUE` is disabled for `MQGET`.
 - b. When you have identified and resolved the error, enable the `SYSTEM.CLUSTER.COMMAND.QUEUE` so that the queue manager can restart successfully.
3. In the unlikely event of the repository running out of storage, storage allocation errors are sent to the queue-manager log or z/OS system console. To fix the storage problem, stop and then restart the queue manager. When the queue manager is restarted, more storage is automatically allocated to hold all the repository information.

What happens if a cluster queue is disabled for MQPUT

All instances of a cluster queue that is being used for workload balancing might be disabled for MQPUT. Applications putting a message to the queue either receive a MQRC_CLUSTER_PUT_INHIBITED or a MQRC_PUT_INHIBITED return code. You might want to modify this behavior.

Problem

When a cluster queue is disabled for MQPUT, its status is reflected in the repository of each queue manager that is interested in that queue. The workload management algorithm tries to send messages to destinations that are enabled for MQPUT. If there are no destinations enabled for MQPUT and no local instance of a queue, an MQOPEN call that specified MQ00_BIND_ON_OPEN returns a return code of MQRC_CLUSTER_PUT_INHIBITED to the application. If MQ00_BIND_NOT_FIXED is specified, or there is a local instance of the queue, an MQOPEN call succeeds but subsequent MQPUT calls fail with return code MQRC_PUT_INHIBITED.

Solution

You can write a user exit program to modify the workload management routines so that messages can be routed to a destination that is disabled for MQPUT.

A message can arrive at a destination that is disabled for MQPUT. The message might have been in flight at the time the queue became disabled, or a workload exit might have chosen the destination explicitly. The workload management routine at the destination queue manager has a number of ways to deal with the message:

- Choose another appropriate destination, if there is one.
- Place the message on the dead-letter queue.
- Return the message to the originator, if there is no dead-letter queue

Resolving problems with undelivered messages

Use the advice given here to help you to resolve problems when messages do not delivered successfully.

- **Scenario:** Messages do not arrive on a queue when you are expecting them.
- **Explanation:** Messages that cannot be delivered for some reason are placed on the dead-letter queue.
- **Solution:** You can check whether the queue contains any messages by issuing an MQSC DISPLAY QUEUE command.

If the queue contains messages, you can use the provided browse sample application (amqsbcg) to browse messages on the queue using the MQGET call. The sample application steps through all the messages on a named queue for a named queue manager, displaying both the message descriptor and the message context fields for all the messages on the named queue.

You must decide how to dispose of any messages found on the dead-letter queue, depending on the reasons for the messages being put on the queue. Problems might occur if you do not associate a dead-letter queue with each queue manager.

For more information about dead-letter queues and handling undelivered messages, see  Handling undelivered messages with the WebSphere MQ dead-letter queue handler (*WebSphere MQ V7.1 Installing Guide*).

Resolving problems with IBM WebSphere MQ MQI clients

This collection of topics contains information about techniques for solving problems in IBM WebSphere MQ MQI client applications.

An application running in the IBM WebSphere MQ MQI client environment receives MQRC_* reason codes in the same way as IBM WebSphere MQ server applications. However, there are additional reason codes for error conditions associated with IBM WebSphere MQ MQI clients. For example:

- Remote machine not responding
- Communications line error
- Invalid machine address

The most common time for errors to occur is when an application issues an MQCONN or MQCONNX and receives the response MQRC_Q_MQR_NOT_AVAILABLE. Look in the client error log for a message explaining the failure. There might also be errors logged at the server, depending on the nature of the failure. Also, check that the application on the IBM WebSphere MQ MQI client is linked with the correct library file.

IBM WebSphere MQ MQI client fails to make a connection

An MQCONN or MQCONNX might fail because there is no listener program running on the server, or during protocol checking.

When the IBM WebSphere MQ MQI client issues an MQCONN or MQCONNX call to a server, socket and port information is exchanged between the IBM WebSphere MQ MQI client and the server. For any exchange of information to take place, there must be a program on the server with the role to 'listen' on the communications line for any activity. If there is no program doing this, or there is one but it is not configured correctly, the MQCONN or MQCONNX call fails, and the relevant reason code is returned to the IBM WebSphere MQ MQI client application.

If the connection is successful, IBM WebSphere MQ protocol messages are exchanged and further checking takes place. During the IBM WebSphere MQ protocol checking phase, some aspects are negotiated while others cause the connection to fail. It is not until all these checks are successful that the MQCONN or MQCONNX call succeeds.

For information about the MQRC_* reason codes, see API reason codes.

Stopping IBM WebSphere MQ MQI clients

Even though a IBM WebSphere MQ MQI client has stopped, it is still possible for the associated process at the server to be holding its queues open. The queues are not closed until the communications layer detects that the partner has gone.

If sharing conversations is enabled, the server channel is always in the correct state for the communications layer to detect that the partner has gone.

Error messages with IBM WebSphere MQ MQI clients

When an error occurs with a IBM WebSphere MQ MQI client system, error messages are put into the IBM WebSphere MQ system error files.

- On UNIX and Linux systems, these files are found in the /var/mqm/errors directory
- On Windows, these files are found in the errors subdirectory of the IBM WebSphere MQ MQI client installation. Usually this directory is C:\Program Files\IBM\WebSphere MQ\errors.
- On IBM i, these files are found in the /QIBM/UserData/mqm/errors directory

Certain client errors can also be recorded in the IBM WebSphere MQ error files associated with the server to which the client was connected.

Troubleshooting IBM WebSphere MQ client for HP Integrity NonStop Server

Provides information to help you to detect and deal with problems when you are using the IBM WebSphere MQ client for HP Integrity NonStop Server.

Toggling between the use of IBM WebSphere MQ and TMF transactions on a single connection

If a IBM WebSphere MQ client for HP Integrity NonStop Server application toggles between the use of IBM WebSphere MQ and TMF transactions on a single connection, then IBM WebSphere MQ operations such as MQPUT and MQGET might fail with a return code of “2072 (0818) (RC2072): MQRC_SYNCPOINT_NOT_AVAILABLE” on page 1413. Errors and a first failure symptom report for the client application are generated in the IBM WebSphere MQ client for HP Integrity NonStop Server errors directory.

This error occurs because mixed TMF and IBM WebSphere MQ transactions on a single connection are not supported.

Use the standard facilities that are supplied with your system to record the problem identifier and to save any generated output files. Use either the IBM WebSphere MQ Support site: <http://www.ibm.com/software/integration/wmq/support/>, or the IBM Support Assistant (ISA): https://www.ibm.com/support/home/product/C100515X13178X21/other_software/ibm_support_assistant to check whether a solution is already available. If you are unable to find a solution, contact your IBM support center. Do not discard these files until the problem is resolved.

Java and JMS troubleshooting

Use the advice that is given here to help you to resolve common problems that can arise when you are using Java or JMS applications.

Troubleshooting JMSSC0108 messages

There are a number of steps that you can take to prevent a JMSSC0108 message from occurring when you are using activation specifications and WebSphere Application Server listener ports that are running in Application Server Facilities (ASF) mode.

When you are using activation specifications and WebSphere Application Server listener ports that are running in ASF mode, which is the default mode of operation, it is possible that the following message might appear in the application server log file:

```
JMSSC0108: The WebSphere MQ classes for JMS had detected a message, ready for asynchronous delivery to an application. When delivery was attempted, the message was no longer available.
```

Use the information in this topic to understand why this message appears, and the possible steps that you can take to prevent it from occurring.

How activation specifications and listener ports detect and process messages

An activation specification or WebSphere Application Server listener port performs the following steps when it starts up:

1. Create a connection to the queue manager that they have been set to use.
2. Open the JMS destination on that queue manager that they have been configured to monitor.

3. Browse that destination for messages.

When a message is detected, the activation specification or listener port performs the following steps:

1. Constructs an internal message reference that represents the message.
2. Gets a server session from its internal server session pool.
3. Loads the server session up with the message reference.
4. Schedules a piece of work with the application server Work Manager to run the server session and process the message.

The activation specification or listener port then goes back to monitoring the destination again, looking for another message to process.

The application server Work Manager runs the piece of work that the activation specification or listener port submitted on a new server session thread. When started, the thread completes the following actions:

- Starts either a local or global (XA) transaction, depending on whether the message-driven bean requires XA transactions or not, as specified in the message-driven bean's deployment descriptor.
- Gets the message from the destination by issuing a destructive MQGET API call.
- Runs the message-driven bean's `onMessage()` method.
- Completes the local or global transaction, once the `onMessage()` method has finished.
- Return the server session back to the server session pool.

Why the JMSCC0108 message occurs, and how to prevent it

The main activation specification or listener port thread browses messages on a destination. It then asks the Work Manager to start a new thread to destructively get the message and process it. This means that it is possible for a message to be found on a destination by the main activation specification or listener port thread, and no longer be available by the time the server session thread attempts to get it. If this happens, then the server session thread writes the following message to the application server's log file:

```
JMSCC0108: The WebSphere MQ classes for JMS had detected a message, ready for asynchronous delivery to an application. When delivery was attempted, the message was no longer available.
```

There are two reasons why the message is no longer on the destination when the server session thread tries to get it:

- Reason 1: The message has been consumed by another application
- Reason 2: The message has expired

Reason 1: The message has been consumed by another application

If two or more activation specifications and/or listener ports are monitoring the same destination, then it is possible that they could detect the same message and try to process it. When this happens:

- A server session thread started by one activation specification or listener port gets the message and delivers it to a message-driven bean for processing.
- The sever session thread started by the other activation specification or listener port tries to get the message, and finds that it is no longer on the destination.

If an activation specification or listener port is connecting to a queue manager in any of the following ways, the messages that the main activation specification or listener port thread detects are marked:

- A queue manager on any platform, using  IBM WebSphere MQ messaging provider normal mode (*WebSphere MQ V7.1 Reference*).
- A queue manager running on z/OS, using  IBM WebSphere MQ messaging provider migration mode (*WebSphere MQ V7.1 Reference*).

Marking a message prevents any other activation specification or listener port from seeing that message, and trying to process it.

By default, messages are marked for five seconds. After the message has been detected and marked, the five second timer starts. During these five seconds, the following steps must be carried out:

- The activation specification or listener port must get a server session from the server session pool.
- The server session must be loaded with details of the message to process.
- The work must be scheduled.
- The Work Manager must process the work request and start the server session thread.
- The server session thread needs to start either a local or global transaction.
- The server session thread needs to destructively get the message.

On a busy system, it might take longer than five seconds for these steps to be carried out. If this happens, then the mark on the message is released. This means that other activation specifications or listener ports can now see the message, and can potentially try to process it, which can result in the JMSCC0108 message being written to the application server's log file.

In this situation, you should consider the following options:

- Increase the value of  the queue manager property Message mark browse interval (MARKINT) (*WebSphere MQ V7.1 Reference*), to give the activation specification or listener port that originally detected the message more time to get it. Ideally, the property should be set to a value greater than the time taken for your message-driven beans to process messages. This means that, if the main activation specification or listener port thread blocks waiting for a server session because all of the server sessions are busy processing messages, then the message should still be marked when a server session becomes available. Note that the MARKINT property is set on a queue manager, and so is applicable to all applications that browse messages on that queue manager.
- Increase the size of the server session pool used by the activation specification or listener port. This would mean that there are more server sessions available to process messages, which should ensure that messages can be processed within the specified mark interval. One thing to note with this approach is that the activation specification or listener port will now be able to process more messages concurrently, which could impact the overall performance of the application server.

If an activation specification or listener port is connecting to a queue manager running on a platform other than z/OS, using  IBM WebSphere MQ messaging provider migration mode (*WebSphere MQ V7.1 Reference*), the marking functionality is not available. This means that it is not possible to prevent two or more activation specifications and/or listener ports from detecting the same message and trying to process it. In this situation, the JMSCC0108 message is expected.

Reason 2: The message has expired

The other reason that a JMSCC0108 message is generated is if the message has expired in between being detected by the activation specification or listener port and being consumed by the server session. If this happens, when the server session thread tries to get the message, it finds that it is no longer there and so reports the JMSCC0108 message.

Increasing the size of the server session pool used by the activation specification or listener port can help here. Increasing the server session pool size means that there are more server sessions available to process messages, which can potentially mean that the message is processed before it expires. It is important to note that the activation specification or listener port is now able to process more messages concurrently, which could impact the overall performance of the application server.

Problem determination for the WebSphere MQ resource adapter

When using the WebSphere MQ resource adapter, most errors cause exceptions to be thrown, and these exceptions are reported to the user in a manner that depends on the application server. The resource adapter makes extensive use of linked exceptions to report problems. Typically, the first exception in a chain is a high-level description of the error, and subsequent exceptions in the chain provide the more detailed information that is required to diagnose the problem.

For example, if the IVT program fails to obtain a connection to a WebSphere MQ queue manager, the following exception might be thrown:

```
javax.jms.JMSEException: MQJCA0001: An exception occurred in the JMS layer.  
    See the linked exception for details.
```

Linked to this exception is a second exception:

```
javax.jms.JMSEException: MQJMS2005: failed to create an MQQueueManager for  
    'localhost:ExampleQM'
```

This exception is thrown by WebSphere MQ classes for JMS and has a further linked exception:

```
com.ibm.mq.MQException: MQJE001: An MQException occurred: Completion Code 2,  
    Reason 2059
```

This final exception indicates the source of the problem. Reason code 2059 is `MQRC_Q_MGR_NOT_AVAILABLE`, which indicates that the queue manager specified in the definition of the `ConnectionFactory` object might not have been started.

If the information provided by exceptions is not sufficient to diagnose a problem, you might need to request a diagnostic trace. For information about how to enable diagnostic tracing, see  Configuration of the WebSphere MQ resource adapter.

Configuration problems commonly occur in the following areas:

Problems in deploying the resource adapter

If the resource adapter fails to deploy, check that JCA resources are configured correctly. If WebSphere MQ is already installed, check that the correct versions of the JCA and WebSphere MQ classes for JMS are in the class path.

Failures in deploying the resource adapter are generally caused by not configuring JCA resources correctly. For example, a property of the `ResourceAdapter` object might not be specified correctly, or the deployment plan required by the application server might not be written correctly. Failures might also occur when the application server attempts to create objects from the definitions of JCA resources and bind the objects into the JNDI namespace, but certain properties are not specified correctly or the format of a resource definition is incorrect.

The resource adapter can also fail to deploy because it loaded incorrect versions of JCA or WebSphere MQ classes for JMS classes from JAR files in the class path. This type of failure can commonly occur on a system where WebSphere MQ is already installed. On such a system, the application server might find existing copies of the WebSphere MQ classes for JMS JAR files and load classes from them in preference to the classes supplied in the WebSphere MQ resource adapter RAR file. If the extended transactional client JAR file, `com.ibm.mqetclient.jar`, cannot be loaded when the resource adapter is deployed, a warning is written to the diagnostic trace, if enabled, but this does not cause deployment to fail.

Problems in deploying MDBs

Failures when the application server attempts to start message delivery to an MDB might be caused by an error in the definition of the associated `ActivationSpec` object, or by missing resources. Deployment might also fail if an MDB is transacted, the connection is in client mode, but distributed transactions are not available.

Failures might occur when the application server attempts to start message delivery to an MDB. This type of failure is typically caused by an error in the definition of the associated `ActivationSpec` object, or because the resources referenced in the definition are not available. For example, the queue manager might not be running, or a specified queue might not exist.

An `ActivationSpec` object attempts to validate its properties when the MDB is deployed. Deployment then fails if the `ActivationSpec` object has any properties that are mutually exclusive or does not have all the required properties. However, not all problems associated with the properties of the `ActivationSpec` object can be detected at this time.

Deployment might also fail if an MDB is transacted and the connection is in client mode, but distributed transactions are not available because the extended transactional client JAR file, `com.ibm.mqetclient.jar`, is not in the class path.

Failures to start message delivery are reported to the user in a manner that depends on the application server. Typically, these failures are reported in the logs and diagnostic trace of the application server. If enabled, the diagnostic trace of the WebSphere MQ resource adapter also records these failures.

Problems in creating connections for outbound communication

A failure in outbound communication can occur if a `ConnectionFactory` object cannot be found, or if the `ConnectionFactory` object is found but a connection cannot be created. There are various reasons for either of these problems.

Failures in outbound communication typically occur when an application attempts to look up and use a `ConnectionFactory` object in a JNDI namespace. A JNDI exception is thrown if the `ConnectionFactory` object cannot be found in the namespace. A `ConnectionFactory` object might not be found for the following reasons:

- The application specified an incorrect name for the `ConnectionFactory` object.
- The application server was not able to create the `ConnectionFactory` object and bind it into the namespace. In this case, the startup logs of the application server typically contain information about the failure.

If the application successfully retrieves the `ConnectionFactory` object from the JNDI namespace, an exception might still be thrown when the application calls the `ConnectionFactory.createConnection()` method. An exception in this context indicates that it is not possible to create a connection to a WebSphere MQ queue manager. Here are some common reasons why an exception might be thrown:

- The queue manager is not available, or cannot be found using the properties of the `ConnectionFactory` object. For example, the queue manager is not running, or the specified host name, IP address, or port number of the queue manager is incorrect.
- The user is not authorized to connect to the queue manager. For a client connection, if the `createConnection()` call does not specify a user name, and the application server supplies no user identity information, the JVM process ID is passed to the queue manager as the user name. For the connection to succeed, this process ID must be a valid user name in the system on which the queue manager is running.
- The application is a transacted EJB and therefore the connection must be transacted, but distributed transactions are not available because the extended transactional client JAR file, `com.ibm.mqetclient.jar`, is not in the class path.

- The ConnectionFactory object has a property called ccdtURL and a property called channel. These properties are mutually exclusive.
- On an SSL connection, the SSL-related properties, or the SSL-related attributes in the server connection channel definition, have not been specified correctly.
- The sslFipsRequired property has different values for different JCA resources. For more information about this limitation, see  Limitations of the WebSphere MQ resource adapter.

Troubleshooting for IBM WebSphere MQ Telemetry

Look for a troubleshooting task to help you solve a problem with running IBM WebSphere MQ Telemetry applications.

Related tasks:

“Tracing the telemetry service” on page 1209

“Tracing the MQTT v3 Java client” on page 1210

“Resolving problem: MQTT client does not connect” on page 1212

“Resolving problem: MQTT client connection dropped” on page 1214

“Resolving problem: Lost messages in an MQTT application” on page 1215

“Resolving problem: Telemetry service does not start” on page 1216

“Resolving problem: JAAS login module not called by the telemetry service” on page 1218

“Resolving problem: Starting or running the daemon” on page 1221

“Resolving problem: MQTT clients not connecting to the daemon” on page 1221

Related reference:

“Location of telemetry logs, error logs, and configuration files”

“MQTT v3 Java client reason codes” on page 1208

“System requirements for using SHA-2 cipher suites with MQTT channels and clients” on page 1211

Related information:

 IBM WebSphere MQ Telemetry (*WebSphere MQ V7.1 Product Overview Guide*)

Location of telemetry logs, error logs, and configuration files

Find the logs, error logs, and configuration files used by WebSphere MQ Telemetry.

Note: The examples are coded for Windows systems. Change the syntax to run the examples on AIX or Linux systems.

Server-side logs

The installation wizard for WebSphere MQ Telemetry writes messages to its installation log:

```
WMQ program directory\mqxr
```

The telemetry service writes messages to the WebSphere MQ queue manager error log, and FDC files to the WebSphere MQ error directory:

```
WMQ data directory\Qmgrs\qMgrName\errors\AMQERR01.LOG
```

```
WMQ data directory\errors\AMQnnn.n.FDC
```

It also writes a log for the telemetry (MQXR) service. The log displays the properties the service started with, and errors it has found acting as a proxy for an MQTT client. For example, unsubscribing from a subscription that the client did not create. The log path is:

```
WMQ data directory\Qmgrs\qMgrName\errors\mqxr.log
```

The WebSphere MQ telemetry sample configuration created by WebSphere MQ explorer starts the telemetry (MQXR) service using the command **runMQXRService**, which is in *WMQ Telemetry install directory\bin*. This command writes to:

```
WMQ data directory\Qmgrs\qMgrName\mqxr.stdout
WMQ data directory\Qmgrs\qMgrName\mqxr.stderr
```

Modify **runMQXRService** to display the paths configured for the telemetry (MQXR) service, or to echo the initialization before starting the telemetry (MQXR) service.

Server-side configuration files

Telemetry channels and telemetry (MQXR) service

Restriction: The format, location, content, and interpretation of the telemetry channel configuration file might change in future releases. You must use WebSphere MQ Explorer to configure telemetry channels.

WebSphere MQ Explorer saves telemetry configurations in the `mqxr_win.properties` file on Windows systems, and the `mqxr_unix.properties` file on AIX or Linux systems. The properties files are saved in the telemetry configuration directory:

```
WMQ data directory\Qmgrs\qMgrName\mqxr
```

Figure 137. Telemetry configuration directory on Windows

```
/var/mqm/qmgrs/qMgrName/mqxr
```

Figure 138. Telemetry configuration directory on AIX or Linux

JVM Set Java properties that are passed as arguments to the telemetry (MQXR) service in the file, `java.properties`. The properties in the file are passed directly to the JVM running the telemetry (MQXR) service. They are passed as additional JVM properties on the Java command line. Properties set on the command line take precedence over properties added to the command line from the `java.properties` file.

Find the `java.properties` file in the same folder as the telemetry configurations. See Figure 137 and Figure 138.

Modify `java.properties` by specifying each property as a separate line. Format each property exactly as you would to pass the property to the JVM as an argument. For example:

```
-Xmx1024m
-Xms1024m
```

JAAS The JAAS configuration file is described in “Telemetry channel JAAS configuration” on page 145, which includes the sample JAAS configuration file, `JAAS.config`, shipped with WebSphere MQ Telemetry.

If you configure JAAS, you are almost certainly going to write a class to authenticate users to replace the standard JAAS authentication procedures.

To include your Login class in the class path used by the telemetry (MQXR) service class path, provide a WebSphere MQ `service.env` configuration file.

Set the class path for your JAAS LoginModule in `service.env`. You cannot use the variable, `%classpath%` in `service.env`. The class path in `service.env` is added to the class path already set in the telemetry (MQXR) service definition.

Display the class paths that are being used by the telemetry (MQXR) service by adding `echo set classpath` to `runMQXRService.bat`. The output is sent to `mqxr.stdout`.

The default location for the `service.env` file is:

WMQ data directory\service.env

Override these settings with a service.env file for each queue manager in:

WMQ data directory\Qmgrs\qMgrName\service.env

```
CLASSPATH=WMQ Install Directory\mqxr\samples
```

Note: service.env must not contain any variables. Substitute the actual value of *WMQ Install Directory*.

Figure 139. Sample service.env for Windows.

A sample service.env file to use the sample LoginModule.class.

Trace See “Tracing the telemetry service” on page 1209. The parameters to configure trace are stored in two files:

WMQ data directory\Qmgrs\qMgrName\mqxr\trace.config

WMQ data directory\Qmgrs\qMgrName\mqxr\mqxrtrace.properties

Client-side log files

The default file persistence class in the Java SE MQTT client supplied with WebSphere MQ Telemetry creates a folder with the name: *clientIdentifier-tcp*hostNameport or *clientIdentifier-ssl*hostNameport in the client working directory. The folder name tells you the hostName and port used in the connection attempt. The folder contains messages that have been stored by the persistence class. The messages are deleted when they have been delivered successfully.

The folder is deleted when a client, with a clean session, ends.

If client trace is turned on, the unformatted log is, by default, stored in the client working directory. The trace file is called mqttt-*n*.trc

Client-side configuration files

Set trace and SSL properties for the MQTT Java client using Java property files, or set the properties programmatically. Pass the properties to the MQTT Java client using the JVM -D switch: for example,

```
Java -Dcom.ibm.micro.client.mqttv3.trace=c:\\MqttTrace.properties  
-Dcom.ibm.ssl.keyStore=C:\\MyKeyStore.jks
```

See “Tracing the MQTT v3 Java client” on page 1210. For links to client API documentation for the MQTT client libraries, see  MQTT client programming reference.

MQTT v3 Java client reason codes

Look up the causes of reason codes in an MQTT v3 Java client exception or throwable.

Table 130. MQTT v3 Java client reason codes

Reason code	Value	Cause
REASON_CODE_BROKER_UNAVAILABLE	3	
REASON_CODE_CLIENT_ALREADY_CONNECTED	32100	The client is already connected.
REASON_CODE_CLIENT_ALREADY_DISCONNECTED	32101	The client is already disconnected.
REASON_CODE_CLIENT_DISCONNECT_PROHIBITED	32107	Thrown when an attempt to call <code>MqttClient.disconnect</code> has been made from within a method on <code>MqttCallback</code> .
REASON_CODE_CLIENT_DISCONNECTING	32102	The client is currently disconnecting and cannot accept any new work.
REASON_CODE_CLIENT_EXCEPTION	0	Client encountered an exception.
REASON_CODE_CLIENT_NOT_CONNECTED	32104	The client is not connected to the server.
REASON_CODE_CLIENT_TIMEOUT	32000	Client timed out while waiting for a response from the server.
REASON_CODE_FAILED_AUTHENTICATION	4	Authentication with the server has failed, due to a bad user name or password.
REASON_CODE_INVALID_CLIENT_ID	2	The server has rejected the supplied client ID.
REASON_CODE_INVALID_PROTOCOL_VERSION	1	The protocol version requested is not supported by the server.
REASON_CODE_NO_MESSAGE_IDS_AVAILABLE	32001	Internal error, caused by no new message IDs being available.
REASON_CODE_NOT_AUTHORIZED	5	Not authorized to perform the requested operation.
REASON_CODE_SERVER_CONNECT_ERROR	32103	Unable to connect to server.
REASON_CODE_SOCKET_FACTORY_MISMATCH	32105	Server URI and supplied <code>SocketFactory</code> do not match.
REASON_CODE_SSL_CONFIG_ERROR	32106	SSL configuration error.
REASON_CODE_UNEXPECTED_ERROR	6	An unexpected error has occurred.

Tracing the telemetry service

Follow these instructions to start a trace of the telemetry service, set the parameters that control the trace, and find the trace output.

Before you begin

Tracing is a support function. Follow these instructions if an IBM service engineer asks you to trace your telemetry service. The product documentation does not document the format of the trace file, or how to use it to debug a client.

About this task

You can use the IBM WebSphere MQ `strmqtrc` and `endmqtrc` commands to start and stop IBM WebSphere MQ trace. `strmqtrc` captures trace for the telemetry service. When using `strmqtrc`, there is a delay of up to a couple of seconds before the telemetry service trace is started. For further information about IBM WebSphere MQ trace, see [Tracing](#). Alternatively, you can trace the telemetry service by using the following procedure:

Procedure

1. Set the trace options to control the amount of detail and the size of the trace. The options apply to a trace started with either the `strmqtrc` or the `controlMQXRChannel` command.

Set the trace options in the following files:

Trace only works for the WebSphere MQ Telemetry Java client.

About this task

Note: The examples are coded for Windows. Change the syntax to run the examples on Linux²⁷.

Procedure

1. Create a Java properties file containing the trace configuration.
In the properties file specify the following optional properties. If a property key is specified more than once, the last occurrence sets the property.
 - a. `com.ibm.micro.client.mqttv3.trace.outputName`
The directory to write the trace file to. It defaults to the client working directory. The trace file is called `mqtt-n.trc`.
`com.ibm.micro.client.mqttv3.trace.outputName=c:\\MQTT_Trace`
 - b. `com.ibm.micro.client.mqttv3.trace.count`
The number of trace files to write. The default is one file, of unlimited size.
`com.ibm.micro.client.mqttv3.trace.count=5`
 - c. `com.ibm.micro.client.mqttv3.trace.limit`
The maximum size of file to write, the default is 500000. The limit only applies if more than one trace file is requested.
`com.ibm.micro.client.mqttv3.trace.limit=100000`
 - d. `com.ibm.micro.client.mqttv3.trace.client.clientIdentifier.status`
Turn trace on or off, per client. If `clientIdentifier=*`, trace is turned on or off for all clients. By default, trace is turned off for all clients.
`com.ibm.micro.client.mqttv3.trace.client.*.status=on`
`com.ibm.micro.client.mqttv3.trace.client.Client10.status=on`
2. Pass the trace properties file to the JVM using a system property.
`-Dcom.ibm.micro.client.mqttv3.trace=c:\\MqttTrace.properties`
3. Run the client.
4. Convert the trace file from binary encoding to .html.
`java -jar com.ibm.micro.client.mqttv3.trace.jar -i mqtt-0.trc -o mqtt-0.trc.html`

System requirements for using SHA-2 cipher suites with MQTT channels and clients

From Java 6 SR13 onwards, you can use SHA-2 cipher suites to secure your MQTT channels and client apps. However, SHA-2 cipher suites are not enabled by default until Java 7 SR4 onwards, so in earlier versions you must specify the required suite. If you are running an MQTT client with your own JRE, you need to ensure that it supports the SHA-2 cipher suites. For your client apps to use SHA-2 cipher suites, the client must also set the SSL context to a value that supports TLS 1.2.

For Java 7 SR4 onwards, SHA-2 cipher suites are enabled by default. For Java 6 SR13 and later service releases, if you define an MQTT channel without specifying a cipher suite, the channel will not accept connections from a client using a SHA-2 cipher suite. To use SHA-2 cipher suites, you must specify the required suite in the channel definition. This makes the XR Service enable the suite before making connections. It also means that only client apps using the specified suite can connect to this channel.

²⁷. Java uses the correct path delimiter. You can code the delimiter in a property file as `'/'` or `'\\'`; `'\'` is the escape character

There is a similar limitation for the MQTT client for Java. If the client code is running on a Java 1.6 JRE, the required SHA-2 cipher suites must be explicitly enabled. In order to use these suites, the client must also set the SSL context to a value that supports Version 1.2 of the Transport Layer Security (TLS) protocol. For example:

```
MqttConnectOptions mqttConnectOptions = new MqttConnectOptions();
java.util.Properties sslClientProps = new java.util.Properties();
sslClientProps.setProperty("com.ibm.ssl.keyStore", sslKeys.clientKeyStore);
sslClientProps.setProperty("com.ibm.ssl.keyStorePassword", sslKeys.clientStorePassword);
sslClientProps.setProperty("com.ibm.ssl.trustStore", sslKeys.clientKeyStore);
sslClientProps.setProperty("com.ibm.ssl.trustStorePassword", sslKeys.clientStorePassword);
sslClientProps.setProperty("com.ibm.ssl.protocol", "TLSv1.2");
sslClientProps.setProperty("com.ibm.ssl.enabledCipherSuites", "SSL_RSA_WITH_AES_256_CBC_SHA256" );
mqttConnectOptions.setSSLProperties(sslClientProps);
```

As at May 2013, none of the browsers that work with the MQTT messaging client for JavaScript support the TLS 1.2 protocol, so only the MQTT client for Java and the MQTT client for C can be used to make SHA-2 connections.

For a list of the cipher suites that are currently supported, see any of the related links.

Related concepts:

- “MQTT client authentication using SSL” on page 133
- “Telemetry channel authentication using SSL” on page 137

Related reference:

-  DEFINE CHANNEL (MQTT) (*WebSphere MQ V7.1 Reference*)
-  ALTER CHANNEL (MQTT) (*WebSphere MQ V7.1 Reference*)

Resolving problem: MQTT client does not connect

Resolve the problem of an MQTT client program failing to connect to the telemetry service.

Before you begin

Is the problem at the server, at the client, or with the connection? Have you have written your own MQTT v3 protocol handling client, or an MQTT client application using the C or Java WebSphere MQTT clients?

Run the verification application supplied with WebSphere MQ Telemetry on the server, and check that the telemetry channel and telemetry service are running correctly. Then transfer the verification application to the client, and run the verification application there.

About this task

There are a number of reasons why an MQTT client might not connect, or you might conclude it has not connected, to the telemetry server.

Procedure

1. Consider what inferences can be drawn from the reason code that the telemetry service returned to MqttClient.Connect. What type of connection failure is it?

Option	Description
REASON_CODE_INVALID_PROTOCOL_VERSION	Make sure that the socket address corresponds to a telemetry channel, and you have not used the same socket address for another broker.

Option	Description
REASON_CODE_INVALID_CLIENT_ID	Check that the client identifier is no longer than 23 bytes, and contains only characters from the range: A-Z, a-z, 0-9, '._/%
REASON_CODE_SERVER_CONNECT_ERROR	Check that the telemetry service and the queue manager are running normally. Use netstat to check that the socket address is not allocated to another application.

If you have written an MQTT client library rather than use one of the libraries provided by WebSphere MQ Telemetry, look at the CONNACK return code.

From these three errors you can infer that the client has connected to the telemetry service, but the service has found an error.

2. Consider what inferences can be drawn from the reason codes that the client produces when the telemetry service does not respond:

Option	Description
REASON_CODE_CLIENT_EXCEPTION REASON_CODE_CLIENT_TIMEOUT	Look for an FDC file at the server; see “Server-side logs” on page 1206. When the telemetry service detects the client has timed out, it writes a first-failure data capture (FDC) file. It writes an FDC file whenever the connection is unexpectedly broken.

The telemetry service might not have responded to the client, and the timeout at the client expires. The WebSphere MQ Telemetry Java client only hangs if the application has set an indefinite timeout. The client throws one of these exceptions after the timeout set for `MqttClient.Connect` expires with an undiagnosed connection problem.

Unless you find an FDC file that correlates with the connection failure you cannot infer that the client tried to connect to the server:

- a. Confirm that the client sent a connection request.

Check the TCP/IP request with a tool such as **tcpmon**, available from  <https://tcpmon.dev.java.net/>

- b. Does the remote socket address used by the client match the socket address defined for the telemetry channel?

The default file persistence class in the Java SE MQTT client supplied with WebSphere MQ Telemetry creates a folder with the name: `clientIdentifier-tcpHostNameport` or `clientIdentifier-sslHostNameport` in the client working directory. The folder name tells you the `hostName` and `port` used in the connection attempt; see “Client-side log files” on page 1208.

- c. Can you ping the remote server address?
- d. Does **netstat** on the server show the telemetry channel is running on the port the client is connecting too?

3. Check whether the telemetry service found a problem in the client request.

The telemetry service writes errors it detects into `mqr.log`, and the queue manager writes errors into `AMQERR01.LOG`; see

4. Attempt to isolate the problem by running another client.

- Run the MQTT sample application using the same telemetry channel.
- Run the **wmqttSample** GUI client to verify the connection. Get **wmqttSample** by downloading SupportPac  IA92.

Note: Older versions of IA92 do not include the MQTT v3 Java client library.

Run the sample programs on the server platform to eliminate uncertainties about the network connection, then run the samples on the client platform.

5. Other things to check:

- a. Are tens of thousands of MQTT clients trying to connect at the same time?
Telemetry channels have a queue to buffer a backlog of incoming connections. Connections are processed in excess of 10,000 a second. The size of the backlog buffer is configurable using the telemetry channel wizard in WebSphere MQ Explorer. Its default size is 4096. Check that the backlog has not been configured to a low value.
- b. Are the telemetry service and queue manager still running?
- c. Has the client connected to a high availability queue manager that has switched its TCPIP address?
- d. Is a firewall selectively filtering outbound or return data packets?

Resolving problem: MQTT client connection dropped

Find out what is causing a client to throw unexpected `ConnectionLost` exceptions after successfully connecting and running for either a short or long while.

Before you begin

The MQTT client has connected successfully. The client might be up for a long while. If clients are starting with only a short interval between them, the time between connecting successfully and the connection being dropped might be short.

It is not hard to distinguish a dropped connection from a connection that was successfully made, and then later dropped. A dropped connection is defined by the MQTT client calling the `MqttCallback.ConnectionLost` method. The method is only called after the connection has been successfully established. The symptom is different to `MqttClient.Connect` throwing an exception after receiving a negative acknowledgment or timing out.

If the MQTT client application is not using the MQTT client libraries supplied by WebSphere MQ, the symptom depends on the client. In the MQTT v3 protocol, the symptom is a lack of timely response to a request to the server, or the failure of the TCP/IP connection.

About this task

The MQTT client calls `MqttCallback.ConnectionLost` with a throwable exception in response to any server-side problems encountered after receiving a positive connection acknowledgment. When an MQTT client returns from `MqttTopic.publish` and `MqttClient.subscribe` the request is transferred to an MQTT client thread that is responsible for sending and receiving messages. Server-side errors are reported asynchronously by passing a throwable exception to the `ConnectionLost` callback method.

The telemetry service always writes a first-failure data capture file if it drops the connection.

Procedure

1. Has another client started that used the same `ClientIdentifier`?
If a second client is started, or the same client is restarted, using the same `ClientIdentifier`, the first connection to the first client is dropped.
2. Has the client accessed a topic that it is not authorized to publish or subscribe to?
Any actions the telemetry service takes on behalf of a client that return `MQCC_FAIL` result in the service dropping the client connection.
The reason code is not returned to the client.
 - Look for log messages in the `mqxr.log` and `AMQERR01.LOG` files for the queue manager the client is connected to; see “Server-side logs” on page 1206.
3. Has the TCP/IP connection dropped?

A firewall might have a low timeout setting for marking a TCPIP connection as inactive, and dropped the connection.

- Shorten the inactive TCPIP connection time using `MqttConnectOptions.setKeepAliveInterval`.

Resolving problem: Lost messages in an MQTT application

Resolve the problem of losing a message. Is the message non-persistent, sent to the wrong place, or never sent? A wrongly coded client program might lose messages.

Before you begin

How certain are you that the message you sent, was lost? Can you infer that a message is lost because the message was not received? If message is a publication, which message is lost: the message sent by the publisher, or the message sent to the subscriber? Or did the subscription get lost, and the broker is not sending publications for that subscription to the subscriber?

If the solution involves distributed publish/subscribe, using clusters or publish/subscribe hierarchies, there are numerous configuration issues that might result in the appearance of a lost message.

If you sent a message with "At least once" or "At most once" quality of service, it is likely that the message you think is lost was not delivered in the way you expected. It is unlikely that the message has been wrongly deleted from the system. It might have failed to create the publication or the subscription you expected.

The most important step you take in doing problem determination of lost messages is to confirm the message is lost. Recreate the scenario and lose more messages. Use the "At least once" or "At most once" quality of service to eliminate all cases of the system discarding messages.

About this task

There are four legs to diagnosing a lost message.

1. "Fire and forget" messages working as-designed. "Fire and forget" messages are sometimes discarded by the system.
2. Configuration: setting up publish/subscribe with the correct authorities in a distributed environment is not straightforward.
3. Client programming errors: the responsibility for message delivery is not solely the responsibility of code written by IBM.
4. If you have exhausted all these possibilities, you might decide to involve IBM service.

Procedure

1. If the lost message had the "Fire and forget" quality of service, set the "At least once" or "At most once" quality of service. Attempt to lose the message again.
 - Messages sent with "Fire and forget" quality of service are thrown away by WebSphere MQ in a number of circumstances:
 - Communications loss and channel stopped.
 - Queue manager shut down.
 - Excessive number of messages.
 - The delivery of "Fire and forget" messages depends upon the reliability of TCP/IP. TCP/IP continues to send data packets again until their delivery is acknowledged. If the TCP/IP session is broken, messages with the "Fire and forget" quality of service are lost. The session might be broken by the client or server closing down, a communications problem, or a firewall disconnecting the session.
2. Check that client is restarting the previous session, in order to send undelivered messages with "At least once" or "At most once" quality of service again.

- a. If the client application is using the Java SE MQTT client, check that it sets `MqttClient.CleanSession` to `false`
- b. If you are using different client libraries, check that a session is being restarted correctly.
3. Check that the client application is restarting the same session, and not starting a different session by mistake.
To start the same session again, `cleanSession = false`, and the `Mqttclient.clientIdentifier` and the `MqttClient.serverURI` must be the same as the previous session.
4. If a session closes prematurely, check that the message is available in the persistence store at the client to send again.
 - a. If the client application is using the Java SE MQTT client, check that the message is being saved in the persistence folder; see “Client-side log files” on page 1208
 - b. If you are using different client libraries, or you have implemented your own persistence mechanism, check that it is working correctly.
5. Check that no one has deleted the message before it was delivered.
Undelivered messages awaiting delivery to MQTT clients are stored in `SYSTEM.MQTT.TRANSMIT.QUEUE`. Messages awaiting delivery to the telemetry server are stored by the client persistence mechanism; see  Message persistence in MQTT clients (*WebSphere MQ V7.1 Programming Guide*).
6. Check that the client has a subscription for the publication it expects to receive.
List subscriptions using WebSphere MQ Explorer, or by using `runmqsc` or PCF commands. All MQTT client subscriptions are named. They are given a name of the form: **ClientIdentifier:Topic name**
7. Check that the publisher has authority to publish, and the subscriber to subscribe to the publication topic.

```
dspmqaut -m qMgr -n topicName -t topic -p user ID
```


In a clustered publish/subscribe system, the subscriber must be authorized to the topic on the queue manager to which the subscriber is connected. It is not necessary for the subscriber to be authorized to subscribe to the topic on the queue manager where the publication is published. The channels between the queue managers must be correctly authorized to pass on the proxy subscription and forward the publication.
Create the same subscription and publish to it using WebSphere MQ Explorer. Simulate your application client publishing and subscribing by using the client utility. Start the utility from WebSphere MQ Explorer and change its user ID to match the one adopted by your client application.
8. Check that the subscriber has permission to put the publication on the `SYSTEM.MQTT.TRANSMIT.QUEUE`.

```
dspmqaut -m qMgr -n queueName -t queue -p user ID
```
9. Check that the WebSphere MQ point-to-point application has authority to put its message on the `SYSTEM.MQTT.TRANSMIT.QUEUE`.

```
dspmqaut -m qMgr -n queueName -t queue -p user ID
```


See “Sending a message to a client directly” on page 128.

Resolving problem: Telemetry service does not start

Resolve the problem of the telemetry service failing to start. Check the WebSphere MQ Telemetry installation and no files are missing, moved, or have the wrong permissions. Check the paths used by the telemetry service locate the telemetry service programs.

Before you begin

The WebSphere MQ Telemetry feature is installed. The WebSphere MQ Explorer has a Telemetry folder in **IBM WebSphere MQ > Queue Managers > qMgrName > Telemetry**. If the folder does not exist, the installation has failed.

The Telemetry service must have been created for it to start. If the telemetry service has not been created, then run the **Define sample configuration...** wizard in the Telemetry folder.

If the telemetry service has been started before, then additional **Channels** and **Channel Status** folders are created under the Telemetry folder. The Telemetry service, SYSTEM.MQXR.SERVICE, is in the **Services** folder. It is visible if the Explorer radio button to show System Objects is clicked.

Right click SYSTEM.MQXR.SERVICE to start and stop the service, show its status, and display whether your user ID has authority to start the service.

About this task

The SYSTEM.MQXR.SERVICE telemetry service fails to start. A failure to start manifests itself in two different ways:

1. The start command fails immediately.
2. The start command succeeds, and is immediately followed by the service stopping.

Procedure

1. Start the service

Result The service stops immediately. A window displays an error message; for example:

WebSphere MQ cannot process the request because the executable specified cannot be started. (AMQ4160)

Reason

Files are missing from the installation, or the permissions on installed files are set wrongly. The WebSphere MQ Telemetry feature is installed only on one of a pair of highly available queue managers. If the queue manager instance switches over to a standby, it tries to start SYSTEM.MQXR.SERVICE. The command to start the service fails because the telemetry service is not installed on the standby.

Investigation

Look in error logs; see “Server-side logs” on page 1206.

Actions

Install, or uninstall and reinstall the WebSphere MQ Telemetry feature.

2. Start the service; wait for 30 seconds; refresh the Explorer and check the service status.

Result The service starts and then stops.

Reason

SYSTEM.MQXR.SERVICE started the **runMQXRService** command, but the command failed.

Investigation

Look in error logs; see “Server-side logs” on page 1206.

See if the problem occurs with only the sample channel defined. Backup and the clear the contents of the *WMQ data directory*\Qmgrs\qMgrName\mqxr\ directory. Run the sample configuration wizard and try to start the service.

Actions

Look for permission and path problems.

Resolving problem: JAAS login module not called by the telemetry service

Find out if your JAAS login module is not being called by the telemetry service, and configure JAAS to correct the problem.

Before you begin

You have modified *WMQ installation directory*\mqxr\samples\LoginModule.java to create your own authentication class *WMQ installation directory*\mqxr\samples\samples\LoginModule.class.

Alternatively, you have written your own JAAS authentication classes and placed them in a directory of your choosing. After some initial testing with the telemetry service, you suspect that your authentication class is not being called by the telemetry service.

Note: Guard against the possibility that your authentication classes might be overwritten by maintenance being applied to WebSphere MQ. Use your own path for authentication classes, rather than a path within the WebSphere MQ directory tree.

About this task

The task uses a scenario to illustrate how to resolve the problem. In the scenario, a package called *security.jaas* contains a JAAS authentication class called *JAASLogin.class*. It is stored in the path *C:\WMQTelemetryApps\security\jaas*. Refer to “Telemetry channel JAAS configuration” on page 145 for help in configuring JAAS for WebSphere MQ Telemetry. The example, “Example JAAS configuration” is a sample configuration.

Procedure

1. Look in *mqxr.log* for an exception thrown by `javax.security.auth.login.LoginException`.
See “Server-side logs” on page 1206 for the path to *mqxr.log*, and Figure 146 on page 1220 for an example of the exception listed in the log.
2. Correct your JAAS configuration by comparing it with the worked example in “Example JAAS configuration.”
3. Replace your login class by the sample *JAASLoginModule*, after refactoring it into your authentication package and deploy it using the same path. Switch the value of `loggedIn` between `true` and `false`.
If the problem goes away when `loggedIn` is `true`, and appears the same when `loggedIn` is `false`, the problem lies in your login class.
4. Check whether the problem is with authorization rather than authentication.
 - a. Change the telemetry channel definition to perform authorization checking using a fixed user ID. Select a user ID that is a member of the `mqm` group.
 - b. Rerun the client application.
If the problem disappears, the solution lies with the user ID being passed for authorization. What is the user name being passed? Print it to file from your login module. Check its access permissions using WebSphere MQ Explorer, or `dspmqauth`.

Example JAAS configuration

Use the **New telemetry channel** wizard, in WebSphere MQ Explorer, to configure a telemetry channel. The client connects on port 1884, and connects to the `JAASMCUser` telemetry channel. Figure 140 on page 1219 shows an example of the telemetry properties file created by the telemetry wizard. Do not edit this file directly. The channel authenticates using JAAS, using the configuration called `JAASConfig`. Once the client has authenticated, it uses the user ID `Admin` to authorize its access to WebSphere MQ objects.

```
com.ibm.mq.MQXR.channel/JAASMCUser: \  
com.ibm.mq.MQXR.Port=1884;\  
com.ibm.mq.MQXR.JAASConfig=JAASConfig;\  
com.ibm.mq.MQXR.UserName=Admin;\  
com.ibm.mq.MQXR.StartWithMQXRService=true
```

Figure 140. WMQ Installation directory\data\qmgrs\qMgrName\mqxr\mqxr_win.properties

The JAAS configuration file has a stanza named JAASConfig that names the Java class security.jaas.JAASLogin, which JAAS is to use to authenticate clients.

```
JAASConfig {  
    security.jaas.JAASLogin required debug=true;  
};
```

Figure 141. WMQ Installation directory\data\qmgrs\qMgrName\mqxr\jaas.config

When SYSTEM.MQTT.SERVICE starts, it adds the path in Figure 142 to its classpath.

```
CLASSPATH=C:\WMQTelemetryApps;
```

Figure 142. WMQ Installation directory\data\qmgrs\qMgrName\service.env

Figure 143 shows the additional path in Figure 142 added to the classpath that is set up for the telemetry service.

```
CLASSPATH=;C:\IBM\MQ\Program\mqxr\bin\..\lib\MQXRListener.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\lib\WMQCommonServices.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\lib\objectManager.utils.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\lib\com.ibm.micro.xr.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\..\java\lib\com.ibm.mq.jmqi.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\..\java\lib\com.ibm.mqjms.jar;  
C:\IBM\MQ\Program\mqxr\bin\..\..\java\lib\com.ibm.mq.jar;  
C:\WMQTelemetryApps;
```

Figure 143. Classpath output from runMQXRService.bat

The output in Figure 144 shows that the telemetry service has started with the channel definition shown in Figure 140.

```
21/05/2010 15:32:12 [main] com.ibm.mq.MQXRService.MQXRPropertiesFile  
AMQXR2011I: Property com.ibm.mq.MQXR.channel/JAASMCUser value  
com.ibm.mq.MQXR.Port=1884;  
com.ibm.mq.MQXR.JAASConfig=JAASConfig;  
com.ibm.mq.MQXR.UserName=Admin;  
com.ibm.mq.MQXR.StartWithMQXRService=true
```

Figure 144. WMQ Installation directory\data\qmgrs\qMgrName\errors\mqxr.log

When the client application connects to the JAAS channel, if com.ibm.mq.MQXR.JAASConfig=JAASWrongConfig does not match the name of a JAAS stanza in the

jaas.config file, the connection fails, and the client throws an exception with a return code of 0; see Figure 145. The second exception, Client is not connected (32104), was thrown because the client attempted to disconnect when it was not connected.

```
C:\WMQTelemetryApps>java com.ibm.mq.id.PubAsyncRestartable
Starting a clean session for instance "Admin_PubAsyncRestartab"
Publishing "Hello World Fri May 21 17:23:23 BST 2010" on topic "MQTT Example"
for client instance: "Admin_PubAsyncRestartab" using QoS=1 on address
tcp://localhost:1884"
Userid: "Admin", Password: "Password"
Delivery token "528752516" has been received: false
Connection lost on instance "Admin_PubAsyncRestartab" with cause "MqttException"
MqttException (0) - java.io.EOFException
    at com.ibm.micro.client.mqttv3.internal.CommsReceiver.run(
CommsReceiver.java:118)
    at java.lang.Thread.run(Thread.java:801)
Caused by: java.io.EOFException
    at java.io.DataInputStream.readByte(DataInputStream.java:269)
    at com.ibm.micro.client.mqttv3.internal.wire.MqttInputStream.
readMqttWireMessage(MqttInputStream.java:56)
    at com.ibm.micro.client.mqttv3.internal.CommsReceiver.run(CommsReceiver.java:90)
    ... 1 more
Client is not connected (32104)
    at com.ibm.micro.client.mqttv3.internal.ExceptionHelper.createMqttException(
ExceptionHelper.java:33)
    at com.ibm.micro.client.mqttv3.internal.ClientComms.internalSend(
ClientComms.java:100)
    at com.ibm.micro.client.mqttv3.internal.ClientComms.sendNowait(
ClientComms.java:117)
    at com.ibm.micro.client.mqttv3.internal.ClientComms.disconnect(
ClientComms.java:229)
    at com.ibm.micro.client.mqttv3.MqttClient.disconnect(MqttClient.java:385)
    at com.ibm.mq.id.PubAsyncRestartable.main(PubAsyncRestartable.java:49)
```

Figure 145. Exception thrown connecting com.ibm.mq.id.PubAsyncRestartable

mqxr.log contains additional output shown in Figure 145.

The error is detected by JAAS which throws `javax.security.auth.login.LoginException` with the cause `No LoginModules configured for JAAS`. It could be caused, as in Figure 146, by a bad configuration name. It might also be the result of other problems JAAS has encountered loading the JAAS configuration.

If no exception is reported by JAAS, JAAS has successfully loaded the `security.jaas.JAASLogin` class named in the JAASConfig stanza.

```
21/05/2010 12:06:12 [ServerWorker0] com.ibm.mq.MQXRService.MQTTCommunications
AMQXR2050E: Unable to load JAAS config: JAASWrongConfig.
The following exception occurred javax.security.auth.login.LoginException:
No LoginModules configured for JAAS
```

Figure 146. mqxr.log - error loading JAAS configuration

Resolving problem: Starting or running the daemon

Consult the WebSphere MQ Telemetry daemon for devices console log, turn on tracing, or use the symptom table in this topic to troubleshoot problems with the daemon.

Procedure

1. Check the console log.

If the daemon is running in the foreground, the console messages are written to the terminal window.
If the daemon has been started in the background, the console is where you have redirected stdout to.

2. Restart the daemon.

Changes to the configuration file are not activated until the daemon is restarted.

3. Consult Table 131:

Table 131. Symptom table

Problem	Suggested solution
The following message is displayed when you start the daemon on Windows: The system cannot execute the specified program or The application has failed to start because its side-by-side configuration is incorrect.	Install Microsoft Visual C++ 2008 Redistributable Package.
Two or more daemons or MQTT-capable servers are inter-connected by a bridge or bridges, and the processor is showing excessive load.	There is possibly a message loop, with one or more messages being repeatedly passed from one server to another. Examine the topic parameters in the configuration files. Use more specific topics where possible. Broad wildcard characters in both directions are the most common cause of connection loops.
The bridge is unable to connect to a remote MQTT-capable server that other MQTT clients can connect to.	The remote server might be incompatible with attempts to determine if the remote server is also WebSphere MQ Telemetry daemon for devices. Try setting try_private to off to disable special processing to eliminate message loops.
This message is printed when a bridge is configured: Warning: Connect was not first packet on socket 1888, got CONNACK.	You have probably configured a bridge to loop back to the local daemon. Loopback is not supported.

Resolving problem: MQTT clients not connecting to the daemon

Clients are not connecting to the daemon, or the daemon is not connecting to other daemons or to a WebSphere MQ telemetry channel.

About this task

Trace each MQTT packet sent and received by the daemon.

Procedure

Set the **trace_output** parameter to protocol in the daemon configuration file or send a command to the daemon using the amqtdc.upd file.

See  Transfer messages between the WebSphere MQ Telemetry daemon for devices and WebSphere MQ (*WebSphere MQ V7.1 Product Overview Guide*), for an example of using the amqtdc.upd file.

Using the protocol setting, the daemon prints a message to the console describing each MQTT packet it sends and receives.

Troubleshooting channel authentication records

If you are having problems using channel authentication records, check whether the problem is described in the following information.

What address are you presenting to the queue manager?

The address that your channel presents to the queue manager depends on the network adapter being used. For example, if the CONNAME you use to get to the listener is "localhost", you present 127.0.0.1 as your address; if it is the real IP address of your computer, then that is the address you present to the queue manager. You might invoke different authentication rules for 127.0.0.1 and your real IP address.

Using BLOCKADDR with channel names

If you use SET CHLAUTH TYPE(BLOCKADDR), it must have the generic channel name CHLAUTH(*) and nothing else. You must block access from the specified addresses using any channel name.

CHLAUTH(*) on z/OS systems

On z/OS, a channel name including the asterisk (*) must be enclosed in quotation marks. This rule also applies to the use of a single asterisk to match all channel names. Thus, where you would specify CHLAUTH(*) on other platforms, on z/OS you must specify CHLAUTH("*").

Behaviour of SET CHLAUTH command over queue manager restart

If the SYSTEM.CHLAUTH.DATA.QUEUE, has been deleted or altered in a way that it is no longer accessible i.e. PUT(DISABLED), the SET CHLAUTH command will only be partially successful. In this instance, SET CHLAUTH will update the in-memory cache, but will fail when hardening.

This means that although the rule put in place by the SET CHLAUTH command may be operable initially, the effect of the command will not persist over a queue manager restart. The user should investigate, ensuring the queue is accessible and then reissue the command (using ACTION(REPLACE)) before cycling the queue manager.

If the SYSTEM.CHLAUTH.DATA.QUEUE remains inaccessible at queue manager startup, the cache of saved rules cannot be loaded and all channels will be blocked until the queue and rules become accessible.

Maximum size of ADDRESS and ADDRLIST on z/OS systems

On z/OS, the maximum size for the ADDRESS and ADDRLIST fields are 48 characters. Some IPV6 address patterns could be longer than this limit, for example '0000-ffff:0000-ffff:0000-ffff:0000-ffff:0000-ffff:0000-ffff:0000-ffff:0000-ffff'. In this case, you could use '*' instead.

If you want to use a pattern more than 48 characters long, try to express the requirement in a different way. For example, instead of specifying

'0001-fffe:0001-fffe:0001-fffe:0001-fffe:0001-fffe:0001-fffe:0001-fffe:0001-fffe:0001-fffe' as the address pattern for a USERSRC(MAP), you could specify three rules:

- USERSRC(MAP) for all addresses (*)
- USERSRC(NOACCESS) for address '0000:0000:0000:0000:0000:0000:0000:0000'
- USERSRC(NOACCESS) for address 'ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff'

Multicast troubleshooting

The following hints and tips are in no significant order, and might be added to when new versions of the documentation are released. They are subjects that, if relevant to the work that you are doing, might save you time.

Testing multicast applications on a non-multicast network

Use this information to learn how to test IBM WebSphere MQ Multicast applications locally instead of over a multicast network.

When developing or testing multicast applications you might not yet have a multicast enabled network. To run the application locally, you must edit the `mqclient.ini` file as shown in the following example:

Edit the `Interface` parameter in the `Multicast` stanza of the `MQ_DATA_PATH/mqclient.ini`:

```
Multicast:
  Interface           = 127.0.0.1
```

where `MQ_DATA_PATH` is the location of the IBM WebSphere MQ data directory (`/var/mqm/mqclient.ini`).

The multicast transmissions now only use the local loopback adapter.

Setting the appropriate network for multicast traffic

When developing or testing multicast applications, after testing them locally, you might want to test them over a multicast enabled network. If the application only transmits locally, you might have to edit the `MQClient.ini` file as shown later in this section. If the machine setup is using multiple network adapters, or a virtual private network (VPN) for example, the **Interface** parameter in the `MQClient.ini` file must be set to the address of the network adapter you want to use.

If the `Multicast` stanza exists in the `MQClient.ini` file, edit the **Interface** parameter as shown in the following example:

Change:

```
Multicast:
  Interface           = 127.0.0.1
```

To:

```
Multicast:
  Interface           = IPAddress
```

where `IPAddress` is the IP address of the interface on which multicast traffic flows.

If there is no `Multicast` stanza in the `MQClient.ini` file, add the following example:

```
Multicast:
  Interface           = IPAddress
```

where `IPAddress` is the IP address of the interface on which multicast traffic flows.

The multicast applications now run over the multicast network.

Multicast topic string is too long

If your WebSphere MQ Multicast topic string is rejected with reason code MQRC_TOPIC_STRING_ERROR, it might be because the string is too long.

WebSphereMQ Multicast has a 255 character limit for topic strings. This limitation means that care must be taken with the names of nodes and leaf-nodes within the tree; if the names of nodes and leaf-nodes are too long, the topic string might exceed 255 characters and return the “2425 (0979) (RC2425): MQRC_TOPIC_STRING_ERROR” on page 1534 reason code. It is recommended to make topic strings as short as possible because longer topic strings might have a detrimental effect on performance.

Multicast topic topology issues

Use these examples to understand why certain WebSphere MQ Multicast topic topologies are not recommended.

As was mentioned in “WebSphere MQ Multicast topic topology” on page 159, WebSphere MQ Multicast support requires that each subtree has its own multicast group and data stream within the total hierarchy. Do not use a different multicast group address for a subtree and its parent.

The *classful network* IP addressing scheme has designated address space for multicast address. The full multicast range of IP address is 224.0.0.0 to 239.255.255.255, but some of these addresses are reserved. For a list of reserved address either contact your system administrator or see <http://www.iana.org/assignments/multicast-addresses> for more information. It is recommended that you use the locally scoped multicast address in the range of 239.0.0.0 to 239.255.255.255.

Recommended multicast topic topology

This example is the same as the one from “WebSphere MQ Multicast topic topology” on page 159, and shows 2 possible multicast data streams. Although it is a simple representation, it demonstrates the kind of situation that WebSphere MQ Multicast was designed for, and is shown here to contrast the second example:

```
DEF COMMINFO(MC1) GRPADDR(227.20.133.1)
```

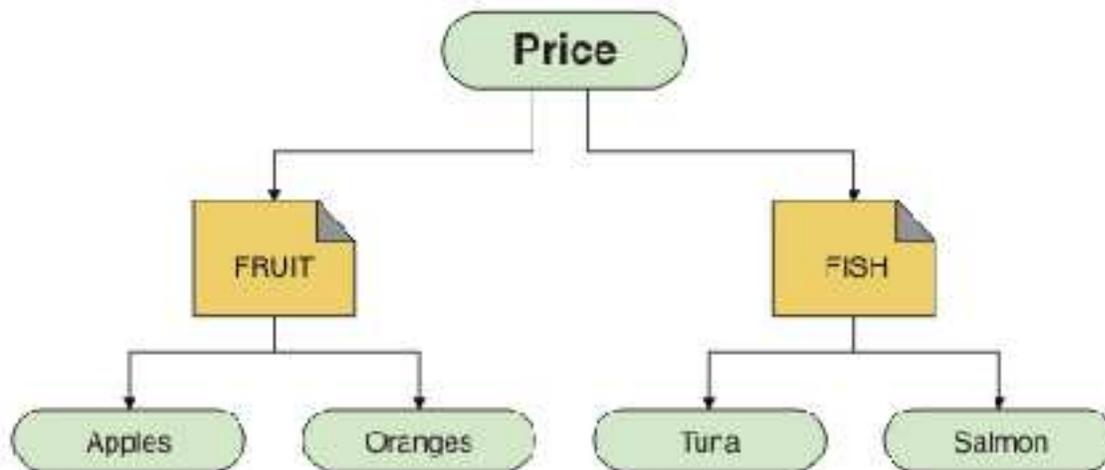
```
DEF COMMINFO(MC2) GRPADDR(227.20.133.2)
```

where 227.20.133.1 and 227.20.133.2 are valid multicast addresses.

These topic definitions are used to create a topic tree as shown in the following diagram:

```
DEFINE TOPIC(FRUIT) TOPICSTRING('Price/FRUIT') MCAST(ENABLED) COMMINFO(MC1)
```

```
DEFINE TOPIC(FISH) TOPICSTRING('Price/FISH') MCAST(ENABLED) COMMINFO(MC2)
```



Each multicast communication information (COMMINFO) object represents a different stream of data because their group addresses are different. In this example, the topic FRUIT is defined to use COMMINFO object MC1, and the topic FISH is defined to use COMMINFO object MC2.

WebSphere MQ Multicast has a 255 character limit for topic strings. This limitation means that care must be taken with the names of nodes and leaf-nodes within the tree; if the names of nodes and leaf-nodes are too long, the topic string might exceed 255 characters and return the MQRC_TOPIC_STRING_ERROR reason code.

Non-recommended multicast topic topology

This example extends the previous example by adding another topic object called ORANGES which is defined to use another COMMINFO object definition (MC3):

```
DEF COMMINFO(MC1) GRPADDR(227.20.133.1)
```

```
DEF COMMINFO(MC2) GRPADDR(227.20.133.2)
```

```
DEF COMMINFO(MC3) GRPADDR(227.20.133.3)
```

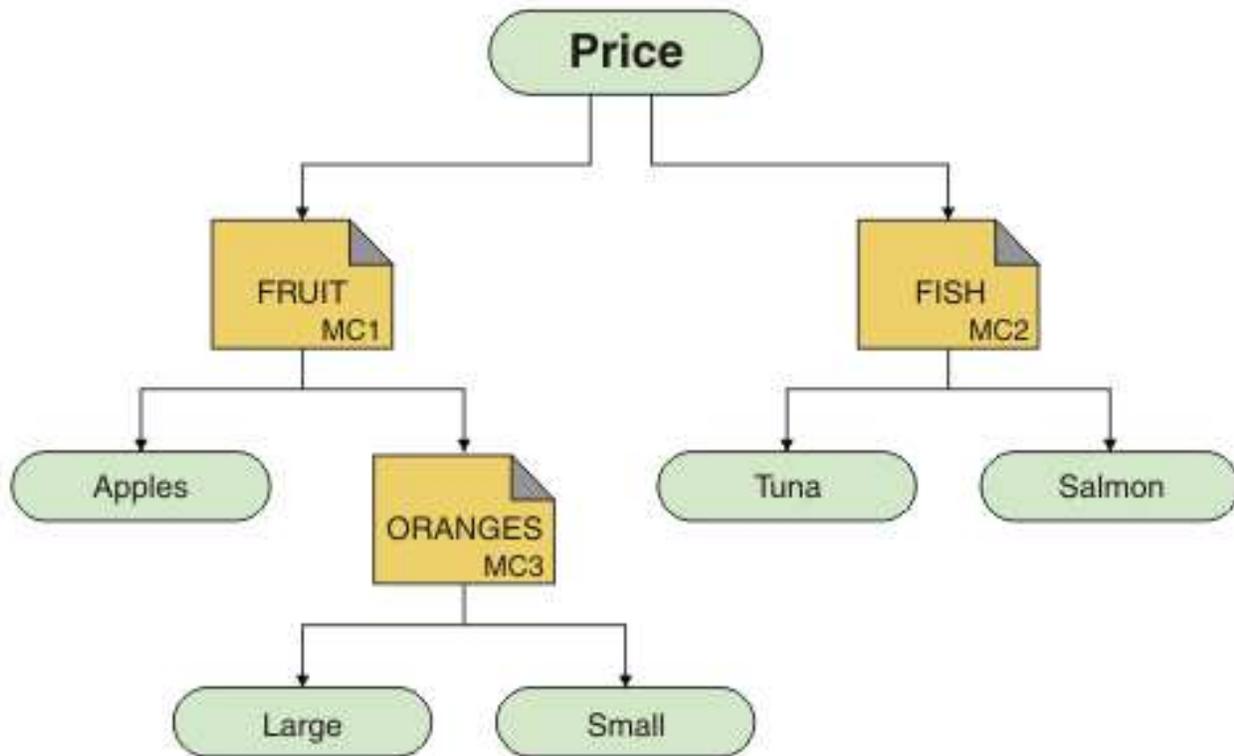
where *227.20.133.1*, *227.20.133.2*, and *227.20.133.3* are valid multicast addresses.

These topic definitions are used to create a topic tree as shown in the following diagram:

```
DEFINE TOPIC(FRUIT) TOPICSTRING('Price/FRUIT') MCAST(ENABLED) COMMINFO(MC1)
```

```
DEFINE TOPIC(FISH) TOPICSTRING('Price/FISH') MCAST(ENABLED) COMMINFO(MC2)
```

```
DEFINE TOPIC(ORANGES) TOPICSTRING('Price/FRUIT/ORANGES') MCAST(ENABLED)
COMMINFO(MC3)
```



While this kind of multicast topology is possible to create, it is not recommended because applications might not receive the data that they were expecting.

An application subscribing on 'Price/FRUIT/#' receives multicast transmission on the COMMINFO MC1 group address. The application expects to receive publications on all topics at or below that point in the topic tree.

However, the messages created by an application publishing on 'Price/FRUIT/ORANGES/Small' are not received by the subscriber because the messages are sent on the group address of COMMINFO MC3.

Using logs

There are various logs that you can use to help with problem determination and troubleshooting.

Use the following links to find out about the logs available for your platform and how to use them:

- "Error logs on Windows, UNIX and Linux systems" on page 1227
- "Error logs on IBM i" on page 1231
- "Error logs on HP Integrity NonStop Server" on page 1230

On z/OS error messages are written to:

- The z/OS system console
- The channel-initiator job log

It is possible to suppress or exclude some messages on both distributed and z/OS systems.

For details of suppressing some messages on distributed systems, see  Queue manager error logs (*WebSphere MQ V7.1 Installing Guide*).

On z/OS, if you are using the z/OS message processing facility to suppress messages, the console messages can be suppressed. For more information, see  WebSphere MQ for z/OS concepts (*WebSphere MQ V7.1 Product Overview Guide*). For information about error messages, console logs, and dumps on IBM WebSphere MQ for z/OS, see WebSphere MQ for z/OS concepts.

Related concepts:

“Troubleshooting and support” on page 1149

“Troubleshooting overview” on page 1149

“First Failure Support Technology (FFST)” on page 1315

“Using trace” on page 1234

Error logs on Windows, UNIX and Linux systems

About error log files, and an example.

At installation time, an errors subdirectory is created in the `/var/mqm` file path under UNIX and Linux systems, and in the installation directory, for example `C:\Program Files\IBM\WebSphere MQ\` file path under Windows systems. The errors subdirectory can contain up to three error log files named:

- AMQERR01.LOG
- AMQERR02.LOG
- AMQERR03.LOG

For more information about directories where log files are stored, see “Error log directories” on page 1229.

After you have created a queue manager, it creates three error log files when it needs them. These files have the same names as those files in the system error log directory. That is, AMQERR01, AMQERR02, and AMQERR03, and each has a default capacity of 2 MB (2 097 152 bytes). The capacity can be altered in the Extended queue manager properties page from the WebSphere MQ Explorer, or in the `QMErrorLog` stanza in the `qm.ini` file. These files are placed in the errors subdirectory in the queue manager data directory that you selected when you installed IBM WebSphere MQ or created your queue manager. The default location for the errors subdirectory is `/var/mqm/qmgrs/qmname` file path under UNIX and Linux systems, and `C:\Program Files\IBM\WebSphere MQ\qmgrs\qmname\errors` file path under Windows systems.

As error messages are generated, they are placed in AMQERR01. When AMQERR01 gets bigger than 2 MB (2 097 152 bytes) it is copied to AMQERR02. Before the copy, AMQERR02 is copied to AMQERR03.LOG. The previous contents, if any, of AMQERR03 are discarded.

The latest error messages are thus always placed in AMQERR01, the other files being used to maintain a history of error messages.

All messages relating to channels are also placed in the appropriate error files belonging to the queue manager, unless the queue manager is unavailable, or its name is unknown. In which case, channel-related messages are placed in the system error log directory.

To examine the contents of any error log file, use your usual system editor.

An example of an error log

Figure 147 on page 1228 shows an extract from a WebSphere MQ error log:

```
17/11/2004 10:32:29 - Process(2132.1) User(USER_1) Program(runmqchi.exe)
Host(HOST_1) Installation(Installation1)
VRMF(7.1.0.0) QMgr (A.B.C)
AMQ9542: Queue manager is ending.
```

EXPLANATION:

The program will end because the queue manager is quiescing.

ACTION:

None.

----- amqrimna.c : 931 -----

Figure 147. Sample WebSphere MQ error log

Operator messages

Operator messages identify normal errors, typically caused directly by users doing things like using parameters that are not valid on a command. Operator messages are national-language enabled, with message catalogs installed in standard locations.

These messages are written to the associated window, if any. In addition, some operator messages are written to the AMQERR01.LOG file in the queue manager directory, and others to the equivalent file in the system error log directory.

Error log access restrictions

Certain error log directories and error logs have access restrictions.

To gain the following access permissions, a user or application must be a member of the mqm group:

- Read and write access to all queue manager error log directories.
- Read and write access to all queue manager error logs.
- Write access to the system error logs.

If an unauthorized user or application attempts to write a message to a queue manager error log directory, the message is redirected to the system error log directory.

Ignoring error codes under UNIX and Linux systems

On UNIX and Linux systems, if you do not want certain error messages to be written to a queue manager error log, you can specify the error codes that are to be ignored using the QMErrorLog stanza.

For more information, see  Queue manager error logs (*WebSphere MQ V7.1 Installing Guide*).

Ignoring error codes under Windows systems

On Windows systems, if an error message has a severity of ERROR, the message is written to both the WebSphere MQ error log and the Windows Application Event Log. If you do not want certain error messages to be written to the Windows Application Event Log, you can specify the error codes that are to be ignored in the Windows registry.

Use the following registry key:

```
HKLM\Software\IBM\WebSphere MQ\Installation\MQ_INSTALLATION_NAME\IgnoredErrorCodes
```

where *MQ_INSTALLATION_NAME* is the installation name associated with a particular installation of IBM WebSphere MQ.

The value that you set it to is an array of strings delimited by the NULL character, with each string value relating to the error code that you want ignored from the error log. The complete list is terminated with a NULL character, which is of type REG_MULTI_SZ.

For example, if you want WebSphere MQ to exclude error codes AMQ3045, AMQ6055, and AMQ8079 from the Windows Application Event Log, set the value to:

```
AMQ3045\0AMQ6055\0AMQ8079\0\0
```

The list of messages you want to exclude is defined for all queue managers on the machine. Any changes you make to the configuration will not take effect until each queue manager is restarted.

Related concepts:

- “Troubleshooting and support” on page 1149
- “Using logs” on page 1226
- “Using trace” on page 1234
- “Problem determination on z/OS” on page 1265

Related reference:

- “Error logs on IBM i” on page 1231

Error log directories

WebSphere MQ uses a number of error logs to capture messages concerning its own operation of WebSphere MQ, any queue managers that you start, and error data coming from the channels that are in use. The location of the error logs depends on whether the queue manager name is known and whether the error is associated with a client.

The location the error logs are stored in depends on whether the queue manager name is known and whether the error is associated with a client. *MQ_INSTALLATION_PATH* represents the high level directory where WebSphere MQ is installed.

- If the queue manager name is known, the location of the error log is shown in Table 132.

Table 132. Queue manager error log directory

Platform	Directory
UNIX and Linux systems	<i>/var/mqm/qmgrs/qmname/errors</i>
Windows systems	<i>MQ_INSTALLATION_PATH\QMGRS\qmname\ERRORS\AMQERR01.LOG</i>

- If the queue manager name is not known, the location of the error log is shown in Table 133.

Table 133. System error log directory

Platform	Directory
UNIX and Linux systems	<i>/var/mqm/errors</i>
Windows systems	<i>MQ_INSTALLATION_PATH\QMGRS\@SYSTEM\ERRORS\AMQERR01.LOG</i>

- If an error has occurred with a client application, the location of the error log on the client is shown in Table 134 on page 1230.

Table 134. Client error log directory

Platform	Directory
UNIX and Linux systems	/var/mqm/errors
Windows systems	MQ_DATA_PATH\ERRORS\AMQERR01.LOG

In WebSphere MQ for Windows, an indication of the error is also added to the Application Log, which can be examined with the Event Viewer application provided with Windows systems.

Early errors

There are a number of special cases where these error logs have not yet been established and an error occurs. WebSphere MQ attempts to record any such errors in an error log. The location of the log depends on how much of a queue manager has been established.

If, because of a corrupt configuration file for example, no location information can be determined, errors are logged to an errors directory that is created at installation time on the root directory (/var/mqm or C:\Program Files\IBM\WebSphere MQ).

If WebSphere MQ can read its configuration information, and can access the value for the Default Prefix, errors are logged in the errors subdirectory of the directory identified by the Default Prefix attribute. For example, if the default prefix is C:\Program Files\IBM\WebSphere MQ, errors are logged in C:\Program Files\IBM\WebSphere MQ\errors.

For further information about configuration files, see  Changing IBM WebSphere MQ and queue manager configuration information (*WebSphere MQ V7.1 Installing Guide*).

Note: Errors in the Windows Registry are notified by messages when a queue manager is started.

Error logs on HP Integrity NonStop Server

Use this information to understand the IBM WebSphere MQ client on HP Integrity NonStop Server error logs, together with an example.

At installation time, an errors subdirectory is created in the <mqpath>/var/mqm file path. The errors subdirectory can contain up to three error log files named:

- AMQERR01.LOG
- AMQERR02.LOG
- AMQERR03.LOG

As error messages are generated, they are written to AMQERR01.LOG. When AMQERR01.LOG gets bigger than 2 MB (2 097 152 bytes), it is copied to AMQERR02.LOG. Before the copy, AMQERR02.LOG is copied to AMQERR03.LOG. The previous contents, if any, of AMQERR03.LOG are discarded.

The latest error messages are therefore always placed in AMQERR01.LOG. The other log files are used to maintain a history of error messages.

To examine the contents of any error log file, use your system editor. The contents of the log files can read by any user, but write access requires the user to be a member of the mqm group.

An example of an error log

Figure 148 on page 1231 shows an extract from a WebSphere MQ error log:

```

04/30/13 06:18:22 - Process(320406477.1) User(MYUSER) Program(nssfcp_s_c)
                    Host(myhost)
                    VRMF(7.1.0.0)
AMQ9558: The remote channel 'SYSTEM.DEF.SVRCONN' on host 'hostname
(x.x.x.x)(1414)' is not currently available.

EXPLANATION:
The channel program ended because an instance of channel 'SYSTEM.DEF.SVRCONN'
could not be started on the remote system. This could be for one of the
following reasons:

The channel is disabled.

The remote system does not have sufficient resources to run another instance of
the channel.

In the case of a client-connection channel, the limit on the number of
instances configured for the remote server-connection channel was reached.

ACTION:
Check the remote system to ensure that the channel is able to run. Try the
operation again.
----- cmqxrft.c : 504 -----

```

Figure 148. Sample WebSphere MQ error log

Error logs on IBM i

Use this information to understand the IBM WebSphere MQ for IBM i error logs.

IBM WebSphere MQ uses a number of error logs to capture messages concerning the operation of IBM WebSphere MQ itself, any queue managers that you start, and error data coming from the channels that are in use.

At installation time, a /QIBM/UserData/mqm/errors subdirectory is created in the IFS.

The location of the error logs depends on whether the queue manager name is known.

In the IFS:

- If the queue manager name is known and the queue manager is available, error logs are located in:
/QIBM/UserData/mqm/qmgrs/*qmname*/errors
- If the queue manager is not available, error logs are located in:
/QIBM/UserData/mqm/errors

You can use the system utility EDTF to browse the errors directories and files. For example:

```
EDTF '/QIBM/UserData/mqm/errors'
```

Alternatively, you can use option 23 against the queue manager from the WRKMQM panel.

The errors subdirectory can contain up to three error log files named:

- AMQERR01.LOG
- AMQERR02.LOG
- AMQERR03.LOG

After you have created a queue manager, three error log files are created when they are needed by the queue manager. These files have the same names as the /QIBM/UserData/mqm/errors ones, that is

AMQERR01, AMQERR02, and AMQERR03, and each has a capacity of 2 MB (2 097 152 bytes). The files are placed in the errors subdirectory of each queue manager that you create, that is `/QIBM/UserData/mqm/qmgrs/qmname/errors`.

As error messages are generated, they are placed in AMQERR01. When AMQERR01 gets bigger than 2 MB (2 097 152 bytes), it is copied to AMQERR02. Before the copy, AMQERR02 is copied to AMQERR03.LOG. The previous contents, if any, of AMQERR03 are discarded.

The latest error messages are thus always placed in AMQERR01, the other files being used to maintain a history of error messages.

All messages relating to channels are also placed in the appropriate errors files of the queue manager, unless the name of their queue manager is unknown or the queue manager is unavailable. When the queue manager name is unavailable or its name cannot be determined, channel-related messages are placed in the `/QIBM/UserData/mqm/errors` subdirectory.

To examine the contents of any error log file, use your system editor, EDTF, to view the stream files in the IFS.

Note:

1. Do not change ownership of these error logs.
2. If any error log file is deleted, it is automatically re-created when the next error message is logged.

Early errors

There are a number of special cases where the error logs have not yet been established and an error occurs. IBM WebSphere MQ attempts to record any such errors in an error log. The location of the log depends on how much of a queue manager has been established.

If, because of a corrupted configuration file, for example, no location information can be determined, errors are logged to an errors directory that is created at installation time.

If both the IBM WebSphere MQ configuration file and the `DefaultPrefix` attribute of the `AllQueueManagers` stanza are readable, errors are logged in the errors subdirectory of the directory identified by the `DefaultPrefix` attribute.

Operator messages

Operator messages identify normal errors, typically caused directly by users doing things like using parameters that are not valid on a command. Operator messages are national language enabled, with message catalogs installed in standard locations.

These messages are written to the job log, if any. In addition, some operator messages are written to the `AMQERR01.LOG` file in the queue manager directory, and others to the `/QIBM/UserData/mqm/errors` directory copy of the error log.

An example IBM WebSphere MQ error log

Figure 149 on page 1233 shows a typical extract from a IBM WebSphere MQ error log.

```

*****Beginning of data*****
07/19/02 11:15:56 AMQ9411: Repository manager ended normally.

EXPLANATION:
Cause . . . . . : The repository manager ended normally.
Recovery . . . . : None.
Technical Description . . . . . : None.
-----
07/19/02 11:15:57 AMQ9542: Queue manager is ending.

EXPLANATION:
Cause . . . . . : The program will end because the queue manager is quiescing.
Recovery . . . . : None.
Technical Description . . . . . : None.
----- amqrimna.c : 773 -----
07/19/02 11:16:00 AMQ8004: WebSphere MQ queue manager 'mick' ended.
EXPLANATION:
Cause . . . . . : WebSphere MQ queue manager 'mick' ended.
Recovery . . . . : None.
Technical Description . . . . . : None.
-----
07/19/02 11:16:48 AMQ7163: WebSphere MQ job number 18429 started.

EXPLANATION:
Cause . . . . . : This job has started to perform work for Queue Manager
                  mick, The job's PID is 18429 the CCSID is 37. The job name is
                  582775/MQUSER/AMQZXMA0.
Recovery . . . . : None
-----
07/19/02 11:16:49 AMQ7163: WebSphere MQ job number 18430 started.

EXPLANATION:
Cause . . . . . : This job has started to perform work for Queue Manager
                  mick, The job's PID is 18430 the CCSID is 0. The job name is
                  582776/MQUSER/AMQZFUMA.
Recovery . . . . : None
-----
07/19/02 11:16:49 AMQ7163: WebSphere MQ job number 18431 started.

EXPLANATION:
Cause . . . . . : This job has started to perform work for Queue Manager
                  mick, The job's PID is 18431 the CCSID is 37. The job name is
                  582777/MQUSER/AMQZXMAX.
Recovery . . . . : None
-----
07/19/02 11:16:50 AMQ7163: WebSphere MQ job number 18432 started.

EXPLANATION:
Cause . . . . . : This job has started to perform work for Queue Manager
                  mick, The job's PID is 18432 the CCSID is 37. The job name is
                  582778/MQUSER/AMQALMPX.
Recovery . . . . : None
-----

```

Figure 149. Extract from a IBM WebSphere MQ error log

Related concepts:

“Error logs on Windows, UNIX and Linux systems” on page 1227

“Troubleshooting and support” on page 1149

“Using logs” on page 1226

“Using trace”

“Problem determination on z/OS” on page 1265

Using trace

You can use different types of trace to help you with problem determination and troubleshooting.

Use the following links to find out about the different types of trace, and how to run trace for your platform:

- “Using trace on Windows”
- “Using trace on UNIX and Linux systems” on page 1236
- “Using trace on WebSphere MQ server on IBM i” on page 1240
- “Using trace for problem determination on z/OS” on page 1244
- “Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255
- “Tracing WebSphere MQ classes for JMS programs” on page 1256 (*WebSphere MQ V7.1 Programming Guide*)
- “Tracing WebSphere MQ classes for Java programs” on page 1260 (*WebSphere MQ V7.1 Programming Guide*)
- “Java diagnostics” on page 1262

Related concepts:

“Troubleshooting and support” on page 1149

“Troubleshooting overview” on page 1149

“Using logs” on page 1226

“First Failure Support Technology (FFST)” on page 1315

Related tasks:

“Contacting IBM Software Support” on page 1346

Using trace on Windows

Use the **strmqtrc** and **endmqtrc** commands or the IBM WebSphere MQ Explorer interface to start and end tracing.

Windows uses the following commands for the client trace facility:

strmqtrc

to start tracing

endmqtrc

to end tracing

The output files are created in the MQ_DATA_PATH/trace directory.

Trace files on IBM WebSphere MQ for Windows

Trace files are named AMQppppp.qq.TRC where the variables are:

ppppp The ID of the process reporting the error.

qq A sequence number, starting at 0. If the full file name exists, this value is incremented by one until a unique trace file name is found. A trace file name can exist if a process is reused.

Note:

1. The process identifier can contain fewer, or more, digits than shown in the example.
2. There is one trace file for each process running as part of the entity being traced.

To format or view a trace file, you must be either the creator of the trace file, or a member of the mqm group.

SSL trace files have the names AMQ.SSL.TRC and AMQ.SSL.TRC.1. You cannot format SSL trace files; send them unchanged to IBM support.

How to start and stop a trace

Enable or modify tracing using the **strmqtrc** control command (see  *strmqtrc (WebSphere MQ V7.1 Reference)*). To stop tracing, use the **endmqtrc** control command (see  *endmqtrc (WebSphere MQ V7.1 Reference)*).

In IBM WebSphere MQ for Windows systems, you can also start and stop tracing using the IBM WebSphere MQ Explorer, as follows:

1. Start the IBM WebSphere MQ Explorer from the **Start** menu.
2. In the Navigator View, right-click the **WebSphere MQ** tree node, and select **Trace...** The Trace Dialog is displayed.
3. Click **Start** or **Stop** as appropriate.

Selective component tracing

Use the **-t** and **-x** options to control the amount of trace detail to record. By default, all trace points are enabled. You can specify the points that you do not want to trace using the **-x** option. So if, for example, you want to trace only data flowing over communications networks, use:

```
strmqtrc -x all -t comms
```

For detailed information about the trace command, see  *strmqtrc (WebSphere MQ V7.1 Reference)*.

Selective process tracing

Use the **-p** option of the **strmqtrc** command control to restrict trace generation to specified named processes. For example, to trace all threads that result from any running process called amqxxx.exe, use the following command:

```
strmqtrc -p amqxxx.exe
```

For detailed information about the trace command, see  *strmqtrc (WebSphere MQ V7.1 Reference)*.

Related concepts:

“Using trace on UNIX and Linux systems”

“Using trace on WebSphere MQ server on IBM i” on page 1240

“Using trace for problem determination on z/OS” on page 1244

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

“Java diagnostics” on page 1262

Using trace on UNIX and Linux systems

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

UNIX and Linux systems use the following commands for the WebSphere MQ MQI client trace facility:

strmqtrc

to start tracing

endmqtrc

to end tracing

dspmqtrc <filename>

to display a formatted trace file

The trace facility uses a number of files, which are:

- One file for each entity being traced, in which trace information is recorded
- One additional file on each machine, to provide a reference for the shared memory used to start and end tracing
- One file to identify the semaphore used when updating the shared memory

Files associated with trace are created in a fixed location in the file tree, which is `/var/mqm/trace`.

All client tracing takes place to files in this directory.

You can handle large trace files by mounting a temporary file system over this directory.

On AIX you can use AIX system trace in addition to using the **strmqtrc** and **endmqtrc** commands. For more information, see “Tracing with the AIX system trace” on page 1238.

Trace files on IBM WebSphere MQ for UNIX and Linux systems

Trace files are created in the directory `/var/mqm/trace`.

Note: You can accommodate the production of large trace files by mounting a temporary file system over the directory that contains your trace files. Alternatively, rename the trace directory and create the symbolic link `/var/mqm/trace` to a different directory.

Trace files are named `AMQppppp.qq.TRC` where the variables are:

ppppp The ID of the process reporting the error.

qq A sequence number, starting at 0. If the full file name exists, this value is incremented by one until a unique trace file name is found. A trace file name can exist if a process is reused.

Note:

1. The process identifier can contain fewer, or more, digits than shown in the example.
2. There is one trace file for each process running as part of the entity being traced.

To format or view a trace file, you must be either the creator of the trace file, or a member of the mqm group.

SSL trace files have the names AMQ.SSL.TRC and AMQ.SSL.TRC.1. You cannot format SSL trace files; send them unchanged to IBM support.

How to start and stop a trace

In IBM WebSphere MQ for UNIX and Linux systems, you enable or modify tracing using the **strmqtrc** control command (see  [strmqtrc \(WebSphere MQ V7.1 Reference\)](#)). To stop tracing, you use the **endmqtrc** control command (see  [endmqtrc \(WebSphere MQ V7.1 Reference\)](#)). On IBM WebSphere MQ for Linux (x86 and x86-64 platforms) systems, you can alternatively use the IBM WebSphere MQ Explorer to start and stop tracing. However, you can trace only everything using the function provided, equivalent to using the commands **strmqtrc -e** and **endmqtrc -e**.

Trace output is unformatted; use the **dspmqtrc** control command to format trace output before viewing. For example, to format all trace files in the current directory use the following command:

```
dspmqtrc *.TRC
```

For detailed information about the control command, **dspmqtrc**, see  [dspmqtrc \(WebSphere MQ V7.1 Reference\)](#).

Selective component tracing on WebSphere MQ for UNIX and Linux systems

Use the **-t** and **-x** options to control the amount of trace detail to record. By default, all trace points are enabled. Specify the points you do not want to trace using the **-x** option. If, for example, you want to trace, for queue manager QM1, only output data associated with using Secure Sockets Layer (SSL) channel security, use:

```
strmqtrc -m QM1 -t ssl
```

For detailed information about the trace command, see  [strmqtrc \(WebSphere MQ V7.1 Reference\)](#).

Selective component tracing on WebSphere MQ for AIX

Use the environment variable **MQS_TRACE_OPTIONS** to activate the high detail and parameter tracing functions individually.

Because **MQS_TRACE_OPTIONS** enables tracing to be active without high detail and parameter tracing functions, you can use it to reduce the effect on performance and trace size when you are trying to reproduce a problem with tracing switched on.

Only set the environment variable **MQS_TRACE_OPTIONS** if you have been instructed to do so by your service personnel.

Typically **MQS_TRACE_OPTIONS** must be set in the process that starts the queue manager, and before the queue manager is started, or it is not recognized. Set **MQS_TRACE_OPTIONS** before tracing starts. If it is set after tracing starts it is not recognized.

Selective process tracing on WebSphere MQ for UNIX and Linux systems

Use the **-p** option of the **strmqtrc** command control to restrict trace generation to specified named processes. For example, to trace all threads that result from any running process called **amqxxx**, use the following command:

```
strmqtrc -p amqxxx
```

For detailed information about the trace command, see  `strmqtrc` (*WebSphere MQ V7.1 Reference*).

Related concepts:

“Using trace on WebSphere MQ server on IBM i” on page 1240

“Using trace for problem determination on z/OS” on page 1244

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

“Java diagnostics” on page 1262

Related reference:

“Using trace on Windows” on page 1234

Tracing with the AIX system trace

In addition to the WebSphere MQ trace, WebSphere MQ for AIX users can use the standard AIX system trace.

AIX system tracing is a two-step process:

1. Gathering the data
2. Formatting the results

WebSphere MQ uses two trace hook identifiers:

X'30D' This event is recorded by WebSphere MQ on entry to or exit from a subroutine.

X'30E' This event is recorded by WebSphere MQ to trace data such as that being sent or received across a communications network.

Trace provides detailed execution tracing to help you to analyze problems. IBM service support personnel might ask for a problem to be re-created with trace enabled. The files produced by trace can be **very** large so it is important to qualify a trace, where possible. For example, you can optionally qualify a trace by time and by component.

There are two ways to run trace:

1. Interactively.

The following sequence of commands runs an interactive trace on the program `myprog` and ends the trace.

```
trace -j30D,30E -o trace.file
->!myprog
->q
```

2. Asynchronously.

The following sequence of commands runs an asynchronous trace on the program `myprog` and ends the trace.

```
trace -a -j30D,30E -o trace.file
myprog
trcstop
```

You can format the trace file with the command:

```
trcrpt -t MQ_INSTALLATION_PATH/lib/amqtrc.fmt trace.file > report.file
```

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

`report.file` is the name of the file where you want to put the formatted trace output.

Note: All WebSphere MQ activity on the machine is traced while the trace is active.

Using trace on HP Integrity NonStop Server

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file.

Use the following commands on the IBM WebSphere MQ client for HP Integrity NonStop Server system to use the IBM WebSphere MQ client trace facility:

strmqtrc

To start tracing

endmqtrc

To end tracing

dspmqtrc <filename>

To display a formatted trace file

The trace facility creates a file for each entity that is being traced. The trace files are created in a fixed location, which is `<mqpath>/var/mqm/trace`. You can handle large trace files by mounting a temporary file system over this directory.

Trace files are named `AMQ.nnn.xx.ppp.qq.TRC` where:

nnn The name of the process.

xx The processor number on which the process is running.

ppp The PIN of the process that you are tracing.

qq A sequence number, starting at 0. If the full file name exists, this value is incremented by one until a unique trace file name is found. A trace file name can exist if a process is reused.

Note:

1. Each field can contain fewer, or more, digits than shown in the example.
2. There is one trace file for each process that is running as part of the entity that is being traced.

Trace files are created in a binary format. To format or view a trace file use the **dspmqtrc** command, you must be either the creator of the trace file, or a member of the `mqm` group. For example, to format all trace files in the current directory use the following command:

```
dspmqtrc *.TRC
```

For more information about the control command **dspmqtrc**, see  [dspmqtrc \(WebSphere MQ V7.1 Reference\)](#).

How to start and stop a trace

On IBM WebSphere MQ client for HP Integrity NonStop Server systems, you can enable or modify

tracing by using the **strmqtrc** control command, for more information, see  [strmqtrc \(WebSphere MQ V7.1 Reference\)](#). To stop tracing, use the **endmqtrc** control command, for more information, see

 [endmqtrc \(WebSphere MQ V7.1 Reference\)](#).

The control commands **strmqtrc** and **endmqtrc** affect tracing only for those processes that are running in one specific processor. By default, this processor is the same as the one in your OSS shell. To enable or end tracing for processes that are running in another processor, you must precede the **strmqtrc** or **endmqtrc** commands with `run -cpu=n` at an OSS shell command prompt, where `n` is the processor number. Here is an example of how to enter the **strmqtrc** command at an OSS shell command prompt:

```
run -cpu=2 strmqtrc
```

This command enables tracing for all processes that are running in processor 2.

The `-m` option to select a queue manager is not relevant for use on the IBM WebSphere MQ client for HP Integrity NonStop Server. Specifying the `-m` option produces an error.

Use the `-t` and `-x` options to control the amount of trace detail to record. By default, all trace points are enabled. Specify the points that you do not want to trace by using the `-x` option.

Using trace on WebSphere MQ server on IBM i

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

There are two stages in using trace:

1. Decide whether you want early tracing. Early tracing lets you trace the creation and startup of queue managers. Note, however, that early trace can easily generate large amounts of trace, because it is implemented by tracing all jobs for all queue managers. To enable early tracing, use TRCMQM with the TRCEARLY parameter set to `*YES`.
2. Start tracing work using TRCMQM `*ON`. To stop the trace, you have two options:
 - TRCMQM `*OFF`, to stop collecting trace records for a queue manager. The trace records are written to files in the `/QIBM/UserData/mqm/t/trace` directory.
 - TRCMQM `*END`, to stop collecting trace records for all queue managers and to disable early trace. This option ignores the value of the TRCEARLY parameter.

Specify the level of detail you want, using the TRCLEVEL parameter set to one of the following values:

`*DFT` For minimum-detail level for flow processing trace points.

`*DETAIL`

For high-detail level for flow processing trace points.

`*PARMS`

For default-detail level for flow processing trace points.

Specify the type of trace output you want, using the OUTPUT parameter set to one of the following values:

`*MQM`

Collect binary IBM WebSphere MQ trace output in the directory specified by the TRCDIR parameter. This value is the default value.

`*MQMFMT`

Collect formatted IBM WebSphere MQ trace output in the directory specified by the TRCDIR parameter.

`*PEX` Collect Performance Explorer (PEX) trace output

`*ALL` Collect both IBM WebSphere MQ unformatted trace and PEX trace output

Selective trace

You can reduce the amount of trace data being saved, improving runtime performance, using the command TRCMQM with F4=prompt, then F9 to customize the TRCTYPE and EXCLUDE parameters:

TRCTYPE

Specifies the type of trace data to store in the trace file. If you omit this parameter, all trace points except those trace points specified in EXCLUDE are enabled.

EXCLUDE

Specifies the type of trace data to omit from the trace file. If you omit this parameter, all trace points specified in TRCTYPE are enabled.

The options available on both TRCTYPE and EXCLUDE are:

***ALL (TRCTYPE only)**

All the trace data as specified by the following keywords is stored in the trace file.

trace-type-list

You can specify more than one option from the following keywords, but each option can occur only once.

***API** Output data for trace points associated with the MQI and major queue manager components.

***CMTRY**

Output data for trace points associated with comments in the IBM WebSphere MQ components.

***COMMS**

Output data for trace points associated with data flowing over communications networks.

***CSDATA**

Output data for trace points associated with internal data buffers in common services.

***CSFLOW**

Output data for trace points associated with processing flow in common services.

***LQMDATA**

Output data for trace points associated with internal data buffers in the local queue manager.

***LQMFLOW**

Output data for trace points associated with processing flow in the local queue manager.

***OTHDATA**

Output data for trace points associated with internal data buffers in other components.

***OTHFLOW**

Output data for trace points associated with processing flow in other components.

***RMTDATA**

Output data for trace points associated with internal data buffers in the communications component.

***RMTFLOW**

Output data for trace points associated with processing flow in the communications component.

***SVCDATA**

Output data for trace points associated with internal data buffers in the service component.

***SVCFLOW**

Output data for trace points associated with processing flow in the service component.

***VSNDATA**

Output data for trace points associated with the version of IBM WebSphere MQ running.

Wrapping trace

Use the MAXSTG parameter to wrap trace, and to specify the maximum size of storage to be used for the collected trace records.

The options are:

***DFT** Trace wrapping is not enabled. For each job, trace data is written to a file with the suffix .TRC until tracing is stopped.

maximum-K-bytes

Trace wrapping is enabled. When the trace file reaches its maximum size, it is renamed with the suffix .TRS, and a new trace file with suffix .TRC is opened. Any existing .TRS file is deleted. Specify a value in the range 1 through 16 000.

Formatting trace output

To format any trace output:

- Enter the QShell
- Enter the command

```
/QSYS.LIB/QMQM.LIB/DSPMQTRC.PGM [-t Format] [-h] [-s]  
[-o OutputFileName] InputFileName
```

where:

InputFileName

Is a required parameter specifying the name of the file containing the unformatted trace. For example /QIBM/UserData/mqm/trace/AMQ12345.TRC.

-t *FormatTemplate*

Specifies the name of the template file containing details of how to display the trace. The default value is /QIBM/ProdData/mqm/lib/amqtrc.fmt.

-h Omit header information from the report.

-s Extract trace header and put to stdout.

-o *output_filename*

The name of the file into which to write formatted data.

You can also specify `dspmqtrc *` to format all trace.

Related concepts:

“Using trace on UNIX and Linux systems” on page 1236

“Using trace for problem determination on z/OS” on page 1244

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

“Java diagnostics” on page 1262

Related reference:

“Using trace on Windows” on page 1234

Using trace on WebSphere MQ client on IBM i

On IBM i, there is no Control Language (CL) command to capture the trace when using a stand-alone WebSphere MQ MQI client. STRMQTRC and ENDMQTRC programs can be used to enable and disable the trace.

Example for start trace:

```
CALL PGM(QMQM/STRMQTRC) PARM('-e' '-t' 'all' '-t' 'detail')
```

Where `-e` option requests early tracing of all the process `-t` option for trace type

To end the trace

```
CALL PGM(QMQM/ENDMQTRC) PARM('-e')
```

- Optional parameters:

-t *TraceType*

The points to trace and the amount of trace detail to record. By default all trace points are enabled and a default-detail trace is generated.

Alternatively, you can supply one or more of the options in Table 1. For each *TraceType* value you specify, including `-t all`, specify either `-t parms` or `-t detail` to obtain the appropriate level of trace detail. If you do not specify either `-t parms` or `-t detail` for any particular trace type, only a default-detail trace is generated for that trace type.

If you supply multiple trace types, each must have its own `-t` flag. You can include any number of `-t` flags, if each has a valid trace type associated with it.

It is not an error to specify the same trace type on multiple `-t` flags.

See the following table for allowed values for *TraceType*.

Table 135. *TraceType* values

Value	Description
all	Output data for every trace point in the system (the default). Using <i>all</i> activates tracing at default detail level.
api	Output data for trace points associated with the message queue interface (MQI) and major queue manager components.
commentary	Output data for trace points associated with comments in the WebSphere MQ components.
comms	Output data for trace points associated with data flowing over communications networks.
csdata	Output data for trace points associated with internal data buffers in common services.
csflows	Output data for trace points associated with processing flow in common services.
detail	Activate tracing at high-detail level for flow processing trace points.
lqmdata	Output data for trace points associated with internal data buffers in the local queue manager.
lqmflows	Output data for trace points associated with processing flow in the local queue manager.
otherdata	Output data for trace points associated with internal data buffers in other components.
otherflows	Output data for trace points associated with processing flow in other components.
parms	Activate tracing at default-detail level for flow processing trace points.
remotedata	Output data for trace points associated with internal data buffers in the communications component.
remoteflows	Output data for trace points associated with processing flow in the communications component.
servicedata	Output data for trace points associated with internal data buffers in the service component.
serviceflows	Output data for trace points associated with processing flow in the service component.
versiondata	Output data for trace points associated with the version of WebSphere MQ running.

-x *TraceType*

The points not to trace. By default all trace points are enabled and a default-detail trace is generated. The *TraceType* values you can specify are the same as the values listed for the `-t` flag in Table 1.

You can use the `-x` flag with *TraceType* values to exclude those trace points you do not want to record. Excluding specified trace points is useful in reducing the amount of trace produced.

If you supply multiple trace types, each must have its own `-x` flag. You can include any number of `-x` flags, if each has a valid *TraceType* associated with it.

-s Reports the tracing options that are currently in effect. You must use this parameter on its own with no other parameters.

A limited number of slots are available for storing trace commands. When all slots are in use, then no more trace commands can be accepted unless they replace an existing slot. Slot numbers are not fixed, so if the command in slot number 0 is removed, for example by an `endmqtrc` command, then all the other slots move up, with slot 1 becoming slot 0, for example. An asterisk (*) in a field means that no value is defined, and is equivalent to the asterisk wildcard.

-l *MaxSize*

The maximum size of a trace file (AMQppppp.qq.TRC) in megabytes (MB). For example, if you specify a *MaxSize* of 1, the size of the trace is limited to 1 MB.

When a trace file reaches the specified maximum, it is renamed to AMQppppp.qq.TRS and a new AMQppppp.qq.TRC file is started. If a previous copy of an AMQppppp.qq.TRS file exists, it is deleted.

The highest value that *MaxSize* can be is 2048 MB.

-e Requests early tracing of all processes

For more details see the  **strmqtrc** command (*WebSphere MQ V7.1 Reference*)

- To end the trace:

```
/QSYS.LIB/QMQM.LIB/ENDMQTRC.PGM [-e] [-a]
```

where:

-e Ends early tracing of all processes.

Using **endmqtrc** with no parameters has the same effect as **endmqtrc -e**. You cannot specify the **-e** flag with the **-m** flag, the **-i** flag, or the **-p** flag.

-a Ends all tracing.

For more details see the  **endmqtrc** command (*WebSphere MQ V7.1 Reference*)

- To display a formatted trace file:

```
/QSYS.LIB/QMQM.LIB/DSPMQTRC.pgm
```

To examine FFST files, see the First Failure Support Technology (FFST) for WebSphere MQ for IBM i.

Related concepts:

“Using trace on UNIX and Linux systems” on page 1236

“Using trace for problem determination on z/OS”

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

“Java diagnostics” on page 1262

Related reference:

“Using trace on Windows” on page 1234

Using trace for problem determination on z/OS

There are different trace options that can be used for problem determination with WebSphere MQ. Use this topic to understand the different options and how to control trace.

The trace facilities available with WebSphere MQ for z/OS are:

- The user parameter (or API) trace
- The IBM internal trace used by the support center
- The channel initiator trace
- The line trace

Use the following links to find out how to collect and interpret the data produced by the user parameter trace, and describes how to produce the IBM internal trace for use by the IBM support center. There is also information about the other trace facilities that you can use with WebSphere MQ.

- Controlling the GTF for your z/OS system
- Controlling the WebSphere MQ trace for each queue manager subsystem for which you want to collect data
- “Formatting and identifying the control block information” on page 1247

- “Interpreting the trace information” on page 1248
- “The Log and Trace Analyzer tool” on page 1250

If trace data is not produced, check the following:

- Was the GTF started correctly, specifying EIDs 5E9, 5EA, and 5EE on the USRP option?
- Was the START TRACE(GLOBAL) command entered correctly, and were the relevant classes specified?

For more information about other trace options available on z/OS, see “Other types of trace” on page 1250.

Related concepts:

“Using trace on UNIX and Linux systems” on page 1236

“Using trace on WebSphere MQ server on IBM i” on page 1240

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

“Java diagnostics” on page 1262

Related reference:

“Using trace on Windows” on page 1234

The MQI call and user parameter, and z/OS generalized trace facility (GTF)

Use this topic to understand how to control GTF and WebSphere MQ trace.

You can obtain information about MQI calls and user parameters passed by some WebSphere MQ calls on entry to, and exit from, WebSphere MQ. To do this, use the global trace in conjunction with the z/OS generalized trace facility (GTF).

Controlling the GTF:

Use this topic to understand how to start and stop the GTF.

- Starting the GTF
- Stopping the GTF

Starting the GTF

When you start the GTF, specify the USRP option. You are prompted to enter a list of event identifiers (EIDs). The EIDs used by WebSphere MQ are:

5E9 To collect information about control blocks on entry to WebSphere MQ

5EA To collect information about control blocks on exit from WebSphere MQ

Sometimes, if an error occurs that you cannot solve yourself, you might be asked by your IBM support center to supply other, internal, trace information for them to analyze. The additional type of trace is:

5EE To collect information internal to WebSphere MQ

You can also use the JOBNAMEP option, specifying the batch, CICS, IMS, or TSO job name, to limit the trace output to specific jobs. Figure 150 on page 1246 illustrates sample startup for the GTF, specifying the four EIDs, and a jobname. The lines shown in bold type **like this** are the commands that you enter at the console; the other lines are prompts and responses.

For more information about starting the GTF trace, see the *MVS Diagnosis: Tools and Service Aids* manual.

```

START GTFxx.yy
  EHASP100 GTFxx.yy ON STCINRDR
  EHASP373 GTFxx.yy STARTED
*01 AHL100A SPECIFY TRACE OPTIONS
  R 01,TRACE=JOBNAMEP,USRP
  TRACE=JOBNAMEP,USRP
  IEE600I REPLY TO 01 IS;TRACE=JOBNAMEP,USRP
*02 ALH101A SPECIFY TRACE EVENT KEYWORDS - JOBNAME=,USR=
  R 02,JOBNAME=(xxxxMSTR,xxxxCHIN,zzzzzzzz),USR=(5E9,5EA,5EE)
  JOBNAME=(xxxxMSTR,xxxxCHIN,zzzzzzzz),USR=(5E9,5EA,5EE)
  IEE600I REPLY TO 02 IS;JOBNAME=(xxxxMSTR,xxxxCHIN,zzzzzzzz),USR=(5E9,5EA,5EE)
*03 ALH102A CONTINUE TRACE DEFINITION OR REPLY END
  R 03,END
  END
  IEE600I REPLY TO 03 IS;END
  AHL103I TRACE OPTIONS SELECTED-USR=(5E9,5EA,5EE)
  AHL103I JOBNAME=(xxxxMSTR,xxxxCHIN,zzzzzzzz)
*04 AHL125A RESPECIFY TRACE OPTIONS OR REPLY U
  R 04,U
  U
  IEE600I REPLY TO 04 IS;U
  AHL031I GTF INITIALIZATION COMPLETE

```

where:

xx is the name of the GTF procedure to use (optional), and yy is an identifier for this occurrence of GTF trace.

xxxx is the name of the queue manager and zzzzzzzz is a batch job or CICS region name. Up to 5 job names can be listed.

Figure 150. Example startup of GTF to use with the WebSphere MQ trace

When using GTF, specify the primary job name (CHINIT, CICS, or batch) in addition to the queue manager name (xxxxMSTR).

Stopping the GTF

When you stop the GTF, you must specify the additional identifier (**yy**) used at startup. Figure 151 illustrates a sample stop command for the GTF. The commands shown in bold type **like this** are the commands that you enter at the console.

```
STOP yy
```

Figure 151. Example of GTF Stop command to use with the WebSphere MQ trace

Related information:

 [Generating WebSphere MQ GTF trace on IBM z/OS](#)

Controlling the trace within WebSphere MQ:

WebSphere MQ for z/OS trace is controlled using MQSC commands. Use this topic to understand how to control the trace, and the type of trace information that is output.

Use the START TRACE command, specifying type GLOBAL to start writing WebSphere MQ records to the GTF. You must also specify dest(GTF), for example in the following command:

```
/cpf start trace(G)class(2,3)dest(GTF)
```

To define the events that you want to produce trace data for, use one or more of the following classes:

CLASS	Event traced
2	Record the MQI call and MQI parameters when a completion code other than MQRC_NONE is detected.
3	Record the MQI call and MQI parameters on entry to and exit from the queue manager.

After the trace has started, you can display information about, alter the properties of, and stop, the trace with the following commands:

- DISPLAY TRACE
- ALTER TRACE
- STOP TRACE

To use any of the trace commands, you must have one of the following:

- Authority to issue start and stop trace commands (trace authority)
- Authority to issue the display trace command (display authority)

Note:

1. The trace commands can also be entered through the initialization input data sets.
2. The trace information produced will also include details of syncpoint flows - for example PREPARE and COMMIT.

For information about these commands, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) manual.

Formatting and identifying the control block information:

After capturing a trace, the output must be formatted and the WebSphere MQ control blocks identified.

- Formatting the information
- Identifying the control blocks associated with WebSphere MQ
- Identifying the event identifier associated with the control block

Formatting the information

To format the user parameter data collected by the global trace, use either the batch job shown in Figure 152 on page 1248 or the IPCS GTFTRACE USR(*xxx*) command, where *xxx* is:

- 5E9** To format information about control blocks on entry to WebSphere MQ MQI calls
- 5EA** To format information about control blocks on exit from WebSphere MQ MQI calls
- 5EE** To format information about WebSphere MQ internals

You can also specify the JOBNAME(*jobname*) parameter to limit the formatted output to specific jobs.

```

//S1 EXEC PGM=IKJEFT01,DYNAMNBR=20,REGION=4096K
//IPCSPARM DD DSN=SYS1.PARMLIB,DISP=SHR
//IPCSDDIR DD DSN=thlqual.ipcs.dataset.directory,DISP=SHR
//SYSTSPRT DD SYSOUT=*,DCB=(LRECL=137)
//IPCSTOC DD SYSOUT=*
//GTFIN DD DSN=gtf.trace,DISP=SHR
//SYSTSIN DD *
IPCS
SETDEF FILE(GTFIN) NOCONFIRM
GTFTRACE USR(5E9,5EA,5EE)
/*
//STEPLIB DD DSN=thlqual.SCSQAUTH,DISP=SHR

```

Figure 152. Formatting the GTF output in batch. *thlqual* is your high level qualifier for WebSphere MQ data sets, and *gtf.trace* is the name of the data set containing your trace information. You must also specify your IPCS data set directory.

Identifying the control blocks associated with WebSphere MQ

The format identifier for the WebSphere MQ trace is D9. This value appears at the beginning of each formatted control block in the formatted GTF output, in the form:

```
USRD9
```

Identifying the event identifier associated with the control block

The trace formatter inserts one of the following messages at the top of each control block. These indicate whether the data was captured on entry to or exit from WebSphere MQ

- CSQW072I ENTRY: MQ user parameter trace
- CSQW073I EXIT: MQ user parameter trace

Related concepts:

“Controlling the GTF” on page 1245

Interpreting the trace information:

The GTFTRACE produced by WebSphere MQ can be examined to determine possible errors with invalid addresses, invalid control blocks, and invalid data.

When you look at the data produced by the GTFTRACE command, consider the following points:

- If the control block consists completely of zeros, it is possible that an error occurred while copying data from the user's address space. This might be because an invalid address was passed.
- If the first part of the control block contains non-null data, but the rest consists of zeros, it is again possible that an error occurred while copying data from the user's address space, for example, the control block was not placed entirely within valid storage. This might also be due to the control block not being initialized correctly.
- If the error has occurred on exit from WebSphere MQ, it is possible that WebSphere MQ might not write the data to the user's address space. The data displayed is the version that it was attempting to copy to the user's address space.

The following tables show details of the control blocks that are traced.

Table 136 on page 1249 illustrates which control blocks are traced for different MQI calls.

Table 136. Control blocks traced for WebSphere MQ MQI calls

MQI call	Entry	Exit
MQCB	MQCBD, MQMD, MQGMO	MQCBD, MQMD, MQGMO
MQCLOSE	None	None
MQGET	MQMD, MQGMO	MQMD, MQGMO, and the first 256 bytes of message data
MQINQ	Selectors (if <i>SelectorCount</i> is greater than 0)	Selectors (if <i>SelectorCount</i> is greater than 0) Integer attributes (if <i>IntAttrCount</i> is greater than 0) Character attributes (if <i>CharAttrLength</i> is greater than 0)
MQOPEN	MQOD	MQOD
MQPUT	MQMD, MQPMO, and the first 256 bytes of message data	MQMD, MQPMO, and the first 256 bytes of message data
MQPUT1	MQMD, MQOD, MQPMO, and the first 256 bytes of message data	MQMD, MQOD, MQPMO, and the first 256 bytes of message data
MQSET	Selectors (if <i>SelectorCount</i> is greater than 0) Integer attributes (if <i>IntAttrCount</i> is greater than 0) Character attributes (if <i>CharAttrLength</i> is greater than 0)	Selectors (if <i>SelectorCount</i> is greater than 0) Integer attributes (if <i>IntAttrCount</i> is greater than 0) Character attributes (if <i>CharAttrLength</i> is greater than 0)
MQSTAT	MQSTS	MQSTS
MQSUB	MQSD, MQSD.ObjectString, MQSD.SubName, MQSD.SubUserData, MQSD.SelectionString, MQSD.ResObjectString	MQSD, MQSD.ObjectString, MQSD.SubName, MQSD.SubUserData, MQSD.SelectionString, MQSD.ResObjectString
MQSUBRQ	MQSRO	MQSRO

Note: In the special case of an **MQGET** call with the WAIT option, a double entry is seen if there is no message available at the time of the **MQGET** request, but a message subsequently becomes available before the expiry of any time interval specified.

This is because, although the application has issued a single **MQGET** call, the adapter is performing the wait on behalf of the application and when a message becomes available it reissues the call. So in the trace it will appear as a second **MQGET** call.

Information about specific fields of the queue request parameter list is also produced in some circumstances. The fields in this list are identified as follows:

Identifier	Description
BufferL	Buffer length
CompCode	Completion code
CharAttL	Character attributes length
DataL	Data length
Hobj	Object handle
IntAttC	Count of integer attributes
pObjDesc	Object descriptor

Identifier	Description
Options	Options
pBuffer	Address of buffer
pCharAtt	Address of character attributes
pECB	Address of ECB used in get
pGMO	Address of get message options
pIntAtt	Address of integer attributes
pMsgDesc	Address of message descriptor
pPMO	Address of put message options
pSelect	Address of selectors
Reason	Reason code
RSVn	Reserved for IBM
SelectC	Selector count
Thread	Thread
UOWInfo	Information about the unit of work
Userid	CICS or IMS user ID, for batch or TSO this is zero

The Log and Trace Analyzer tool

There is a Log and Trace analyzer tool available for download, and a guide for users.

The Log and Trace Analyzer tool is an interface that allows you to work with logs and traces produced by different components of a deployed system. It provides a single point of contact for importing and analyzing log files or trace files from multiple products and allows you to determine the relationship between the events captured by these products (correlation).

You can find online help for the Log and Trace Analyzer tool in the Autonomic Computing Toolkit User's Guide here: <http://download.boulder.ibm.com/ibmdl/pub/software/dw/autonomic/books/fpu3mst.htm>

Other types of trace

There are other trace facilities available for problem determination. Use this topic to investigate channel initiator trace, line trace, CICS adapter trace, SSL trace, and z/OS trace.

It can be helpful to use the following trace facilities with WebSphere MQ.

- The channel initiator trace
- The line trace
- The CICS adapter trace
- System SSL trace
- z/OS traces

The channel initiator trace

See Figure 157 on page 1278 for information about how to get a dump of the channel initiator address space. Note that dumps produced by the channel initiator do not include trace data space. The trace data space, which is called CSQXTRDS, contains trace information. You can request this by specifying it on a slip trap or when you use the dump command.

To run the trace using the START TRACE command, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) manual. You can set this trace to start automatically using the TRAXSTR queue manager attribute. For more information about how to do this, see the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*).

You can display this trace information by entering the IPCS command:

```
LIST 1000. DSPNAME(CSQXTRDS)
```

You can format it using the command:

```
CTRACE COMP(CSQX $ssnm$ )
```

where $ssnm$ is the subsystem name.

The line trace

A wrap-around line trace exists for each channel. This trace is kept in a 4 KB buffer for each channel in the channel initiator address space. Trace is produced for each channel, so it is ideal for problems where a channel appears to be hung, because information can be collected about the activity of this channel long after the normal trace has wrapped.

The line trace is always active; you cannot turn it off. It is available for both LU 6.2 and TCP channels and should reduce the number of times a communications trace is required.

You can view the trace as unformatted trace that is written to CSQSNAP. You can display the trace by following these steps:

1. Ensure that the CHIN procedure has a SNAP DD statement.
2. Start a CHIN trace, specifying IFCID 202 as follows:

```
START TRACE(CHINIT) CLASS(4) IFCID(202)
```
3. Display the channel status for those channels for which the line trace is required:

```
DISPLAY CHSTATUS(channel) SAVED
```

This dumps the current line for the selected channels to CSQSNAP. See “Snap dumps” on page 1291 for further information.

Note:

- a. The addresses of the storage dump are incorrect because the CSQXFFST mechanism takes a copy of the storage before writing it to CSQSNAP.
- b. The dump to CSQSNAP is only produced the first time you run the DISPLAY CHSTATUS SAVED command. This is to prevent getting dumps each time you run the command.

To obtain another dump of line trace data, you must stop and restart the current trace.

- 1) You can use a selective STOP TRACE command to stop just the trace that was started to gather the line trace data. To do this, note the TRACE NUMBER assigned to the trace as shown in this example:

```
+ssid START TRACE(CHINIT) CLASS(4) IFCID(202)  
CSQW130I +ssid 'CHINIT' TRACE STARTED, ASSIGNED TRACE NUMBER 01
```

- 2) To stop the trace, issue the following command:

```
+ssid STOP TRACE(CHINIT) TNO(01)
```

- 3) You can then enter another START TRACE command with a DISPLAY CHSTATUS SAVED command to gather more line trace data to CSQSNAP.

- The line trace buffer is unformatted. Each entry starts with a clock, followed by a time stamp, and an indication of whether this is an OUTBOUND or INBOUND flow. Use the time stamp information to find the earliest entry.

The CICS adapter trace

The CICS adapter writes entries to the CICS trace if your trace number is set to a value in the range 0 through 199 (decimal), and if either:

- CICS user tracing is enabled, or
- CICS internal/auxiliary trace is enabled

You can enable CICS tracing in one of two ways:

- Dynamically, using the CICS-supplied transaction CETR
- By ensuring that the USERTR parameter in the CICS system initialization table (SIT) is set to YES

For more information about enabling CICS trace, see the *CICS Problem Determination Guide*.

The CICS trace entry originating from the CICS adapter has a value AP0000, where 000 is the hexadecimal equivalent of the decimal value of the CICS adapter trace number you specified.

The trace entries are shown in “CICS adapter trace entries.”

System SSL trace

You can collect System SSL trace using the SSL Started Task. The details of how to set up this task are in the *System Secure Sockets Layer Programming* documentation, SC24-5901. A trace file is generated for each SSLTASK running in the CHINIT address space.

z/OS traces

z/OS traces, which are common to all products operating as formal subsystems of z/OS, are available for use with WebSphere MQ. For information about using and interpreting this trace facility, see the *MVS Diagnosis: Tools and Service Aids* manual.

CICS adapter trace entries:

Use this topic as a reference for CICS adapter trace entries.

The CICS trace entry for these values is AP0xxx (where xxx is the hexadecimal equivalent of the trace number you specified when the CICS adapter was enabled). These trace entries are all issued by CSQCTRUE, except CSQCTEST, which is issued by CSQCRST and CSQCDSP.

Table 137. CICS adapter trace entries

Name	Description	Trace sequence	Trace data
CSQCABNT	Abnormal termination	Before issuing END_THREAD ABNORMAL to WebSphere MQ. This is due to the end of the task and therefore an implicit backout could be performed by the application. A ROLLBACK request is included in the END_THREAD call in this case.	Unit of work information. You can use this information when finding out about the status of work. (For example, it can be verified against the output produced by the DISPLAY_THREAD command, or the log print utility.)
CSQCAUID	Bridge security	Before validating bridge user password or PassTicket.	User ID.

Table 137. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCBACK	Syncpoint backout	Before issuing BACKOUT to WebSphere MQ. This is due to an explicit backout request from the application.	Unit of work information.
CSQCCONX	MQCONN	Before issuing MQCONN to WebSphere MQ.	Connection tag.
CSQCCRC	Completion code and reason code	After unsuccessful return from API call.	Completion code and reason code.
CSQCCOMM	Syncpoint commit	Before issuing COMMIT to WebSphere MQ. This can be due to a single-phase commit request or the second phase of a two-phase commit request. The request is due to an explicit syncpoint request from the application.	Unit of work information.
CSQCDCFF	IBM use only		
CSQCDCIN	IBM use only		
CSQCDCOT	IBM use only		
CSQCEXER	Execute resolve	Before issuing EXECUTE_RESOLVE to WebSphere MQ.	The unit of work information of the unit of work issuing the EXECUTE_RESOLVE. This is the last in-doubt unit of work in the resynchronization process.
CSQCGETW	GET wait	Before issuing CICS wait.	Address of the ECB to be waited on.
CSQCGMGD	GET message data	After successful return from MQGET.	Up to 40 bytes of the message data.
CSQCGMGH	GET message handle	Before issuing MQGET to WebSphere MQ.	Object handle.
CSQCGMGI	Get message ID	After successful return from MQGET.	Message ID and correlation ID of the message.
CSQCHCER	Hconn error	Before issuing any MQ verb.	Connection handle.
CSQCINDL	In-doubt list	After successful return from the second INQUIRE_INDOUBT.	The in-doubt units of work list.
CSQCINDO	IBM use only		
CSQCINDS	In-doubt list size	After successful return from the first INQUIRE_INDOUBT and the in-doubt list is not empty.	Length of the list; divided by 64 gives the number of in-doubt units of work.
CSQCINDW	Syncpoint in doubt	During syncpoint processing, CICS is in doubt as to the disposition of the unit of work.	Unit of work information.
CSQCINQH	INQ handle	Before issuing MQINQ to WebSphere MQ.	Object handle.
CSQCLOSH	CLOSE handle	Before issuing MQCLOSE to WebSphere MQ.	Object handle.

Table 137. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCLOST	Disposition lost	During the resynchronization process, CICS informs the adapter that it has been cold started so no disposition information regarding the unit of work being resynchronized is available.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCNIND	Disposition not in doubt	During the resynchronization process, CICS informs the adapter that the unit of work being resynchronized should not have been in doubt (that is, perhaps it is still running).	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCNORT	Normal termination	Before issuing END_THREAD NORMAL to WebSphere MQ. This is due to the end of the task and therefore an implicit syncpoint commit might be performed by the application. A COMMIT request is included in the END_THREAD call in this case.	Unit of work information.
CSQCOPNH	OPEN handle	After successful return from MQOPEN .	Object handle.
CSQCOPNO	OPEN object	Before issuing MQOPEN to WebSphere MQ.	Object name.
CSQCPMGD	PUT message data	Before issuing MQPUT to WebSphere MQ.	Up to 40 bytes of the message data.
CSQCPMGH	PUT message handle	Before issuing MQPUT to WebSphere MQ.	Object handle.
CSQCPMGI	PUT message ID	After successful MQPUT from WebSphere MQ.	Message ID and correlation ID of the message.
CSQCPREP	Syncpoint prepare	Before issuing PREPARE to WebSphere MQ in the first phase of two-phase commit processing. This call can also be issued from the distributed queuing component as an API call.	Unit of work information.
CSQCP1MD	PUTONE message data	Before issuing MQPUT1 to WebSphere MQ.	Up to 40 bytes of data of the message.
CSQCP1MI	PUTONE message ID	After successful return from MQPUT1 .	Message ID and correlation ID of the message.
CSQCP1ON	PUTONE object name	Before issuing MQPUT1 to WebSphere MQ.	Object name.
CSQCRBAK	Resolved backout	Before issuing RESOLVE_ROLLBACK to WebSphere MQ.	Unit of work information.
CSQCRCMT	Resolved commit	Before issuing RESOLVE_COMMIT to WebSphere MQ.	Unit of work information.

Table 137. CICS adapter trace entries (continued)

Name	Description	Trace sequence	Trace data
CSQCRMIR	RMI response	Before returning to the CICS RMI (resource manager interface) from a specific invocation.	Architected RMI response value. Its meaning depends of the type of the invocation. To determine the type of invocation, look at previous trace entries produced by the CICS RMI component.
CSQCRSYN	Resync	Before the resynchronization process starts for the task.	Unit of work ID known to CICS for the unit of work being resynchronized.
CSQCSETH	SET handle	Before issuing MQSET to WebSphere MQ.	Object handle.
CSQCTASE	IBM use only		
CSQCTEST	Trace test	Used in EXEC CICS ENTER TRACE call to verify the trace number supplied by the user or the trace status of the connection.	No data.

Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions

How to request iKeyman and iKeycmd tracing.

To request iKeyman tracing, execute the iKeyman command for your platform with the following **-D** flags.

For Windows UNIX and Linux systems:

```
strmqikm -Dkeyman.debug=true -Dkeyman.jnitracng=0N
```

To request iKeycmd tracing, run the iKeycmd command for your platform with the following **-D** flags.

For Windows UNIX and Linux systems:

```
runmqckm -Dkeyman.debug=true -Dkeyman.jnitracng=0N
```

iKeyman and iKeycmd write three trace files to the directory from which you start them, so consider starting iKeyman or iKeycmd from the trace directory to which the runtime SSL trace is written: `/var/mqm/trace` on UNIX and Linux systems and `MQ_INSTALLATION_PATH/trace` on Windows. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed. The trace files that iKeyman and iKeycmd generate are:

ikmgdbg.log

Java related trace

ikmjdbg.log

JNI related trace

ikmcdbg.log

C related trace

These trace files are binary, so they must be transferred in binary transfer mode when they are transferred from system to system using FTP. The trace files are typically approximately 1 MB each.

On UNIX, Linux, and Windows systems, you can independently request trace information for iKeyman, iKeycmd, the runtime SSL functions, or a combination of these.

The runtime SSL trace files have the names AMQ.SSL.TRC and AMQ.SSL.TRC.1. You cannot format any of the SSL trace files; send them unchanged to IBM support. The SSL trace files are binary files and, if they are transferred to IBM support via FTP, they must be transferred in binary transfer mode.

Related concepts:

“Using trace on UNIX and Linux systems” on page 1236

“Using trace on WebSphere MQ server on IBM i” on page 1240

“Using trace for problem determination on z/OS” on page 1244

“Java diagnostics” on page 1262

Related reference:

“Using trace on Windows” on page 1234

Tracing WebSphere MQ classes for JMS programs

The trace facility in IBM WebSphere MQ classes for JMS is provided to help IBM staff to diagnose customer problems. Various properties control its behavior.

Trace is turned off by default, because the output rapidly becomes large, and is unlikely to be of use in normal circumstances.

Except where otherwise stated, all the properties are set in the WebSphere MQ classes for JMS configuration file. For information about this file, see  The WebSphere MQ classes for JMS configuration file.

If you are asked to provide trace output, turn tracing on by setting the property `com.ibm.msg.client.commonservices.trace.status` to ON. To turn tracing off, set the property `com.ibm.msg.client.commonservices.trace.status` to OFF.

You can also turn tracing on and off using the IBM WebSphere MQ commands **strmqtrc** and **endmqtrc** on all platforms except z/OS or IBM i. To use the commands, the JAR file `com.ibm.mq.commonservices.jar` must be present on the class path. The class is not prerequisite for the IBM WebSphere MQ classes for JMS, and so might not be present, for example on an application server. Enable the **strmqtrc** and **endmqtrc** commands by starting the JVM that is running the IBM WebSphere MQ classes for JMS with the system property:

```
-Dcom.ibm.msg.client.commonservices.trace.status=wmqtrc
```

Configure the trace output using the following properties:

com.ibm.msg.client.commonservices.trace.outputName=*traceOutputName*

traceOutputName is the directory and file name to which trace output is sent.

traceOutputName defaults to a file named `mqjms_PID.trc` in the current working directory where *PID* is the current process ID. If a process ID is unavailable, a random number is generated and prefixed with the letter *f*. To include the process ID in a file name you specify, use the string `%PID%`.

If you specify an alternative directory, it must exist, and you must have write permission for this directory. If you do not have write permission, the trace output is written to `System.err`.

com.ibm.msg.client.commonservices.trace.include=*includeList*

includeList is a list of packages and classes that are traced, or the special values ALL or NONE.

Separate package or class names with a semicolon, `;`. *includeList* defaults to ALL, and traces all packages and classes in IBM WebSphere MQ classes for JMS.

Note: You can include a package but then exclude subpackages of that package. For example, if you include package `a.b` and exclude package `a.b.x`, the trace includes everything in `a.b.y` and `a.b.z`, but not `a.b.x` or `a.b.x.1`.

com.ibm.msg.client.commonservices.trace.exclude=excludeList

excludeList is a list of packages and classes that are not traced, or the special values ALL or NONE.

Separate package or class names with a semicolon, ;. *excludeList* defaults to NONE, and therefore excludes no packages and classes in IBM WebSphere MQ classes for JMS from being traced.

Note: You can exclude a package but then include subpackages of that package. For example, if you exclude package a.band include package a.b.x, the trace includes everything in a.b.x and a.b.x.1, but not a.b.y or a.b.z.

Any package or class that is specified, at the same level, as both included and excluded is included.

com.ibm.msg.client.commonservices.trace.maxBytes=maxArrayBytes

maxArrayBytes is the maximum number of bytes that are traced from any byte arrays.

If *maxArrayBytes* is set to a positive integer, it limits the number of bytes in a byte-array that are written out to the trace file. It truncates the byte array after writing *maxArrayBytes* out. Setting *maxArrayBytes* reduces the size of the resulting trace file, and reduces the effect of tracing on the performance of the application.

A value of 0 for this property means that none of the contents of any byte arrays are sent to the trace file.

The default value is -1, which removes any limit on the number of bytes in a byte array that are sent to the trace file.

com.ibm.msg.client.commonservices.trace.limit=maxTraceBytes

maxTraceBytes is the maximum number of bytes that are written to a trace output file.

maxTraceBytes works with *traceCycles*. If the number of bytes of trace written is near to the limit, the file is closed, and a new trace output file is started.

A value of 0 means that a trace output file has zero length. The default value is -1, which means that the amount of data to be written to a trace output file is unlimited.

com.ibm.msg.client.commonservices.trace.count=traceCycles

traceCycles is the number of trace output files to cycle through.

If the current trace output file reaches the limit specified by *maxTraceBytes*, the file is closed. Further trace output is written to the next trace output file in sequence. Each trace output file is distinguished by a numeric suffix appended to the file name. The current or most recent trace output file is mqjms.trc.0, the next most recent trace output file is mqjms.trc.1. Older trace files follow the same numbering pattern up to the limit.

The default value of *traceCycles* is 1. If *traceCycles* is 1, when the current trace output file reaches its maximum size, the file is closed and deleted. A new trace output file with the same name is started. Therefore, only one trace output file exists at a time.

com.ibm.msg.client.commonservices.trace.parameter=traceParameters

traceParameters controls whether method parameters and return values are included in the trace.

traceParameters defaults to TRUE. If *traceParameters* is set to FALSE, only method signatures are traced.

com.ibm.msg.client.commonservices.trace.startup=startup

There is an initialization phase of IBM WebSphere MQ classes for JMS during which resources are allocated. The main trace facility is initialized during the resource allocation phase.

If *startup* is set to TRUE, startup trace is used. Trace information is produced immediately and includes the setup of all components, including the trace facility itself. Startup trace information can be used to diagnose configuration problems. Startup trace information is always written to System.err.

startup defaults to FALSE.

startup is checked before initialization is complete. For this reason, only specify the property on the command line as a Java system property. Do not specify it in the IBM WebSphere MQ classes for JMS configuration file.

com.ibm.msg.client.commonservices.trace.compress=*compressedTrace*

Set *compressedTrace* to TRUE to compress trace output.

The default value of *compressedTrace* is FALSE.

If *compressedTrace* is set to TRUE, trace output is compressed. The default trace output file name has the extension *.trz*. If compression is set to FALSE, the default value, the file has the extension *.trc* to indicate it is uncompressed. However if the file name for the trace output has been specified in *traceOutputName* that name is used instead; no suffix is applied to the file.

Compressed trace output is smaller than uncompressed. Because there is less I/O, it can be written out faster than uncompressed trace. Compressed tracing has less effect on the performance of IBM WebSphere MQ classes for JMS than uncompressed tracing.

If *maxTraceBytes* and *traceCycles* are set, multiple compressed trace files are created in place of multiple flat files.

If IBM WebSphere MQ classes for JMS ends in an uncontrolled manner, a compressed trace file might not be valid. For this reason, trace compression must only be used when IBM WebSphere MQ classes for JMS closes down in a controlled manner. Only use trace compression if the problems being investigated do not cause the JVM itself to stop unexpectedly. Do not use trace compression when diagnosing problems that can result in System.Halt() shutdowns or abnormal, uncontrolled JVM terminations.

com.ibm.msg.client.commonservices.trace.level=*traceLevel*

traceLevel specifies a filtering level for the trace. The defined trace levels are as follows:

TRACE_NONE	0
TRACE_EXCEPTION	1
TRACE_WARNING	3
TRACE_INFO	6
TRACE_ENTRYEXIT	8
TRACE_DATA	9
TRACE_ALL	Integer.MAX_VALUE

Each trace level includes all lower levels. For example, if trace level is set at TRACE_INFO, then any trace point with a defined level of TRACE_EXCEPTION, TRACE_WARNING, or TRACE_INFO is written to the trace. All other trace points are excluded.

com.ibm.msg.client.commonservices.trace.standalone=*standaloneTrace*

standaloneTrace controls whether the IBM WebSphere MQ JMS client tracing service is used in a WebSphere Application Server environment.

If *standaloneTrace* is set to TRUE, the IBM WebSphere MQ JMS client tracing properties are used to determine the trace configuration.

If *standaloneTrace* is set to FALSE, and the IBM WebSphere MQ JMS client is running in an WebSphere Application Server container, the WebSphere Application Server trace service is used. The trace information that is generated depends upon the trace settings of the application server.

The default value of *standaloneTrace* is FALSE.

To dynamically enable or disable trace from within an application, or to change the trace level, use the methods of the `com.ibm.msg.client.services.Trace` class.

setOn()

Turns on the trace facility.

setOff()

Turns off the trace facility.

setStatus(boolean traceOn)

Turns the trace facility on or off, depending on the value of *traceOn*.

isOn() Checks whether the trace facility is on.

setTraceLevel(int newTraceLevel)

Sets the tracing detail level.

getTraceLevel()

Returns the tracing detail level.

If a severe or unrecoverable error occurs, First Failure Support Technology™ (FFST™) information is recorded in a file with a name of the format JMSCCxxx.FDC. xxxx is a four-digit number. It is incremented to differentiate .FDC files.

.FDC files are always written to a subdirectory called FFDC. The subdirectory is in one of two locations, depending on whether trace is active:

Trace is active, and *traceOutputName* is set

The FFDC directory is created as a subdirectory of the directory to which the trace file is being written.

Trace is not active or *traceOutputName* is not set

The FFDC directory is created as a subdirectory of the current working directory.

For more information about FFST in WebSphere MQ classes for JMS, see  First failure support technology (FFST) in WebSphere MQ classes for JMS (*WebSphere MQ V7.1 Programming Guide*).

The JSE common services uses `java.util.logging` as its trace and logging infrastructure. The root object of this infrastructure is the `LogManager`. The log manager has a `reset` method, which closes all handlers and sets the log level to `null` - in effect turning off all the trace. If your application or application server calls `java.util.logging.LogManager.getLogManager().reset()`, it closes all trace, which might prevent you from diagnosing any problems. To avoid closing all trace, create a `LogManager` class with an overridden `reset()` method that does nothing, as in the following example.

```
package com.ibm.javaut.tests;
import java.util.logging.LogManager;
public class JmsLogManager extends LogManager {
    // final shutdown hook to ensure that the trace is finally shutdown
    // and that the lock file is cleaned-up
    public class ShutdownHook extends Thread{
        public void run(){
            doReset();
        }
    }
    public JmsLogManager(){
        // add shutdown hook to ensure final cleanup
        Runtime.getRuntime().addShutdownHook(new ShutdownHook());
    }
    public void reset() throws SecurityException {
        // does nothing
    }
}
```

```

public void doReset(){
    super.reset();
}
}

```

The shutdown hook is necessary to ensure that trace is properly shutdown when the JVM finishes. To use the modified log manager instead of the default one, add a system property to the JVM startup:

```
java -Djava.util.logging.manager=com.mycompany.logging.LogManager ...
```

Note: When trace is activated, trace creates a file named `mqjms.trc.lck`. If you use a version of Java earlier than Java 5, `mqjms.trc.lck` is not removed when trace ends, due to a defect in the Java class libraries. Delete `mqjms.trc.lck` manually after the trace file has been closed.

Tracing using MQJMS_TRACE_LEVEL

To maintain backwards compatibility, the trace parameters used by Version 6.0 of IBM WebSphere MQ classes for JMS are still supported. `MQJMS_TRACE_LEVEL` is deprecated for any new application.

In version 6.0, the Java property `MQJMS_TRACE_LEVEL` turned on JMS trace. It has three values:

- on** Traces IBM WebSphere MQ classes for JMS calls only.
- base** Traces both IBM WebSphere MQ classes for JMS calls and the underlying IBM WebSphere MQ classes for Java calls.
- off** Disables tracing.

Setting `MQJMS_TRACE_LEVEL` to `on` or `base` produces the same results as setting the `com.ibm.msg.client.commonservices.trace.status` property to `on`.

Setting the property, `MQJMS_TRACE_DIR` to `somepath/tracedir` is equivalent to setting the `com.ibm.msg.client.commonservices.trace.outputName` property to `somepath/tracedir/mqjms_%PID%.trc`.

Tracing WebSphere MQ classes for Java programs

The WebSphere MQ classes for Java include a trace facility, which you can use to produce diagnostic messages if you suspect that there might be a problem with the code.

There are two ways to collect trace:

- By using the WebSphere MQ common services trace facility
- By using the WebSphere MQ messaging client trace mechanism

The preferred method of collecting trace is by using the WebSphere MQ messaging client trace mechanism. This trace mechanism is also used by the WebSphere MQ classes for JMS, and generates a more detailed trace than the trace that is generated by the WebSphere MQ common services trace facility.

Using the WebSphere MQ common services trace facility

You can collect a WebSphere MQ classes for Java trace by using the WebSphere MQ common services trace facility.

Before collecting a trace by using the WebSphere MQ common services trace facility, you need to configure the system to collect WebSphere MQ Java diagnostics, as described in “Java diagnostics” on page 1262, and create a WebSphere MQ common services properties file as described in “Using `com.ibm.mq.commonservices`” on page 1263.

When you have created a WebSphere MQ common services properties file for your application to use, start your application by using a **java** command with the following format:

```
java -Dcom.ibm.mq.commonservices=<common services properties file> MyApp
```

The following example shows a **java** command:

```
java -Dcom.ibm.mq.commonservices=C:\mydir\mqcommonservices.properties MyApp
```

When an application starts, the WebSphere MQ classes for Java read the contents of the properties file and store the specified properties in an internal property store. If the **java** command does not identify a properties file, or if the properties file cannot be found, the WebSphere MQ classes for Java use the default values for all the properties. If required, you can override any property in the configuration file by specifying it as a system property on the **java** command.

Trace is written to a file called AMQ*.trc in the directory that is specified by the `Diagnostics.Java.Trace.Destination.Pathname` entry in the WebSphere MQ common services properties file. If this entry has not been specified, then the trace file is written to the current working directory for the application.

If an application has been started by using a **java** command that includes the **-Dcom.ibm.mq.commonservices** property, then it is possible to control WebSphere MQ classes for Java trace within your application by using the methods:

- `MQEnvironment.enableTrace(int level)`
- `MQEnvironment.enableTrace(int level, OutputStream stream)`
- `MQEnvironment.disableTrace()`

The following example shows how to do this:

```
MQEnvironment.enableTracing(int); // start trace
...                               // these commands will be traced
MQEnvironment.disableTracing();  // turn tracing off again
```

Important: When trace is enabled in this way, the parameters that are passed into the `MQEnvironment.enableTrace(int level)` and `MQEnvironment.enableTrace(int level, OutputStream stream)` methods are ignored. Trace is always:

- Written to a file called AMQ*.trc in either the directory that is specified by `Diagnostics.Java.Trace.Destination.Pathname` entry in the common services properties file, or the current working directory.
- Collected at the trace level that is specified by the `Diagnostics.Java.Trace.Detail` entry in the common services properties file.

Related concepts:

“Using the WebSphere MQ messaging client trace mechanism”

Using the WebSphere MQ messaging client trace mechanism

It is possible to use the WebSphere MQ messaging client trace mechanism to collect a WebSphere MQ classes for Java trace. This trace mechanism, which is primarily used to collect diagnostic information from the WebSphere MQ classes for JMS, generates a detailed trace that contains more information than traces collected by using the WebSphere MQ common services trace facility, and is the preferred way to collect trace.

To collect a trace of a WebSphere MQ classes for Java application by using the WebSphere MQ messaging client trace mechanism, start your application by using a **java** command with the following format:

```
java
  -Dcom.ibm.mq.integrateJMSTrace=true
  -Dcom.ibm.msg.client.commonservices.trace.status=ON
  -Dcom.ibm.msg.client.commonservices.trace.outputName=<trace file name> MyApp
```

The following example shows a **java** command:

```
java
-Dcom.ibm.mq.integrateJMSTrace=true
-Dcom.ibm.msg.client.commonservices.trace.status=ON
-Dcom.ibm.msg.client.commonservices.trace.outputName=C:\mqjava_trc.txt MyApp
```

When the WebSphere MQ classes for Java start, they start writing detailed trace to the file specified by the Java system property **-Dcom.ibm.msg.client.commonservices.trace.outputName**. Note that if this property is omitted, trace is written to a file called `mqjms_<PID>.trc`, in the current working directory for the application, where `<PID>` is the process identifier of the WebSphere MQ classes for Java application. The reason that the file starts with the `mqjms` prefix is because the Java system property **-Dcom.ibm.mq.integrateJMSTrace=true** causes the trace information that is generated by the WebSphere MQ classes for Java to be routed through the WebSphere MQ classes for JMS trace mechanism.

Related concepts:

“Using the WebSphere MQ common services trace facility” on page 1260

Java diagnostics

For Java components of WebSphere MQ, for example the WebSphere MQ Explorer and the Java implementation of WebSphere MQ Transport for SOAP, diagnostic information is output using the standard WebSphere MQ diagnostic facilities or by Java diagnostic classes.

Diagnostic information in this context consists of trace, first-failure data capture (FFDC) and error messages.

You can choose to have this information produced using WebSphere MQ facilities or the facilities of WebSphere MQ classes for Java or WebSphere MQ classes for JMS, as appropriate. Generally use the WebSphere MQ diagnostic facilities if they are available on the local system.

You might want to use the Java diagnostics in the following circumstances:

- On a system on which queue managers are available, if the queue manager is managed separately from the software you are running.
- To reduce performance effect of WebSphere MQ trace.

To request and configure diagnostic output, two system properties are used when starting a WebSphere MQ Java process:

- System property `com.ibm.mq.commonservices` specifies a standard Java property file, which contains a number of lines which are used to configure the diagnostic outputs. Each line of code in the file is free-format, and is terminated by a new line character.
- System property `com.ibm.mq.commonservices.diagid` associates trace and FFDC files with the process which created them.

For information about using the `com.ibm.mq.commonservices` properties file to configure diagnostics information, see “Using `com.ibm.mq.commonservices`” on page 1263.

For instructions on locating trace information and FFDC files, see “Java trace and FFDC files” on page 1264.

Related concepts:

“Using trace on UNIX and Linux systems” on page 1236

“Using trace on WebSphere MQ server on IBM i” on page 1240

“Using trace for problem determination on z/OS” on page 1244

“Tracing Secure Sockets Layer (SSL) iKeyman and iKeycmd functions” on page 1255

Related reference:

“Using trace on Windows” on page 1234

Using com.ibm.mq.commonservices

The com.ibm.mq.commonservices properties file contains the following entries relating to the output of diagnostics from the Java components of WebSphere MQ.

Note that case is significant in all these entries:

Diagnostics.MQ=enabled | disabled

Are WebSphere MQ diagnostics to be used? If Diagnostics.MQ is enabled, diagnostic output is as for other WebSphere MQ components; trace output is controlled by the parameters in the **strmqtrc** and **endmqtrc** control commands, or the equivalent. The default is *enabled*.

Diagnostics.Java=options

Which components are traced using Java trace. Options are one or more of *explorer*, *soap*, and *wmqjavaclasses*, separated by commas, where "explorer" refers to the diagnostics from the WebSphere MQ Explorer, "soap" refers to the diagnostics from the running process within WebSphere MQ Transport for SOAP, and "wmqjavaclasses" refers to the diagnostics from the underlying WebSphere MQ Java classes. By default no components are traced.

Diagnostics.Java.Trace.Detail=high | medium | low

Detail level for Java trace. The *high* and *medium* detail levels match those used in WebSphere MQ tracing but *low* is unique to Java trace. This property is ignored if Diagnostics.Java is not set. The default is *medium*.

Diagnostics.Java.Trace.Destination.File=enabled | disabled

Whether Java trace is written to a file. This property is ignored if Diagnostics.Java is not set. The default is *disabled*.

Diagnostics.Java.Trace.Destination.Console=enabled | disabled

Whether Java trace is written to the system console. This property is ignored if Diagnostics.Java is not set. The default is *disabled*.

Diagnostics.Java.Trace.Destination.Pathname=dirname

The directory to which Java trace is written. This property is ignored if Diagnostics.Java is not set or Diagnostics.Java.Trace.Destination.File=disabled. On UNIX and Linux systems, the default is */var/mqm/trace* if it is present, otherwise the Java console (System.err). On Windows, the default is the system console.

Diagnostics.Java.FFDC.Destination.Pathname=dirname

The directory to which Java FFDC output is written. The default is the current working directory.

Diagnostics.Java.Errors.Destination.Filename=filename

The fully qualified file name to which Java error messages are written. The default is *AMQJAVA.LOG* in the current working directory.

An example of a com.ibm.mq.commonservices properties file is given in Figure 153 on page 1264. Lines beginning with the number sign (#) are treated as comments.

```

#
# Base WebSphere MQ diagnostics are disabled
#
Diagnostics.MQ=disabled
#
# Java diagnostics for WebSphere MQ Transport for SOAP
# and the WebSphere MQ Java Classes are both enabled
#
Diagnostics.Java=soap,wmqjavaclasses
#
# High detail Java trace
#
Diagnostics.Java.Trace.Detail=high
#
# Java trace is written to a file and not to the console.
#
Diagnostics.Java.Trace.Destination.File=enabled
Diagnostics.Java.Trace.Destination.Console=disabled
#
# Directory for Java trace file
#
Diagnostics.Java.Trace.Destination.Pathname=c:\\tracedir
#
# Directory for First Failure Data Capture
#
Diagnostics.Java.FFDC.Destination.Pathname=c:\\ffdcdir
#
# Directory for error logging
#
Diagnostics.Java.Errors.Destination.FileName=c:\\errorsdir\\SOAPERRORS.LOG
#

```

Figure 153. Sample `com.ibm.mq.commonservices.properties` file

A sample properties file, `WMQSoap_RAS.properties`, is also supplied as part of the "Java messaging and SOAP transport" install option.

Java trace and FFDC files

File name conventions for Java trace and FFDC files.

When Java trace is generated for the IBM WebSphere MQ Explorer or for IBM WebSphere MQ Transport for SOAP it is written to a file with a name of the format `AMQ.diagid.counter.TRC`. Here, *diagid* is the value of the system property `com.ibm.mq.commonservices.diagid` associated with this Java process, as described earlier in this section, and *counter* is an integer greater than or equal to 0. All letters in the name are in uppercase, matching the naming convention used for normal IBM WebSphere MQ trace.

If `com.ibm.mq.commonservices.diagid` is not specified, the value of *diagid* is the current time, in the format `YYYYMMDDhhmmssmm`.

The IBM WebSphere MQ Java classes trace file has a name based on the equivalent IBM WebSphere MQ Explorer or SOAP Java trace file. The name differs in that it has the string `.JC` added before the `.TRC` string, giving a format of `AMQ.diagid.counter.JC.TRC`.

When Java FFDC is generated for the IBM WebSphere MQ Explorer or for IBM WebSphere MQ Transport for SOAP it is written to a file with a name of the format `AMQ.diagid.counter.FDC` where *diagid* and *counter* are as described for Java trace files.

Java error message output for the IBM WebSphere MQ Explorer and for IBM WebSphere MQ Transport for SOAP is written to the file specified by *Diagnostics.Java.Errors.Destination.FileName* for the appropriate Java process. The format of these files matches closely the format of the standard IBM WebSphere MQ error logs.

When a process is writing trace information to a file, it appends to a single trace output file for the lifetime of the process. Similarly, a single FFDC output file is used for the lifetime of a process.

All trace output is in the UTF-8 character set.

None of the preceding information applies to output of FFDC data or error messages for the IBM WebSphere MQ Java classes. This occurs as part of exception logging as detailed in  Using WebSphere MQ classes for Java (*WebSphere MQ V7.1 Programming Guide*). The IBM WebSphere MQ Java class trace is also detailed in  Using WebSphere MQ classes for Java (*WebSphere MQ V7.1 Programming Guide*).

Problem determination on z/OS

WebSphere MQ for z/OS, CICS, Db2, and IMS produce diagnostic information which can be used for problem determination.

This section contains information about the following topics:

- The recovery actions attempted by the queue manager when a problem is detected.
- WebSphere MQ for z/OS abends, and the information produced when an abend occurs.
- The diagnostic information produced by WebSphere MQ for z/OS, and additional sources of useful information.

The type of information provided to help with problem determination and application debugging depends on the type of error encountered, and the way your subsystem is set up.

See the following links for more information about problem determination and diagnostic information on IBM WebSphere MQ for z/OS:

- “IBM WebSphere MQ for z/OS performance constraints” on page 1266
- “WebSphere MQ for z/OS recovery actions” on page 1268
- “WebSphere MQ for z/OS abends” on page 1269
- “Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272
- “Other sources of information” on page 1274
- “Diagnostic aids for CICS” on page 1275
- “Diagnostic aids for IMS” on page 1275
- “Diagnostic aids for Db2” on page 1276
- “WebSphere MQ dumps” on page 1276
- “Dealing with performance problems on z/OS” on page 1293
- “Dealing with incorrect output” on page 1299

Related concepts:

“Troubleshooting overview” on page 1149

“Using trace” on page 1234

“Using logs” on page 1226

“First Failure Support Technology (FFST)” on page 1315

IBM WebSphere MQ for z/OS performance constraints

Use this topic to investigate z/OS resources that can cause performance constraints.

There are a number of decisions to be made when customizing IBM WebSphere MQ for z/OS that can affect the way your systems perform. These decisions include:

- The size and placement of data sets
- The allocation of buffers
- The distribution of queues among page sets, and Coupling Facility structures
- The number of tasks that you allow to access the queue manager at any one time

Log buffer pools

Insufficient log buffers can cause applications to wait until a log buffer is available, which can affect IBM WebSphere MQ performance. RMF reports might show heavy I/O to volumes that hold log data sets.

There are three parameters you can use to tune log buffers. The most important is OUTBUFF. If the log manager statistic QJSTWTB is greater than 0, increase the size of the log buffer. This parameter controls the number of buffers to be filled before they are written to the active log data sets (in the range 1 - 256). Commits and out-of-syncpoint processing of persistent messages cause log buffers to be written out to the log. As a result this parameter might have little effect except when processing large messages, and the number of commits or out of sync point messages is low. These parameters are specified in the

CSQ6LOGP macro (see  *Configuring z/OS (WebSphere MQ V7.1 Installing Guide)* for details), and the significant ones are:

OUTBUFF

This parameter controls the size of the output buffer (in the range 40 KB through 4000 KB).

WRTHRSH

This parameter controls the number of buffers to be filled before they are written to the active log data sets (in the range 1 through 256).

You must also be aware of the LOGLOAD parameter of the CSQ6SYSP macro. This parameter specifies the number of log records that are written between checkpoint records. The range is 200 through 16 000 000 but a typical value for a large system is 500 000. If a value is too small you receive frequent checkpoints, which consume processor time and can cause additional disk I/O.

Buffer pool size

There is a buffer pool associated with each page set. You can specify the number of buffers in the buffer pool using the DEFINE BUFFPOOL command. See  *Configuring z/OS (WebSphere MQ V7.1 Installing Guide)* for more information.

Incorrect specification of buffer pool size can adversely affect IBM WebSphere MQ performance. The smaller the buffer pool, the more frequently physical I/O is required. RMF might show heavy I/O to volumes that hold page sets. For buffer pools with only short-lived messages the buffer manager statistics QPSTSLA, QPSTSOS, and QPSTRIO must typically be zero. For other buffer pools, QPSTSOS and QPSTSTLA must be zero.

Distribution of data sets on available DASD

The distribution of page data sets on DASD can have a significant effect on the performance of IBM WebSphere MQ.

Place log data sets on low usage volumes with log n and log $n+1$ on different volumes. Ensure that dual logs are placed on DASD on different control units and that the volumes are not on the same physical disk.

Distribution of queues on page sets

The distribution of queues on page sets can affect performance. This change in performance can be indicated by poor response times experienced by transactions using specific queues that reside on heavily used page sets. RMF reports might show heavy I/O to volumes containing the affected page sets.

You can assign queues to specific page sets by defining storage class (STGCLASS) objects specifying a particular page set, and then defining the STGCLASS parameter in the queue definition. It is a good idea to define heavily used queues on different page sets in this way.

Distribution of queues on Coupling Facility structures

The distribution of queues on Coupling Facility structures can affect performance.

A queue-sharing group can connect to up to 64 Coupling Facility structures, one of which must be the administration structure. You can use the remaining 63 Coupling Facility structures for IBM WebSphere MQ data with each structure holding up to 512 queues. If you need more than one Coupling Facility structure, separate the queues across several structures based on the function of the queue.

There are some steps you can take to maximize efficiency:

- Delete any Coupling Facility structures you no longer require.
- Place all the queues used by an application on the same Coupling Facility to make application processing efficient.
- If work is particularly performance sensitive, choose a faster Coupling Facility structure.

Consider that if you lose a Coupling Facility structure, you lose any non-persistent messages stored in it. The loss of these non-persistent messages can cause consistency problems if queues are spread across various Coupling Facility structures. To use persistent messages, you must define the Coupling Facility structures with at least CFLEVEL(3) and RECOVER(YES).

Limitation of concurrent threads

The number of tasks accessing the queue manager can also affect performance, particularly if there are other constraints, such as storage, or there are many tasks accessing a few queues. The symptoms can be heavy I/O against one or more page sets, or poor response times from tasks known to access the same queues. The number of threads in IBM WebSphere MQ is limited to 32767 for both TSO and Batch.

In a CICS environment, you can use CICS MAXTASK to limit concurrent access.

Using the IBM WebSphere MQ trace for administration

Although you might have to use specific traces on occasion, using the trace facility has a negative effect on the performance of your systems.

Consider what destination you want your trace information sent to. Using the internal trace table saves I/O, but it is not large enough for traces that produce large volumes of data.

The statistics trace gathers information at intervals. The intervals are controlled by the STATIME parameter of the CSQ6SYSP macro, described in  [Configuring z/OS \(WebSphere MQ V7.1 Installing Guide\)](#). An accounting trace record is produced when the task or channel ends, which might be after many days.

You can limit traces by class, resource manager identifier (RMID), and instrumentation facility identifier (IFCID) to reduce the volume of data collected. See  [WebSphere MQ Script \(MQSC\) Command Reference \(WebSphere MQ V7.1 Reference\)](#) for more information.

WebSphere MQ for z/OS recovery actions

Use this topic to understand some of the recovery actions for user detected and queue manager detected errors.

WebSphere MQ for z/OS can recover from program checks caused by incorrect user data. A completion and reason code are issued to the caller. These codes are documented in  [IBM WebSphere MQ for z/OS messages, completion, and reason codes \(WebSphere MQ V7.1 Reference\)](#).

Program errors

Program errors might be associated with user application program code or WebSphere MQ code, and fall into two categories:

- User detected errors
- Subsystem detected errors

User detected errors

User detected errors are detected by the user (or a user-written application program) when the results of a service request are not as expected (for example, a nonzero completion code). The collection of problem determination data cannot be automated because detection occurs after the WebSphere MQ function has completed. Rerunning the application with the WebSphere MQ user parameter trace facility activated can provide the data needed to analyze the problem. The output from this trace is directed to the *generalized trace facility* (GTF).

You can turn the trace on and off using an operator command. See “Using trace for problem determination on z/OS” on page 1244 for more information.

Queue manager detected errors

The queue manager detects errors such as:

- A program check
- A data set filling up
- An internal consistency error

WebSphere MQ analyzes the error and takes the following actions:

- If the problem was caused by a user or application error (such as an invalid address being used), the error is reflected back to the application by completion and reason codes.
- If the problem was not caused by a user or application error (for example, all available DASD has been used, or the system detected an internal inconsistency), WebSphere MQ recovers if possible, either by sending completion and reason codes to the application, or if this is not possible, by stopping the application.

- If WebSphere MQ cannot recover, it terminates with a specific reason code. An SVC dump is typically taken recording information in the *system diagnostic work area* (SDWA) and *variable recording area* (VRA) portions of the dump, and an entry is made in SYS1.LOGREC.

WebSphere MQ for z/OS abends

Abends can occur in WebSphere for z/OS or other z/OS systems. Use this topic to understand the WebSphere MQ system abend codes and how to investigate abends which occur in CICS, IMS, and z/OS.

WebSphere MQ for z/OS uses two system abend completion codes, X'5C6' and X'6C6'. These codes identify:

- Internal errors encountered during operation
- Diagnostic information for problem determination
- Actions initiated by the component involved in the error

X'5C6'

An X'5C6' abend completion code indicates that WebSphere MQ has detected an internal error and has terminated an internal task (TCB) or a user-connected task abnormally. Errors associated with a X'5C6' abend completion code might be preceded by a z/OS system code, or by internal errors.

Examine the diagnostic material generated by the X'5C6' abend to determine the source of the error that actually resulted in a subsequent task or subsystem termination.

X'6C6'

An X'6C6' abend completion code indicates that WebSphere MQ has detected a severe error and has terminated the queue manager abnormally. When a X'6C6' is issued, WebSphere MQ has determined that continued operation could result in the loss of data integrity. Errors associated with a X'6C6' abend completion code might be preceded by a z/OS system error, one or more X'5C6' abend completion codes, or by error message CSQV086E indicating abnormal termination of WebSphere MQ.

Table 138 summarizes the actions and diagnostic information available to WebSphere MQ for z/OS when these abend completion codes are issued. Different pieces of this information are relevant in different error situations. The information produced for a particular error depends upon the specific problem. For more information about the z/OS services that provide diagnostic information, see "Diagnostic information produced on IBM WebSphere MQ for z/OS" on page 1272.

Table 138. Abend completion codes

	X'5C6'	X'6C6'
Explanation	<ul style="list-style-type: none"> • Error during WebSphere MQ normal operation 	<ul style="list-style-type: none"> • Severe error; continued operation might jeopardize data integrity
System action	<ul style="list-style-type: none"> • Internal WebSphere MQ task is abended • Connected user task is abended 	<ul style="list-style-type: none"> • The entire WebSphere MQ subsystem is abended • User task with an active WebSphere MQ connection might be abnormally terminated with a X'6C6' code • Possible MEMTERM (memory termination) of connected allied address space
Diagnostic information	<ul style="list-style-type: none"> • SVC dump • SYS1.LOGREC entry • VRA data entries 	<ul style="list-style-type: none"> • SYS1.LOGREC • VRA data entries

Table 138. Abend completion codes (continued)

	X'5C6'	X'6C6'
Associated reason codes	<ul style="list-style-type: none"> • WebSphere MQ abend reason code • Associated z/OS system codes 	<ul style="list-style-type: none"> • Subsystem termination reason code • z/OS system completion codes and X'5C6' codes that precede the X'6C6' abend
Location of accompanying codes	<ul style="list-style-type: none"> • SVC dump title • Message CSQW050I • Register 15 of SDWA section 'General Purpose Registers at Time of Error' • SYS1.LOGREC entries • VRA data entries 	<ul style="list-style-type: none"> • SYS1.LOGREC • VRA data entries • Message CSQV086E, which is sent to z/OS system operator

Related concepts:

“Dealing with program abends”

“CICS, IMS, and z/OS abends” on page 1271

“Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272

“WebSphere MQ dumps” on page 1276

Dealing with program abends

Abends can occur with applications and other z/OS systems. Use this topic to investigate the various types of abends that can occur when using IBM WebSphere MQ for z/OS.

Types of abend

Program abends can be caused by applications failing to check, and respond to, reason codes from IBM WebSphere MQ. For example, if a message has not been received, using fields that would have been set up in the message for calculation might cause X'0C4' or X'0C7' abends (ASRA abends in CICS).

The following pieces of information indicate a program abend:

- Error messages from IBM WebSphere MQ in the console log
- CICS error messages
- CICS transaction dumps
- IMS region dumps
- IMS messages on user or master terminal
- Program dump information in batch or TSO output
- Abend messages in batch job output
- Abend messages on the TSO screen

If you have an abend code, see one of the following manuals for an explanation of the cause of the abend:

- For IBM WebSphere MQ for z/OS abends (abend codes X'5C6' and X'6C6'), see  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)
- For batch abends, the *MVS System Codes* manual
- For CICS abends, the *CICS Messages and Codes* manual
- For IMS abends, the *IMS/ESA Messages and Codes* manual
- For Db2 abends, the *Db2 Messages and Codes* manual
- For RRS abends, the *MVS System Messages* manual
- For XES abends, the *MVS System Messages* manual

Batch abends

Batch abends cause an error message containing information about the contents of registers to be displayed in the syslog. TSO abends cause an error message containing similar information to be produced on the TSO screen. A SYSUDUMP is taken if there is a SYSUDUMP DD statement for the step (see “WebSphere MQ dumps” on page 1276).

CICS transaction abends

CICS transaction abends are recorded in the CICS CSMT log, and a message is produced at the terminal (if there is one). A CICS AICA abend indicates a possible loop. See “Dealing with loops” on page 1298 for more information. If you have a CICS abend, using CEDF and the CICS trace might help you to find the cause of the problem. See the *CICS Problem Determination Guide* for more information.

IMS transaction abends

IMS transaction abends are recorded on the IMS master terminal, and an error message is produced at the terminal (if there is one). If you have an IMS abend, see the *IMS/ESA Diagnosis Guide and Reference* manual.

CICS, IMS, and z/OS abends

Use this topic to investigate abends from CICS, IMS, and z/OS.

CICS abends

A CICS abend message is sent to the terminal, if the application is attached to one, or to the CSMT log. CICS abend codes are explained in the *CICS Messages and Codes* manual.

The CICS adapter issues abend reason codes beginning with the letter Q (for example, QDCL). These codes are documented in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

IMS abends

An IMS application might abend in one of the following circumstances:

- A normal abend.
- An IMS pseudo abend, with an abend code such as U3044 resulting from an error in an ESAF exit program.
- Abend 3051 or 3047, when the REO (region error option) has been specified as "Q" or "A", and an IMS application attempts to reference a non-operational external subsystem, or when resources are unavailable at the time when a thread is created.

An IMS message is sent to the user terminal or job output, and the IMS master terminal. The abend might be accompanied by a region dump.

z/OS abends

During WebSphere MQ operation, an abend might occur with a z/OS system completion code. If you receive a z/OS abend, see the appropriate z/OS publication.

Diagnostic information produced on IBM WebSphere MQ for z/OS

Use this topic to investigate some of the diagnostic information produced by z/OS that can be useful in problem determination and understand how to investigate error messages, dumps, console logs, job output, symptom strings, and queue output.

WebSphere MQ for z/OS functional recovery routines use z/OS services to provide diagnostic information to help you in problem determination.

The following z/OS services provide diagnostic information:

SVC dumps

The WebSphere MQ abend completion code X'5C6' uses the z/OS SDUMP service to create SVC dumps. The content and storage areas associated with these dumps vary, depending on the specific error and the state of the queue manager at the time the error occurred.

SYS1.LOGREC

Entries are requested in the SYS1.LOGREC data set at the time of the error using the z/OS SETRP service. The following are also recorded in SYS1.LOGREC:

- Subsystem abnormal terminations
- Secondary abends occurring in a recovery routine
- Requests from the recovery termination manager

Variable recording area (VRA) data

Data entries are added to the VRA of the SDWA by using a z/OS VRA defined key. VRA data includes a series of diagnostic data entries common to all WebSphere MQ for z/OS abend completion codes. Additional information is provided during initial error processing by the invoking component recovery routine, or by the recovery termination manager.

WebSphere MQ for z/OS provides unique messages that, together with the output of dumps, are aimed at providing sufficient data to allow diagnosis of the problem without having to try to reproduce it. This is known as first failure data capture.

Error messages

WebSphere MQ produces an error message when a problem is detected. WebSphere MQ diagnostic messages begin with the prefix CSQ. Each error message generated by WebSphere MQ is unique; that is, it is generated for one and only one error. Information about the error can be found in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*).

The first three characters of the names of WebSphere MQ modules are also usually CSQ. The exceptions to this are modules for C++ (IMQ), and the header files (CMQ). The fourth character uniquely identifies the component. These identifiers are listed in “WebSphere MQ component and resource manager identifiers” on page 1344. Characters five through eight are unique within the group identified by the first four characters.

Make sure that you have some documentation on application messages and codes for programs that were written at your installation, as well as viewing  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

There might be some instances when no message is produced, or, if one is produced, it cannot be communicated. In these circumstances, you might have to analyze a dump to isolate the error to a particular module. For more information about the use of dumps, see “WebSphere MQ dumps” on page 1276.

Dumps

Dumps are an important source of detailed information about problems. Whether they are as the result of an abend or a user request, they allow you to see a snapshot of what was happening at the moment the dump was taken. “WebSphere MQ dumps” on page 1276 contains guidance about using dumps to locate problems in your WebSphere MQ system. However, because they only provide a snapshot, you might need to use them with other sources of information that cover a longer period of time, such as logs.

Snap dumps are also produced for specific types of error in handling MQI calls. The dumps are written to the CSQSNAP DD.

Console logs and job output

You can copy console logs into a permanent data set, or print them as required. If you are only interested in specific events, you can select which parts of the console log to print.

Job output includes output produced from running the job, as well as that from the console. You can copy this output into permanent data sets, or print it as required. You might need to collect output for all associated jobs, for example CICS, IMS, and WebSphere MQ.

Symptom strings

Symptom strings display important diagnostic information in a structured format. When a symptom string is produced, it is available in one or more of the following places:

- On the z/OS system console
- In SYS1.LOGREC
- In any dump taken

Figure 154 shows an example of a symptom string.

```
PIDS/5655R3600 RIDS/CSQMAIN1 AB/S6C6 PRCS/0E30003
```

Figure 154. Sample symptom string

The symptom string provides a number of keywords that you can use to search the IBM software support database. If you have access to one of the optional search tools, you can search the database yourself. If you report a problem to the IBM support center, you are often asked to quote the symptom string. (For more information about searching the IBM software support database, See “Searching the IBM database for similar problems, and solutions” on page 1327.)

Although the symptom string is designed to provide keywords for searching the database, it can also give you a lot of information about what was happening at the time the error occurred, and it might suggest an obvious cause or a promising area to start your investigation. See “Building a keyword string” on page 1331 for more information about keywords.

Queue information

You can display information about the status of queues by using the operations and control panels. Alternatively you can enter the DISPLAY QUEUE and DISPLAY QSTATUS commands from the z/OS console.

Note: If the command was issued from the console, the response is copied to the console log, allowing the documentation to be kept together compactly.

Related concepts:

“Using trace for problem determination on z/OS” on page 1244

“Other sources of information”

“Diagnostic aids for CICS” on page 1275

“Diagnostic aids for IMS” on page 1275

“Diagnostic aids for Db2” on page 1276

Other sources of information

Use this topic to investigate other sources of information for problem determination.

You might find the following items of documentation useful when solving problems with WebSphere MQ for z/OS.

- Your own documentation
- Documentation for the products you are using
- Source listings and link-edit maps
- Change log
- System configuration charts
- Information from the DISPLAY CONN command

Your own documentation

Your own documentation is the collection of information produced by your organization about what your system and applications should do, and how they are supposed to do it. How much of this information you need depends on how familiar you are with the system or application in question, and could include:

- Program descriptions or functional specifications
- Flowcharts or other descriptions of the flow of activity in a system
- Change history of a program
- Change history of your installation
- Statistical and monitoring profile showing average inputs, outputs, and response times

Documentation for the products you are using

The documentation for the product you are using are the InfoCenters in the WebSphere MQ library, and in the libraries for any other products you use with your application.

Make sure that the level of any documentation you refer to matches the level of the system you are using. Problems often arise through using either obsolete information, or information about a level of a product that is not yet installed.

Source listings and link-edit maps

Include the source listings of any applications written at your installation with your set of documentation. (They can often be the largest single element of documentation.) Make sure that you include the relevant output from the linkage editor with your source listings to avoid wasting time trying to find your way through a load module with an out-of-date link map. Be sure to include the JCL at the beginning of your listings, to show the libraries that were used and the load library the load module was placed in.

Change log

The information in the change log can tell you of changes made in the data processing environment that might have caused problems with your application program. To get the most out of your change log, include the data concerning hardware changes, system software (such as z/OS and WebSphere MQ) changes, application changes, and any modifications made to operating procedures.

System configuration charts

System configuration charts show what systems are running, where they are running, and how the systems are connected to each other. They also show which WebSphere MQ, CICS, or IMS systems are test systems and which are production systems.

Information from the DISPLAY CONN command

The DISPLAY CONN command provides information about which applications are connected to a queue manager, and information to help you to diagnose those that have a long-running unit of work. You could collect this information periodically and check it for any long-running units of work, and display the detailed information about that connection.

Diagnostic aids for CICS

You can use the CICS diagnostic transactions to display information about queue manager tasks, and MQI calls. Use this topic to investigate these facilities.

You can use the CKQC transaction (the CICS adapter control panels) to display information about queue manager tasks, and what state they are in (for example, a GET WAIT). See “Administering IBM WebSphere MQ for z/OS” on page 237 for more information about CKQC.

The application development environment is the same as for any other CICS application, and so you can use any tools normally used in that environment to develop WebSphere MQ applications. In particular, the *CICS execution diagnostic facility* (CEDF) traps entry to and exit from the CICS adapter for each MQI call, as well as trapping calls to all CICS API services. Examples of the output produced by this facility are given in  Examples of CEDF output (*WebSphere MQ V7.1 Reference*).

The CICS adapter also writes trace entries to the CICS trace. These entries are described in “CICS adapter trace entries” on page 1252.

Additional trace and dump data is available from the CICS region. These entries are as described in the *CICS Problem Determination Guide*.

Diagnostic aids for IMS

Use this topic to investigate IMS diagnostic facilities.

The application development environment is the same as for any other IMS application, and so any tools normally used in that environment can be used to develop WebSphere MQ applications.

Trace and dump data is available from the IMS region. These entries are as described in the *IMS/ESA Diagnosis Guide and Reference* manual.

Diagnostic aids for Db2

Use this topic to investigate references for Db2 diagnostic tools.

Refer to the following manuals for help in diagnosing Db2 problems:

- *Db2 for z/OS Diagnosis Guide and Reference*
- *Db2 Messages and Codes*

WebSphere MQ dumps

Use this topic for information about the use of dumps in problem determination. It describes the steps you should take when looking at a dump produced by a WebSphere MQ for z/OS address space.

How to use dumps for problem determination

When solving problems with your WebSphere MQ for z/OS system, you can use dumps in two ways:

- To examine the way WebSphere MQ processes a request from an application program.
To do this, you typically need to analyze the whole dump, including control blocks and the internal trace.
- To identify problems with WebSphere MQ for z/OS itself, under the direction of IBM support center personnel.

Use the instructions in the following topics to get and process a dump:

- “Getting a dump” on page 1277
- “Using the z/OS DUMP command” on page 1277
- “Processing a dump using the WebSphere MQ for z/OS dump display panels” on page 1279
- “Processing a dump using line mode IPCS” on page 1283
- “Processing a dump using IPCS in batch” on page 1288

The dump title might provide sufficient information in the abend and reason codes to resolve the problem. You can see the dump title in the console log, or by using the z/OS command `DISPLAY DUMP,TITLE`. The format of the dump title is explained in “Analyzing the dump and interpreting dump titles” on page 1288. For information about the WebSphere MQ for z/OS abend codes, see “WebSphere

MQ for z/OS abends” on page 1269, and abend reason codes are documented in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*).

If there is not enough information about your problem in the dump title, format the dump to display the other information contained in it.

See the following topics for information about different types of dumps:

- “SYSUDUMP information” on page 1290
- “Snap dumps” on page 1291
- “SYS1.LOGREC information” on page 1292
- “SVC dumps” on page 1292

Related concepts:

“Using trace for problem determination on z/OS” on page 1244

“WebSphere MQ for z/OS abends” on page 1269

“Diagnostic information produced on IBM WebSphere MQ for z/OS” on page 1272

Getting a dump

Use this topic to understand the different dump types for problem determination.

The following table shows information about the types of dump used with WebSphere MQ for z/OS and how they are initiated. It also shows how the dump is formatted:

Table 139. Types of dump used with WebSphere MQ for z/OS

Dump type	Data set	Output type	Formatted by	Caused by
SVC	Defined by system	Machine readable	IPCS in conjunction with a WebSphere MQ for z/OS verb exit	z/OS or WebSphere MQ for z/OS functional recovery routine detecting error, or the operator entering the z/OS DUMP command
SYSUDUMP	Defined by JCL (SYSOUT=A)	Formatted	Normally SYSOUT=A	An abend condition (only taken if there is a SYSUDUMP DD statement for the step)
Snap	Defined by JCL CSQSNAP (SYSOUT=A)	Formatted	Normally SYSOUT=A	Unexpected MQI call errors reported to adapters, or FFST information from the channel initiator
Stand-alone	Defined by installation (tape or disk)	Machine readable	IPCS in conjunction with a WebSphere MQ for z/OS verb exit	Operator IPL of the stand-alone dump program

WebSphere MQ for z/OS recovery routines request SVC dumps for most X'5C6' abends. The exceptions are listed in “SVC dumps” on page 1292. SVC dumps issued by WebSphere MQ for z/OS are the primary source of diagnostic information for problems.

If the dump is initiated by the WebSphere MQ subsystem, information about the dump is put into area called the *summary portion*. This contains information that the dump formatting program can use to identify the key components.

For more information about SVC dumps, see the *MVS Diagnosis: Tools and Service Aids* manual.

Using the z/OS DUMP command

To resolve a problem, IBM can ask you to create a dump file of the queue manager address space, channel initiator address space, or coupling facilities structures. Use this topic to understand the commands to create these dump files.

You might be asked to create dump file for any or several of the following items for IBM to resolve the problem:

- Main WebSphere MQ address space
- Channel initiator address space
- Coupling facility application structure
- Coupling facility administration structure for your queue-sharing group

Figure 155 on page 1278 through to Figure 159 on page 1279 show examples of the z/OS commands to do this, assuming a subsystem name of CSQ1.

```

DUMP COMM=(MQ QUEUE MANAGER DUMP)
*01 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
R 01,JOBNAME=(CSQ1MSTR,BATCH),CONT
*02 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 01 IS;JOBNAME=CSQ1MSTR,CONT
R 02,SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),END
  IEE600I REPLY TO 02 IS;SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),END
IEA794I SVC DUMP HAS CAPTURED: 869
DUMPID=001 REQUESTED BY JOB (*MASTER*)
DUMP TITLE=MQ QUEUE MANAGER MAIN DUMP

```

Figure 155. Dumping the WebSphere MQ queue manager and application address spaces

```

DUMP COMM=(MQ QUEUE MANAGER DUMP)
*01 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
R 01,JOBNAME=(CSQ1MSTR),CONT
*02 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 01 IS;JOBNAME=CSQ1MSTR,CONT
R 02,SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),END
  IEE600I REPLY TO 02 IS;SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),END
IEA794I SVC DUMP HAS CAPTURED: 869
DUMPID=001 REQUESTED BY JOB (*MASTER*)
DUMP TITLE=MQ QUEUE MANAGER DUMP

```

Figure 156. Dumping the WebSphere MQ queue manager address space

```

DUMP COMM=(MQ CHIN DUMP)
*01 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
R 01,JOBNAME=CSQ1CHIN,CONT
*02 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 01 IS;JOBNAME=CSQ1CHIN,CONT
R 02,SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),CONT
*03 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 02 IS;SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),CONT
R 03,DSPNAME=('CSQ1CHIN'.CSQXTRDS),END
  IEE600I REPLY TO 03 IS;DSPNAME='CSQ1CHIN'.CSQXTRDS,END
IEA794I SVC DUMP HAS CAPTURED: 869
DUMPID=001 REQUESTED BY JOB (*MASTER*)
DUMP TITLE=MQ CHIN DUMP

```

Figure 157. Dumping the channel initiator address space

```

DUMP COMM=(MQ MSTR & CHIN DUMP)
*01 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
R 01,JOBNAME=(CSQ1MSTR,CSQ1CHIN),CONT
*02 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 01 IS;JOBNAME=(CSQ1MSTR,CSQ1CHIN),CONT
R 02,SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),CONT
*03 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
  IEE600I REPLY TO 02 IS;SDATA=(CSA,RGN,PSA,SQA,LSQA,TRT,SUM),CONT
R 03,DSPNAME=('CSQ1CHIN'.CSQXTRDS),END
  IEE600I REPLY TO 03 IS;DSPNAME=('CSQ1CHIN'.CSQXTRDS),END
IEA794I SVC DUMP HAS CAPTURED: 869
DUMPID=001 REQUESTED BY JOB (*MASTER*)
DUMP TITLE=MQ MSTR & CHIN DUMP

```

Figure 158. Dumping the WebSphere MQ queue manager and channel initiator address spaces

```

DUMP COMM=('MQ APPLICATION STRUCTURE 1 DUMP')
01 IEE094D SPECIFY OPERAND(S) FOR DUMP COMMAND
R 01,STRLIST=(STRNAME=QSG1APPLICATION1,(LISTNUM=ALL,ADJUNCT=CAPTURE,ENTRYDATA=UNSER))
IEE600I REPLY TO 01 IS;STRLIST=(STRNAME=QSG1APPLICATION1,(LISTNUM=
IEA794I SVC DUMP HAS CAPTURED: 677
DUMPID=057 REQUESTED BY JOB (*MASTER*)
DUMP TITLE='MQ APPLICATION STRUCTURE 1 DUMP'

```

Figure 159. Dumping a coupling facility structure

Processing a dump using the WebSphere MQ for z/OS dump display panels

You can use commands available through IPCS panels to process dumps. Use this topic to understand the IPCS options.

WebSphere MQ for z/OS provides a set of panels to help you process dumps. The following section describes how to use these panels:

1. From the IPCS PRIMARY OPTION MENU, select **ANALYSIS - Analyze dump contents** (option 2).
The IPCS MVS ANALYSIS OF DUMP CONTENTS panel is displayed.
2. Select **COMPONENT - MVS component data** (option 6).
The IPCS MVS DUMP COMPONENT DATA ANALYSIS panel is displayed. The appearance of the panel depends on the products installed at your installation, but will be similar to the panel shown in IPCS MVS Dump Component Data Analysis panel:

```

----- IPCS MVS DUMP COMPONENT DATA ANALYSIS -----
OPTION ==>                                     SCROLL ==

```

To display information, specify "S option name" or enter S to the left of the option desired. Enter ? to the left of an option to display help regarding the component support.

Name	Abstract
ALCWAIT	Allocation wait summary
AOMDATA	AOM analysis
ASMCHECK	Auxiliary storage paging activity
ASMDATA	ASM control block analysis
AVMDATA	AVM control block analysis
COMCHECK	Operator communications data
CSQMAIN	WebSphere MQ dump formatter panel interface
CSQWDMP	WebSphere MQ dump formatter
CTRACE	Component trace summary
DAEDATA	DAE header data
DIVDATA	Data-in-virtual storage

Figure 160. IPCS MVS Dump Component Data Analysis panel

3. Select **CSQMAIN WebSphere MQ dump formatter panel interface** by typing s beside the line and pressing Enter.

If this option is not available, it is because the member CSQ7IPCS is not present; you should see



Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*) for more information about installing the WebSphere MQ for z/OS dump formatting member.

Note: If you have already used the dump to do a preliminary analysis, and you want to reexamine it, select **CSQWDMP WebSphere MQ dump formatter** to display the formatted contents again, using the default options.

4. The IBM WebSphere MQ for z/OS - DUMP ANALYSIS menu is displayed. Use this menu to specify the action that you want to perform on a system dump.

-----IBM WebSphere MQ for z/OS - DUMP ANALYSIS-----

COMMAND ===>

- 1 Display all dump titles 00 through 99
- 2 Manage the dump inventory
- 3 Select a dump

- 4 Display address spaces active at time of dump
- 5 Display the symptom string
- 6 Display the symptom string and other related data
- 7 Display LOGREC data from the buffer in the dump
- 8 Format and display the dump

- 9 Issue IPCS command or CLIST

(c) Copyright IBM Corporation 1993, 2019. All rights reserved.

F1=Help F3=Exit F12=Cancel

5. Before you can select a particular dump for analysis, the dump you require must be present in the dump inventory. To ensure that this is so, perform the following steps:
 - a. If you do not know the name of the data set containing the dump, specify option 1 - **Display all dump titles xx through xx**.

This displays the dump titles of all the dumps contained in the SYS1.DUMP data sets (where xx is a number in the range 00 through 99). You can limit the selection of data sets for display by using the xx fields to specify a range of data set numbers.

If you want to see details of all available dump data sets, set these values to 00 and 99.

Use the information displayed to identify the dump you want to analyze.
 - b. If the dump has not been copied into another data set (that is, it is in one of the SYS1.DUMP data sets), specify option 2 - **Manage the dump inventory**.

The dump inventory contains the dump data sets that you have used. Because the SYS1.DUMP data sets are reused, the name of the dump that you identified in step 5a might be in the list displayed. However, this entry refers to the previous dump that was stored in this data set, so delete it by typing DD next to it and pressing Enter. Then press F3 to return to the DUMP ANALYSIS MENU.
6. Specify option 3 - **Select a dump**, to select the dump that you want to work with. Type the name of the data set containing the dump in the Source field, check that NOPRINT and TERMINAL are specified in the Message Routing field (this is to ensure that the output is directed to the terminal), and press Enter. Press F3 to return to the DUMP ANALYSIS MENU.
7. Having selected a dump to work with, you can now use the other options on the menu to analyze the data in different parts of the dump:
 - To display a list of all address spaces active at the time the dump was taken, select option 4.
 - To display the symptom string, select option 5. If you want to use the symptom string to search the RETAIN database for solutions to similar problems, refer to "Searching the IBM database for similar problems, and solutions" on page 1327.
 - To display the symptom string and other serviceability information, including the variable recording area of the system diagnostic work area (SDWA), select option 6.
 - To format and display the data contained in the in-storage LOGREC buffer, select option 7.

It could be that the abend that caused the dump was not the original cause of the error, but was caused by an earlier problem. To determine which LOGREC record relates to the cause of the problem, go to the bottom of the data set, type FIND ERRORID: PREV, and press Enter. The header of the latest LOGREC record is displayed, for example:

```
JOBNAME: NONE-FRR
ERRORID: SEQ=00081 CPU=0040 ASID=0033 TIME=14:42:47.1

SEARCH ARGUMENT ABSTRACT

PIDS/5655R3600 RIDS/CSQRLLM1#L RIDS/CSQRRHSL AB/S05C6
PRCS/00D10231 REGS/0C1F0 RIDS/CSQVEUS2#R

SYMPTOM          DESCRIPTION
-----          -
PIDS/5655R3600   PROGRAM ID: 5655R3600
.
.
.
```

Note the program identifier (if it is not 5655R3600, the problem was not caused by WebSphere MQ for z/OS and you could be looking at the wrong dump). Also note the value of the TIME field. Repeat the command to find the previous LOGREC record, and note the value of the TIME field again. If the two values are close to each other (say, within about one or two tenths of a second), they could both relate to the same problem.

You can use the symptom string from the LOGREC record related to the error to search the RETAIN database for solutions to similar problems (see “Searching the IBM database for similar problems, and solutions” on page 1327).

- To format and display the dump, select option 8. The FORMAT AND DISPLAY THE DUMP panel is displayed:

```
-----IBM WebSphere MQ for z/OS - FORMAT AND DISPLAY DUMP-----
COMMAND ==>

  1 Display the control blocks and trace
  2 Display just the control blocks
  3 Display just the trace

Options:

Use the summary dump? . . . . . __ 1 Yes
                                     2 No

Subsystem name (required if summary dump not used) ____

Address space identifier or ALL. . . . . ALL_

F1=Help   F3=Exit   F12=Cancel
```

- Use this panel to format your selected system dump. You can choose to display control blocks, data produced by the internal trace, or both, which is the default.

Note: You cannot do this for dumps from the channel initiator, or for dumps of coupling facility structures.

- To display the whole of the dump, that is:
 - The dump title
 - The variable recording area (VRA) diagnostic information report
 - The save area trace report
 - The control block summary
 - The trace table

select option 1.

- To display the information listed for option 1, without the trace table, select option 2.
- To display the information listed for option 1, without the control blocks, select option 3.

You can also use the following options:

- **Use the Summary Dump?**

Use this field to specify whether you want WebSphere MQ to use the information contained in the summary portion when formatting the selected dump. The default setting is YES.

Note: If a summary dump has been taken, it might include data from more than one address space.

- **Subsystem name**

Use this field to identify the subsystem with the dump data you want to display. This is only required if there is no summary data (for example, if the operator requested the dump), or if you have specified NO in the **Use the summary dump?** field.

If you do not know the subsystem name, type `IPCS SELECT ALL` at the command prompt, and press Enter to display a list of all the jobs running at the time of the error. If one of the jobs has the word ERROR against it in the SELECTION CRITERIA column, make a note of the name of that job. The job name is of the form `xxxxMSTR`, where `xxxx` is the subsystem name.

```

IPCS OUTPUT STREAM -----
COMMAND ==>
  ASID JOBNAME  ASCBADDR  SELECTION CRITERIA
  -----
0001 *MASTER*  00FD4D80  ALL
0002 PCAUTH    00F8AB80  ALL
0003 RASP     00F8C100  ALL
0004 TRACE    00F8BE00  ALL
0005 GRS      00F8BC00  ALL
0006 DUMPSRV  00F8DE00  ALL
0008 CONSOLE  00FA7E00  ALL
0009 ALLOCAS  00F8D780  ALL
000A SMF      00FA4A00  ALL
000B VLF      00FA4800  ALL
000C LLA      00FA4600  ALL
000D JESM     00F71E00  ALL
001F MQM1MSTR  00FA0680  ERROR ALL
  
```

If no job has the word ERROR against it in the SELECTION CRITERIA column, select option 0 - DEFAULTS on the main IPCS Options Menu panel to display the IPCS Default Values panel. Note the address space identifier (ASID) and press F3 to return to the previous panel. Use the ASID to determine the job name; the form is `xxxxMSTR`, where `xxxx` is the subsystem name.

The following command shows which ASIDs are in the dump data set:

```
LDMP DSN('SYS1.DUMPxx') SELECT(DUMPED) NOSUMMARY
```

This shows the storage ranges dumped for each address space.

Press F3 to return to the FORMAT AND DISPLAY THE DUMP panel, and type this name in the **Subsystem name** field.

– **Address space identifier**

Use this field if the data in a dump comes from more than one address space. If you only want to look at data from a particular address space, specify the identifier (ASID) for that address space.

The default value for this field is ALL, which displays information about all the address spaces relevant to the subsystem in the dump. Change this field by typing the 4-character ASID over the value displayed.

Note: Because the dump contains storage areas common to all address spaces, the information displayed might not be relevant to your problem if you specify the address space identifier incorrectly. In this case, return to this panel, and enter the correct address space identifier.

Related concepts:

“Processing a dump using line mode IPCS”

“Processing a dump using IPCS in batch” on page 1288

“Analyzing the dump and interpreting dump titles” on page 1288

Processing a dump using line mode IPCS

Use the IPCS commands to format a dump.

To format the dump using line mode IPCS commands, select the dump required by issuing the command:

```
SETDEF DSN('SYS1.DUMPxx')
```

(where SYS1.DUMPxx is the name of the data set containing the dump). You can then use IPCS subcommands to display data from the dump.

See the following topics for information on how to format different types of dumps using IPCS commands:

- “Formatting a WebSphere MQ for z/OS dump”
- “Formatting a dump from the channel initiator” on page 1287

Related concepts:

“Processing a dump using the WebSphere MQ for z/OS dump display panels” on page 1279

“Processing a dump using IPCS in batch” on page 1288

“Analyzing the dump and interpreting dump titles” on page 1288

Formatting a WebSphere MQ for z/OS dump:

Use this topic to understand how to format a queue manager dump using line mode IPCS commands.

The IPCS VERBEXIT CSQWDMP invokes the WebSphere MQ for z/OS dump formatting program (CSQWDPRD), and enables you to format an SVC dump to display WebSphere MQ data. You can restrict the amount of data that is displayed by specifying parameters.

IBM Service Personnel might require dumps of your coupling facility administration structure and application structures for your queue-sharing group, with dumps of queue managers in the queue-sharing group, to aid problem diagnosis. For information on formatting a coupling facility list structure, and the STRDATA subcommand, see the *MVS IPCS Commands* book.

Note: This section describes the parameters required to extract the necessary data. Separate operands by commas, not blanks. A blank that follows any operand in the control statement terminates the operand list, and any subsequent operands are ignored. Table 140 on page 1284 explains each keyword that you

can specify in the control statement for formatting dumps.

Table 140. Keywords for the WebSphere MQ for z/OS dump formatting control statement

Keyword	Description
SUBSYS= <i>aaaa</i>	Use this keyword if the summary dump portion is not available, or not to be used, to give the name of the subsystem to format information for. <i>aaaa</i> is a 1 through 4-character subsystem name.
ALL (default)	All control blocks and the trace table.
AA	Data is displayed for all WebSphere MQ for z/OS control blocks in all address spaces.
DIAG=Y	Print diagnostic information. Use only under guidance from IBM service personnel. DIAG=N (suppresses the formatting of diagnostic information) is the default.
EB= <i>nnnnnnnn</i>	Only the trace points associated with this EB thread are displayed (the format of this keyword is EB= <i>nnnnnnnn</i> where <i>nnnnnnnn</i> is the 8-digit address of an EB thread that is contained in the trace). You must use this in conjunction with the TT keyword.
LG	All control blocks.
PTF=Y, LOAD= <i>load module name</i>	A list of PTFs at the front of the report (from MEPL). PTF=N (suppresses the formatting of such a list) is the default. The optional load subparameter allows you to specify the name of a load module, up to a maximum of 8 characters, for which to format a PTF report.
SA= <i>hhhh</i>	The control blocks for a specified address space. Use either of the following formats: <ul style="list-style-type: none"> • SA=<i>hh</i> or • SA=<i>hhhh</i> where <i>h</i> represents a hexadecimal digit.
SG	A subset of system-wide control blocks.
TT	The trace table only.

Table 141 details the dump formatting keywords that you can use to format the data relating to individual resource managers. You cannot use these keywords in conjunction with any of the keywords in Table 140.

Table 141. Resource manager dump formatting keywords

Keyword	What is formatted
BMC=1	Buffer manager data. BMC=1 formats control blocks of all buffers.
BMC=2(<i>buffer pool number</i>)	BMC=2 formats data relating to the buffer identified in the 2-digit <i>buffer pool number</i> .
BMC=3(<i>xx/yyyyy</i>)	BMC=3 and BMC=4 display a page from a pageset, if the page is present in a buffer. (The difference between BMC=3 and BMC=4 is the route taken to the page.)
BMC=4(<i>xx/yyyyy</i>)	
CFS=1	coupling facility manager report.
DB2=1	Db2 manager report.
DMC=1, ONAM= <i>object name</i>	Data manager data. The optional ONAM subparameter allows you to specify the object name, up to a maximum of 20 characters, to limit data printed to objects starting with characters in ONAM.

Table 142. Summary dump keywords for the WebSphere MQ for z/OS dump formatting control statement

Keyword	Description
SUBSYS=aaaa	Use this keyword if the summary dump portion is not available, or not to be used, to give the name of the subsystem to format information for. aaaa is a 1 through 4-character subsystem name.
SUMDUMP=NO	Use this keyword if the dump has a summary portion, but you do not want to use it. (You would usually only do this if so directed by your IBM support center.)

The following list shows some examples of how to use these keywords:

- For default formatting of all address spaces, using information from the summary portion of the dump, use:
VERBX CSQWDMP
- To display the trace table from a dump of subsystem named MQMT, which was initiated by an operator (and so does not have a summary portion) use:
VERBX CSQWDMP 'TT,SUBSYS=MQMT'
- To display all the control blocks and the trace table from a dump produced by a subsystem abend, for an address space with ASID (address space identifier) 1F, use:
VERBX CSQWDMP 'TT,LG,SA=1F'
- To display the portion of the trace table from a dump associated with a particular EB thread, use:
VERBX CSQWDMP 'TT,EB=nnnnnnnn'
- To display message manager 1 report for local non-shared queue objects with a name begins with 'ABC' use:
VERBX CSQWDMP 'MMC=1,ONAM=ABC,Obj=MQLO'

Table 143 shows some other commands that are used frequently for analyzing dumps. For more information about these subcommands, see the *MVS IPCS Commands* manual.

Table 143. IPCS subcommands used for dump analysis

Subcommand	Description
STATUS	To display data usually examined during the initial part of the problem determination process.
STRDATA LISTNUM(ALL) ENTRYPOS(ALL) DETAIL	To format coupling facility structure data.
VERBEXIT LOGDATA	To format the in-storage LOGREC buffer records present before the dump was taken. LOGDATA locates the LOGREC entries that are contained in the LOGREC recording buffer and invokes the EREP program to format and print the LOGREC entries. These entries are formatted in the style of the normal detail edit report.
VERBEXIT TRACE	To format the system trace entries for all address spaces.
VERBEXIT SYMPTOM	To format the symptom strings contained in the header record of a system dump such as stand-alone dump, SVC dump, or an abend dump requested with a SYSUDUMP DD statement.
VERBEXIT GRSTRACE	To format diagnostic data from the major control blocks for global resource serialization.
VERBEXIT SUMDUMP	To locate and display the summary dump data that an SVC dump provides.
VERBEXIT DAEDATA	To format the dump analysis and elimination (DAE) data for the dumped system.

Related concepts:

“Formatting a dump from the channel initiator”

Formatting a dump from the channel initiator:

Use this topic to understand how to format a channel initiator dump using line mode IPCS commands.

The IPCS VERBEXIT CSQXDPRD enables you to format a channel initiator dump. You can select the data that is formatted by specifying keywords.

This section describes the keywords that you can specify.

Table 144 describes the keywords that you can specify with CSQXDPRD.

Table 144. Keywords for the IPCS VERBEXIT CSQXDPRD

Keyword	What is formatted
SUBSYS=aaaa	The control blocks of the channel initiator associated with the named subsystem. It is required for all new formatted dumps.
CHST=1, CNAM=channel name, DUMP=S F C	All channel information. The optional CNAM subparameter allows you to specify the name of a channel, up to a maximum of 20 characters, for which to format details. The optional DUMP subparameter allows you to control the extent of formatting, as follows: <ul style="list-style-type: none">• Specify DUMP=S (for “short”) to format the first line of the hexadecimal dump of the channel data.• Specify DUMP=F (for “full”) to format all lines of the data.• Specify DUMP=C (for “compressed”) to suppress the formatting of all duplicate lines in the data containing only X'00'. This is the default option
CHST=2, CNAM=channel name,	A summary of all channels, or of the channel specified by the CNAM keyword. See CHST=1 for details of the CNAM subparameter.
CHST=3, CNAM=channel name,	Data provided by CHST=2 and a program trace, line trace and formatted semaphore table print of all channels in the dump. See CHST=1 for details of the CNAM subparameter.
CLUS=1	Cluster report including the cluster repository known on the queue manager.
CLUS=2	Cluster report showing cluster registrations.
CTRACE=S F, DPRO=nnnnnnnn, TCB=nnnnnnnn	Select either a short (CTRACE=S) or full (CTRACE=F) CTRACE. The optional DPRO subparameter allows you to specify a CTRACE for the DPRO specified. The optional TCB subparameter allows you to specify a CTRACE for the job specified.
DISP=1, DUMP=S F C	Dispatcher report See CHST=1 for details of the DUMP subparameter.
BUF=1	Buffer report

Related concepts:

“Formatting a WebSphere MQ for z/OS dump” on page 1283

Processing a dump using IPCS in batch

Use this topic to understand how WebSphere MQ dumps can be formatted by IPCS commands in batch mode.

To use IPCS in batch, insert the required IPCS statements into your batch job stream (see Figure 161).

Change the data set name (DSN=) on the DUMP00 statement to reflect the dump you want to process, and insert the IPCS subcommands that you want to use.

```
//*****  
//*   RUNNING IPCS IN A BATCH JOB           *  
//*****  
//MQMDMP EXEC PGM=IKJEFT01,REGION=5120K  
//STEPLIB DD DSN=mqm.library-name,DISP=SHR  
//SYSTSPRT DD SYSOUT=*  
//IPCSPRNT DD SYSOUT=*  
//IPCSDDIR DD DSN=dump.directory-name,DISP=OLD  
//DUMP00 DD DSN=dump.name,DISP=SHR  
//SYSTSIN DD *  
    IPCS NOPARM TASKLIB(SCSQLOAD)  
    SETDEF PRINT TERMINAL DDNAME(DUMP00) NOCONFIRM  
*****  
* INSERT YOUR IPCS COMMANDS HERE, FOR EXAMPLE: *  
VERBEXIT LOGDATA  
VERBEXIT SYMPTOM  
VERBEXIT CSQWDMP 'TT,SUBSYS=QMGR'  
*****  
  
CLOSE ALL  
END  
/*
```

Figure 161. Sample JCL for printing dumps through IPCS in the z/OS environment

Related concepts:

“Processing a dump using the WebSphere MQ for z/OS dump display panels” on page 1279

“Processing a dump using line mode IPCS” on page 1283

“Analyzing the dump and interpreting dump titles”

Analyzing the dump and interpreting dump titles

Use this topic to understand how dump titles are formatted, and how to analyze a dump.

- Analyzing the dump
- Dump title variation with PSW and ASID

Analyzing the dump

The dump title includes the abend completion and reason codes, the failing load module and CSECT names, and the release identifier. For more information on the dump title see Dump title variation with PSW and ASID

The formats of SVC dump titles vary slightly, depending on the type of error.

Figure 162 on page 1289 shows an example of an SVC dump title. Each field in the title is described after the figure.

```
ssnm,ABN=5C6-00D303F2,U=AUSER,C=R3600.710.LOCK-CSQL1GET,  
M=CSQGFRCV,LOC=CSQLLPLM.CSQL1GET+0246
```

Figure 162. Sample SVC dump title

ssnm,ABN=compltn-reason

- `ssnm` is the name of the subsystem that issued the dump.
- `compltn` is the 3-character hexadecimal abend completion code (in this example, X'5C6'), prefixed by U for user abend codes.
- `reason` is the 4-byte hexadecimal reason code (in this example, X'00D303F2').

Note: The abend and reason codes might provide sufficient information to resolve the problem.

See the  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) for an explanation of the reason code.

U=userid

- `userid` is the user identifier of the user (in this example, AUSER). This field is not present for channel initiators.

C=compid.release.comp-function

- `compid` is the last 5 characters of the component identifier (explained in “The component-identifier keyword” on page 1333). The value R3600 uniquely identifies WebSphere MQ for z/OS.
- `release` is a 3-digit code indicating the version, release, and modification level of WebSphere MQ for z/OS (in this example, 710).
- `comp` is an acronym for the component in control at the time of the abend (in this example, LOCK).
- `function` is the name of a function, macro, or routine in control at the time of abend (in this example, CSQL1GET). This field is not always present.

M=module

- `module` is the name of the FRR or ESTAE recovery routine (in this example, CSQGFRCV). This field is not always present.

Note: This is not the name of the module where the abend occurred; that is given by LOC.

LOC=loadmod.csect+csect_offset

- `loadmod` is the name of the load module in control at the time of the abend (in this example, CSQLLPLM). This might be represented by an asterisk if it is unknown.
- `csect` is the name of the CSECT in control at the time of abend (in this example, CSQL1GET).
- `csect_offset` is the offset within the failing CSECT at the time of abend (in this example, 0246).

Note: The value of `csect_offset` might vary if service has been applied to this CSECT, so do not use this value when building a keyword string to search the IBM software support database.

Dump title variation with PSW and ASID

Some dump titles replace the load module name, CSECT name, and CSECT offset with the PSW (program status word) and ASID (address space identifier). Figure 163 on page 1290 illustrates this format.

```
ssnm,ABN=compltn-reason,U=userid,C=compid.release.comp-function,  
M=module,PSW=psw_contents,ASID=address_space_id
```

Figure 163. Dump title with PSW and ASID

psw_contents

- The PSW at the time of the error (for example, X'077C100000729F9C').

address_space_id

- The address space in control at the time of the abend (for example, X'0011'). This field is not present for a channel initiator.

Related concepts:

“Processing a dump using the WebSphere MQ for z/OS dump display panels” on page 1279

“Processing a dump using line mode IPCS” on page 1283

“Processing a dump using IPCS in batch” on page 1288

SYSUDUMP information

The z/OS system can create SYSUDUMPs, which can be used as part of problem determination. This topic shows a sample SYSUDUMP output and gives a reference to the tools for interpreting SYSUDUMPs.

SYSUDUMP dumps provide information useful for debugging batch and TSO application programs. For more information about SYSUDUMP dumps, see the *MVS Diagnosis: Tools and Service Aids* manual.

Figure 164 on page 1291 shows a sample of the beginning of a SYSUDUMP dump.

```

JOB MQMBXBA1  STEP TSUSER  TIME 102912  DATE 001019  ID = 000  CUID = 632202333081  PAGE 00000001
COMPLETION CODE      SYSTEM = 0C1      REASON CODE = 00000001
PSW AT ENTRY TO ABEND 078D1000 000433FC      ILC 2  INTC 000D
PSW LOAD MODULE = BXBAAB01 ADDRESS = 000433FC OFFSET = 0000A7F4

ASCB: 00F56400
+0000 ASCB..... ASCB      FWDP..... 00F60180  BWD..... 0047800  CMSF..... 019D5A30  SVRB..... 008FE9E0
+0014 SYNC..... 00000D6F  IOSP..... 00000000  TNEW..... 00D18F0  CPUS..... 00000001  ASID..... 0066
+0026 R026..... 0000      LL5..... 00      HLHI..... 01      DPHI..... 00      DP..... 9D
+002C TRQP..... 80F5D381  LDA..... 7FF154E8  RSMF..... 00      R035..... 0000      TRQI..... 42
+0038 CSCB..... 00F4D048  TSB..... 00B61938  EJST..... 00000001  8C257E00

+0048 EWST..... 9CCDE747  76A09480      JSTL..... 00141A4  ECB..... 808FEF78  UBET..... 9CCDE740
.
.
ASSB: 01946600
+0000 ASSB..... ASSB      VAFN..... 00000000  EVST..... 00000000  00000000

+0010 VFAT..... 00000000  00000000      RSV..... 000      XMCC..... 0000      XMCT.....00000000
+0020 VSC..... 00000000  NVSC..... 0000004C  ASRR..... 00000000  R02C..... 00000000  00000000 00000000
+0038      00000000  00000000

*** ADDRESS SPACE SWITCH EVENT MASK OFF (ASTESSEM = 0) ***

TCB: 008D18F0
+0000 RBP..... 008FE7D8  PIE..... 00000000  DEB..... 00B1530  TIO..... 008D4000  CMP.....805C6000
+0014 TRN..... 40000000  MSS..... 7FFF7418  PKF..... 80      FLGS..... 01000000  00
+0022 LMP..... FF      DSP..... FE      LLS..... 00D1A88  JLB..... 00011F18  JPQ.....00000000
+0030 GPR0-3... 00001000  008A4000  00000000  00000000
+0040 GPR4-7... 00FDC730  008A50C8  00000002  80E73F04
+0050 GPR8-11.. 81CC4360  008A6754  008A67B4  00000008

```

Figure 164. Sample beginning of a SYSUDUMP

Snap dumps

Snap dump data sets are controlled by z/OS JCL command statements. Use this topic to understand the CSQSNAP DD statement.

Snap dumps are always sent to the data set defined by the CSQSNAP DD statement. They can be issued by the adapters or the channel initiator.

- Snap dumps are issued by the batch, CICS, IMS, or RRS adapter when an unexpected error is returned by the queue manager for an MQI call. A full dump is produced containing information about the program that caused the problem.

For a snap dump to be produced, the CSQSNAP DD statement must be in the batch application JCL, CICS JCL, or IMS dependent region JCL.

- Snap dumps are issued by the channel initiator in specific error conditions instead of a system dump. The dump contains information relating to the error. Message CSQX053E is also issued at the same time.

To produce a snap dump, the CSQSNAP DD statement must be in the channel initiator started-task procedure.

SYS1.LOGREC information

Use this topic to understand how the z/OS SYS1.LOGREC information can assist with problem determination.

WebSphere MQ for z/OS and SYS1.LOGREC

The SYS1.LOGREC data set records various errors that different components of the operating system encounter. For more information about using SYS1.LOGREC records, see the *MVS Diagnosis: Tools and Service Aids* manual.

WebSphere MQ for z/OS recovery routines write information in the *system diagnostic work area* (SDWA) to the SYS1.LOGREC data set when retry is attempted, or when percolation to the next recovery routine occurs. Multiple SYS1.LOGREC entries can be recorded, because two or more retries or percolations might occur for a single error.

The SYS1.LOGREC entries recorded near the time of abend might provide valuable historical information about the events leading up to the abend.

Finding the applicable SYS1.LOGREC information

To obtain a SYS1.LOGREC listing, either:

- Use the EREP service aid, described in the *MVS Diagnosis: Tools and Service Aids* manual to format records in the SYS1.LOGREC data set.
- Specify the VERBEXIT LOGDATA keyword in IPCS.
- Use option 7 on the DUMP ANALYSIS MENU (refer to “Processing a dump using the WebSphere MQ for z/OS dump display panels” on page 1279).

Only records available in storage when the dump was requested are included. Each formatted record follows the heading *****LOGDATA*****.

SVC dumps

Use this topic to understand how to suppress SVC dumps, and reasons why SVC dumps are not produced.

When SVC dumps are not produced

Under some circumstances, SVC dumps are not produced. Generally, dumps are suppressed because of time or space problems, or security violations. The following list summarizes other reasons why SVC dumps might not be produced:

- The z/OS *serviceability level indication processing* (SLIP) commands suppressed the abend.
The description of IEACMD00 in the *MVS Initialization and Tuning Reference* manual lists the defaults for SLIP commands executed at IPL time.
- The abend reason code was one that does not require a dump to determine the cause of abend.
- SDWACOMU or SDWAEAS (part of the system diagnostic work area, SDWA) was used to suppress the dump.

Suppressing WebSphere MQ for z/OS dumps using z/OS DAE

You can suppress SVC dumps that duplicate previous dumps. The *MVS Diagnosis: Tools and Service Aids* manual gives details about using z/OS *dump analysis and elimination* (DAE).

To support DAE, WebSphere MQ for z/OS defines two *variable recording area* (VRA) keys and a minimum symptom string. The two VRA keys are:

- KEY VRADAE (X'53'). No data is associated with this key.

- KEY VRAMINSC (X'52') DATA (X'08')

WebSphere MQ for z/OS provides the following data for the minimum symptom string in the *system diagnostic work area* (SDWA):

- Load module name
- CSECT name
- Abend code
- Recovery routine name
- Failing instruction area
- REG/PSW difference
- Reason code
- Component identifier
- Component subfunction

Dumps are considered duplicates for the purpose of suppressing duplicate dumps if eight (the X'08' from the VRAMINSC key) of the nine symptoms are the same.

Dealing with performance problems on z/OS

Use this topic to investigate performance problems in more detail.

Performance problems are characterized by the following:

- Poor response times in online transactions
- Batch jobs taking a long time to complete
- The transmission of messages is slow

Performance problems can be caused by many factors, from a lack of resource in the z/OS system as a whole, to poor application design.

The following topics present problems and suggested solutions, starting with problems that are relatively simple to diagnose, such as DASD contention, through problems with specific subsystems, such as IBM WebSphere MQ and CICS or IMS.

- “IBM WebSphere MQ for z/OS system considerations”
- “CICS constraints” on page 1294
- “Dealing with applications that are running slowly or have stopped on z/OS” on page 1294

Remote queuing problems can be due to network congestion and other network problems. They can also be caused by problems at the remote queue manager.

Related concepts:

“Dealing with problems” on page 1186

“Dealing with incorrect output” on page 1299

IBM WebSphere MQ for z/OS system considerations

The z/OS system is an area that requires examination when investigating performance problems.

You might already be aware that your z/OS system is under stress because these problems affect many subsystems and applications.

You can use the standard monitoring tools such as Resource Monitoring Facility (RMF) to monitor and diagnose these problems. They might include:

- Constraints on storage (paging)
- Constraints on processor cycles

- Constraints on DASD
- Channel path usage

Use normal z/OS tuning techniques to resolve these problems.

CICS constraints

CICS constraints can also have an adverse effect on WebSphere MQ performance. Use this topic for further information about CICS constraints.

Performance of WebSphere MQ tasks can be affected by CICS constraints. For example, your system might have reached MAXTASK, forcing transactions to wait, or the CICS system might be short on storage. For example, CICS might not be scheduling transactions because the number of concurrent tasks has been reached, or CICS has detected a resource problem. If you suspect that CICS is causing your performance problems (for example because batch and TSO jobs run successfully, but your CICS tasks time out, or have poor response times), see the *CICS Problem Determination Guide* and the *CICS Performance Guide*.

Note: CICS I/O to transient data extrapartition data sets uses the z/OS RESERVE command. This could affect I/O to other data sets on the same volume.

Dealing with applications that are running slowly or have stopped on z/OS

Waits and loops can exhibit similar symptoms. Use the links in this topic to help differentiate between waits and loops.

Waits and loops are characterized by unresponsiveness. However, it can be difficult to distinguish between waits, loops, and poor performance.

Any of the following symptoms might be caused by a wait or a loop, or by a badly tuned or overloaded system:

- An application that appears to have stopped running (if IBM WebSphere MQ for z/OS is still responsive, this problem is probably caused by an application problem)
- An MQSC command that does not produce a response
- Excessive use of processor time

To perform the tests shown in these topics, you need access to the z/OS console, and to be able to issue operator commands.

- “Distinguishing between waits and loops”
- “Dealing with waits” on page 1296
- “Dealing with loops” on page 1298

Related concepts:

“Dealing with problems” on page 1186

Distinguishing between waits and loops:

Waits and loops on WebSphere MQ can present similar symptoms. Use this topic to help determine if you are experiencing a wait or a loop.

Because waits and loops can be difficult to distinguish, in some cases you need to carry out a detailed investigation before deciding which classification is right for your problem.

This section gives you guidance about choosing the best classification, and advice on what to do when you have decided on a classification.

Waits

For problem determination, a wait state is regarded as the state in which the execution of a task has been suspended. That is, the task has started to run, but has been suspended without completing, and has subsequently been unable to resume.

A problem identified as a wait in your system could be caused by any of the following:

- A wait on an MQI call
- A wait on a CICS or IMS call
- A wait for another resource (for example, file I/O)
- An ECB wait
- The CICS or IMS region waiting
- TSO waiting
- WebSphere MQ for z/OS waiting for work
- An apparent wait, caused by a loop
- Your task is not being dispatched by CICS or MVS due to higher priority work
- Db2 or RRS are inactive

Loops

A loop is the repeated execution of some code. If you have not planned the loop, or if you have designed it into your application but it does not terminate for some reason, you get a set of symptoms that vary depending on what the code is doing, and how any interfacing components and products react to it. In some cases, at first, a loop might be diagnosed as a wait or performance problem, because the looping task competes for system resources with other tasks that are not involved in the loop. However, a loop consumes resources but a wait does not.

An apparent loop problem in your system could be caused by any of the following:

- An application doing a lot more processing than usual and therefore taking much longer to complete
- A loop in application logic
- A loop with MQI calls
- A loop with CICS or IMS calls
- A loop in CICS or IMS code
- A loop in WebSphere MQ for z/OS

Symptoms of waits and loops

Any of the following symptoms could be caused by a wait, a loop, or by a badly tuned or overloaded system:

- Timeouts on MQGET WAITs
- Batch jobs suspended
- TSO session suspended
- CICS task suspended
- Transactions not being started because of resource constraints, for example CICS MAX task
- Queues becoming full, and not being processed
- System commands not accepted, or producing no response

Related concepts:

“Dealing with waits”

“Dealing with loops” on page 1298

Dealing with waits:

Waits can occur in batch or TSO applications, CICS transactions, and other components. Use this topic to determine where waits can occur.

When investigating what appears to be a problem with tasks or subsystems waiting, it is necessary to take into account the environment in which the task or subsystem is running.

It might be that your z/OS system is generally under stress. In this case, there can be many symptoms. If there is not enough real storage, jobs experience waits at paging interrupts or swap-outs. Input/output (I/O) contention or high channel usage can also cause waits.

You can use standard monitoring tools, such as *Resource Monitoring Facility (RMF)* to diagnose such problems. Use normal z/OS tuning techniques to resolve them.

Is a batch or TSO program waiting?

Consider the following points:

Your program might be waiting on another resource

For example, a VSAM control interval (CI) that another program is holding for update.

Your program might be waiting for a message that has not yet arrived

This condition might be normal behavior if, for example, it is a server program that constantly monitors a queue.

Alternatively, your program might be waiting for a message that has arrived, but has not yet been committed.

Issue the DIS CONN(*) TYPE(HANDLE) command and examine the queues in use by your program.

If you suspect that your program has issued an MQI call that did not involve an MQGET WAIT, and control has not returned from WebSphere MQ, take an SVC dump of both the batch or TSO job, and the WebSphere MQ subsystem before canceling the batch or TSO program.

Also consider that the wait state might be the result of a problem with another program, such as an abnormal termination (see “Messages do not arrive when expected” on page 1300), or in WebSphere MQ itself (see “Is WebSphere MQ waiting for z/OS?” on page 1297). Refer to “WebSphere MQ dumps” on page 1276 (specifically Figure 155 on page 1278) for information about obtaining a dump.

If the problem persists, refer to “Searching the IBM database for similar problems, and solutions” on page 1327 and “Contacting IBM Software Support” on page 1346 for information about reporting the problem to IBM.

Is a CICS transaction waiting?

Consider the following points:

CICS might be under stress

This might indicate that the maximum number of tasks allowed (MAXTASK) has been reached, or a short on storage (SOS) condition exists. Check the console log for messages that might explain this (for example, SOS messages), or see the *CICS Problem Determination Guide*.

The transaction might be waiting for another resource

For example, this might be file I/O. You can use CEMT INQ TASK to see what the task is waiting for. If the resource type is MQSERIES your transaction is waiting on WebSphere MQ (either in an **MQGET WAIT** or a task switch). Otherwise see the *CICS Problem Determination Guide* to determine the reason for the wait.

The transaction might be waiting for WebSphere MQ for z/OS

This might be normal, for example, if your program is a server program that waits for messages to arrive on a queue. Otherwise it might be the result of a transaction abend, for example (see “Messages do not arrive when expected” on page 1300). If so, the abend is reported in the CSMT log.

The transaction might be waiting for a remote message

If you are using distributed queuing, the program might be waiting for a message that has not yet been delivered from a remote system (for further information, refer to “Problems with missing messages when using distributed queuing” on page 1301).

If you suspect that your program has issued an MQI call that did not involve an **MQGET WAIT** (that is, it is in a task switch), and control has not returned from WebSphere MQ, take an SVC dump of both the CICS region, and the WebSphere MQ subsystem before canceling the CICS transaction. Refer to “Dealing with loops” on page 1298 for information about waits. Refer to “WebSphere MQ dumps” on page 1276 (specifically Figure 155 on page 1278) for information about obtaining a dump.

If the problem persists, refer to “Searching the IBM database for similar problems, and solutions” on page 1327 and “Contacting IBM Software Support” on page 1346 for information about reporting the problem to IBM.

Is Db2 waiting?

If your investigations indicate that Db2 is waiting, check the following:

1. Use the Db2 **-DISPLAY THREAD(*)** command to determine if any activity is taking place between the queue manager and the Db2 subsystem.
2. Try and determine whether any waits are local to the queue manager subsystems or are across the Db2 subsystems.

Is RRS active?

- Use the **D RRS** command to determine if RRS is active.

Is WebSphere MQ waiting for z/OS?

If your investigations indicate that WebSphere MQ itself is waiting, check the following:

1. Use the **DISPLAY THREAD(*)** command to check if anything is connected to WebSphere MQ.
2. Use **SDSF DA**, or the z/OS command **DISPLAY A,xxxxMSTR** to determine whether there is any processor usage (as shown in “Has your application or WebSphere MQ for z/OS stopped processing work?” on page 1174).
 - If WebSphere MQ is using some processor time, reconsider other reasons why WebSphere MQ might be waiting, or consider whether this is actually a performance problem.
 - If there is no processor activity, check whether WebSphere MQ responds to commands. If you can get a response, reconsider other reasons why WebSphere MQ might be waiting.
 - If you cannot get a response, check the console log for messages that might explain the wait (for example, WebSphere MQ might have run out of active log data sets, and be waiting for offload processing).

If you are satisfied that WebSphere MQ has stalled, use the **STOP QMGR** command in both **QUIESCE** and **FORCE** mode to terminate any programs currently being executed.

If the STOP QMGR command fails to respond, cancel the queue manager with a dump, and restart. If the problem recurs, refer to “Contacting IBM Software Support” on page 1346 for further guidance.

Related concepts:

“Distinguishing between waits and loops” on page 1294

“Dealing with loops”

Dealing with loops:

Loops can occur in different areas of a z/OS system. Use this topic to help determine where a loop is occurring.

The following topics describe the various types of loop that you might encounter, and suggest some responses.

Is a batch application looping?

If you suspect that a batch or TSO application is looping, use the console to issue the z/OS command DISPLAY JOBS,A (for a batch application) or DISPLAY TS,A (for a TSO application). Note the CT values from the data displayed, and repeat the command.

If any task shows a significant increase in the CT value, it might be that the task is looping. You could also use SDSF DA, which shows you the percentage of processor that each address space is using.

Is a batch job producing a large amount of output?

An example of this behavior might be an application that browses a queue and prints the messages. If the browse operation has been started with BROWSE FIRST, and subsequent calls have not been reset to BROWSE NEXT, the application browses, and prints the first message on the queue repeatedly.

You can use SDSF DA to look at the output of running jobs if you suspect that it might be causing a problem.

Does a CICS region show heavy processor activity?

It might be that a CICS application is looping, or that the CICS region itself is in a loop. You might see AICA abends if a transaction goes into a tight (unyielding) loop.

If you suspect that CICS, or a CICS application is looping, see the *CICS Problem Determination Guide*.

Does an IMS region show heavy processor activity?

It might be that an IMS application is looping. If you suspect this behavior, see *IMS Diagnosis Guide and Referencel*.

Is the queue manager showing heavy processor activity?

Try to enter an MQSC DISPLAY command from the console. If you get no response, it is possible that the queue manager is looping. Follow the procedure shown in “Has your application or WebSphere MQ for z/OS stopped processing work?” on page 1174 to display information about the processor time being used by the queue manager. If this command indicates that the queue manager is in a loop, take a memory dump, cancel the queue manager and restart.

If the problem persists, see “Searching the IBM database for similar problems, and solutions” on page 1327 and “Contacting IBM Software Support” on page 1346 for information about reporting the problem to IBM.

Is a queue, page set, or Coupling Facility structure filling up unexpectedly?

If so, it might indicate that an application is looping, and putting messages on to a queue. (It might be a batch, CICS, or TSO application.)

Identifying a looping application

In a busy system, it might be difficult to identify which application is causing the problem. If you keep a cross-reference of applications to queues, terminate any programs or transactions that might be putting messages on to the queue. Investigate these programs or transactions before using them again. (The most likely culprits are new, or changed applications; check your change log to identify them.)

Try issuing a DISPLAY QSTATUS command on the queue. This command returns information about the queue that might help to identify which application is looping.

Incorrect triggering definitions

It might be that a getting application has not been triggered because of incorrect object definitions, for example, the queue might be set to NOTRIGGER.

Distributed queuing

Using distributed queuing, a symptom of this problem might be a message in the receiving system indicating that MQPUT calls to the dead-letter queue are failing. This problem might be caused because the dead-letter queue has also filled up. The dead-letter queue message header (dead-letter header structure) contains a reason or feedback code explaining why the message

might not be put on to the target queue. See  MQDLH – Dead-letter header (*WebSphere MQ V7.1 Reference*) for information about the dead-letter header structure.

Allocation of queues to page sets

If a particular page set frequently fills up, there might be a problem with the allocation of queues to page sets. See “IBM WebSphere MQ for z/OS performance constraints” on page 1266 for more information.

Shared queues

Is the Coupling Facility structure full? The z/OS command DISPLAY CF displays information about Coupling Facility storage including the total amount, the total in use, and the total free control and non-control storage. The RMF Coupling Facility Usage Summary Report provides a more permanent copy of this information.

Are a task, and WebSphere MQ for z/OS, showing heavy processor activity?

In this case, a task might be looping on MQI calls (for example, browsing the same message repeatedly).

Related concepts:

“Distinguishing between waits and loops” on page 1294

“Dealing with waits” on page 1296

Dealing with incorrect output

Incorrect output can be missing, unexpected, or corrupted information. Read this topic to investigate further.

The term “incorrect output” can be interpreted in many different ways, and its meaning for problem determination with this product documentation is explained in “Have you obtained incorrect output?” on page 1183.

The following topics contains information about the problems that you could encounter with your system and classify as incorrect output:

- Application messages that do not arrive when you are expecting them

- Application messages that contain the wrong information, or information that has been corrupted
- Additional problems that you might encounter if your application uses distributed queues are also described.

- “Messages do not arrive when expected”
- “Problems with missing messages when using distributed queuing” on page 1301
- “Problems with getting messages when using message grouping” on page 1303
- “Finding messages sent to a cluster queue” on page 1303
- “Finding messages sent to the IBM WebSphere MQ-IMS bridge” on page 1303
- “Messages contain unexpected or corrupted information” on page 1304

Related concepts:

“Dealing with problems” on page 1186

“Dealing with performance problems on z/OS” on page 1293

Messages do not arrive when expected

Missing messages can have different causes. Use this topic to investigate the causes further.

If messages do not arrive on the queue when you are expecting them, check for the following:

Has the message been put onto the queue successfully?

Did WebSphere MQ issue a return and reason code for the MQPUT, for example:

- Has the queue been defined correctly, for example is MAXMSGL large enough? (reason code 2030).
- Can applications put messages on to the queue (is the queue enabled for MQPUT calls)? (reason code 2051).
- Is the queue already full? This could mean that an application could not put the required message on to the queue (reason code 2053).

Is the queue a shared queue?

- Have Coupling Facility structures been defined successfully in the CFRM policy data set? Messages held on shared queues are stored inside a Coupling Facility.
- Have you activated the CFRM policy?

Is the queue a cluster queue?

If it is, there might be multiple instances of the queue on different queue managers. This means that the messages could be on a different queue manager.

- Did you want the message to go to a cluster queue?
- Is your application designed to work with cluster queues?
- Did the message get put to a different instance of the queue from that expected?

Check any cluster-workload exit programs to see that they are processing messages as intended.

Do your gets fail?

- Does the application need to take a syncpoint?

If messages are being put or got within syncpoint, they are not available to other tasks until the unit of recovery has been committed.

- Is the time interval on the MQGET long enough?

If you are using distributed processing, you should allow for reasonable network delays, or problems at the remote end.

- Was the message you are expecting defined as persistent?

If not, and the queue manager has been restarted, the message will have been deleted. Shared queues are an exception because nonpersistent messages survive a queue manager restart.

- Are you waiting for a specific message that is identified by a message or correlation identifier (*MsgId* or *CorrelId*)?
Check that you are waiting for a message with the correct *MsgId* or *CorrelId*. A successful MQGET call sets both these values to that of the message got, so you might need to reset these values to get another message successfully.
Also check if you can get other messages from the queue.
- Can other applications get messages from the queue?
If so, has another application already retrieved the message?
If the queue is a shared queue, check that applications on other queue managers are not getting the messages.

If you cannot find anything wrong with the queue, and the queue manager itself is running, make the following checks on the process that you expected to put the message on to the queue:

- Did the application get started?
If it should have been triggered, check that the correct trigger options were specified.
- Is a trigger monitor running?
- Was the trigger process defined correctly (both to WebSphere MQ for z/OS and CICS or IMS)?
- Did it complete correctly?
Look for evidence of an abend, for example, in the CICS log.
- Did the application commit its changes, or were they backed out?
Look for messages in the CICS log indicating this.

If multiple transactions are serving the queue, they might occasionally conflict with one another. For example, one transaction might issue an MQGET call with a buffer length of zero to find out the length of the message, and then issue a specific MQGET call specifying the *MsgId* of that message. However, while this is happening, another transaction might have issued a successful MQGET call for that message, so the first application receives a completion code of MQRC_NO_MSG_AVAILABLE. Applications that are expected to run in a multi-server environment must be designed to cope with this situation.

Have any of your systems suffered an outage? For example, if the message you were expecting should have been put on to the queue by a CICS application, and the CICS system went down, the message might be in doubt. This means that the queue manager does not know whether the message should be committed or backed out, and so has locked it until this is resolved when resynchronization takes place.

Note: The message is deleted after resynchronization if CICS decides to back it out.

Also consider that the message could have been received, but that your application failed to process it in some way. For example, did an error in the expected format of the message cause your program to reject it? If so, refer to “Messages contain unexpected or corrupted information” on page 1304.

Problems with missing messages when using distributed queuing

Use this topic to understand possible causes of missing messages when using distributed queuing.

If your application uses distributed queuing, consider the following points:

Has distributed queuing been correctly installed on both the sending and receiving systems?

Ensure that the instructions about installing the distributed queue management facility in



Configuring z/OS (*WebSphere MQ V7.1 Installing Guide*) have been followed correctly.

Are the links available between the two systems?

Check that both systems are available, and connected to IBM WebSphere MQ for z/OS. Check that the LU 6.2 or TCP/IP connection between the two systems is active or check the connection definitions on any other systems that you are communicating with.

See “Monitoring and performance” on page 829 for more information about trace-route messaging in a network.

Is the channel running?

- Issue the following command for the transmission queue:
DISPLAY QUEUE (qname) IPPROCS
If the value for IPPROCS is 0, this means that the channel serving this transmission queue is not running.
- Issue the following command for the channel:
DISPLAY CHSTATUS (channel-name) STATUS MSGS
Use the output produced by this command to check that the channel is serving the correct transmission queue and that it is connected to the correct target machine and port. You can determine whether the channel is running from the STATUS field. You can also see if any messages have been sent on the channel by examining the MSGS field.
If the channel is in RETRYING state, this is probably caused by a problem at the other end. Check that the channel initiator and listener have been started, and that the channel has not been stopped. If somebody has stopped the channel, you need to start it manually.

Is triggering set on in the sending system?

Check that the channel initiator is running.

Does the transmission queue have triggering set on?

If a channel is stopped under specific circumstances, triggering can be set off for the transmission queue.

Is the message you are waiting for a reply message from a remote system?

Check the definitions of the remote system, as previously described, and check that triggering is activated in the remote system. Also check that the LU 6.2 connection between the two systems is not single session (if it is, you cannot receive reply messages).

Check that the queue on the remote queue manager exists, is not full, and accepts the message length. If any of these criteria are not fulfilled, the remote queue manager tries to put the message on the dead-letter queue. If the message length is longer than the maximum length that the channel permits, the sending queue manager tries to put the message on its dead-letter queue.

Is the queue already full?

This could mean that an application could not put the required message on to the queue. If this is so, check if the message has been put on to the dead-letter queue.

The dead-letter queue message header (dead-letter header structure) contains a reason or feedback code explaining why the message could not be put on to the target queue. See

 MQDLH – Dead-letter header (*WebSphere MQ V7.1 Reference*) for more information about the dead-letter header structure.

Is there a mismatch between the sending and receiving queue managers?

For example, the message length could be longer than the receiving queue manager can handle. Check the console log for error messages.

Are the channel definitions of the sending and receiving channels compatible?

For example, a mismatch in the wrap value of the sequence number stops the channel. See

 Concepts of intercommunication (*WebSphere MQ V7.1 Product Overview Guide*) for more information about distributed queuing.

Has data conversion been performed correctly?

If a message has come from a different queue manager, are the CCSIDs and encoding the same, or does data conversion need to be performed.

Has your channel been defined for fast delivery of nonpersistent messages?

If your channel has been defined with the NPMSPEED attribute set to FAST (the default), and the channel has stopped for some reason and then been restarted, nonpersistent messages might have been lost. See  Nonpersistent message speed (NPMSPEED) (*WebSphere MQ V7.1 Reference*) for more information about fast messages.

Is a channel exit causing the messages to be processed in an unexpected way?

For example, a security exit might prevent a channel from starting, or an *ExitResponse* of MQXCC_CLOSE_CHANNEL might terminate a channel.

Problems with getting messages when using message grouping

Use this topic to understand some of the issues with getting messages when using message grouping.

Is the application waiting for a complete group of messages?

Ensure all the messages in the group are on the queue. If you are using distributed queuing, see “Problems with missing messages when using distributed queuing” on page 1301. Ensure the last message in the group has the appropriate *MsgFlags* set in the message descriptor to indicate that it is the last message. Ensure the message expiry of the messages in the group is set to a long enough interval that they do not expire before they are retrieved.

If messages from the group have already been retrieved, and the get request is not in logical order, turn off the option to wait for a complete group when retrieving the other group messages.

If the application issues a get request in logical order for a complete group, and midway through retrieving the group it cannot find a message:

Ensure that no other applications are running against the queue and getting messages. Ensure that the message expiry of the messages in the group is set to a long enough interval that they do not expire before they are retrieved. Ensure that no one has issued the CLEAR QUEUE command. You can retrieve incomplete groups from a queue by getting the messages by group ID, without specifying the logical order option.

Finding messages sent to a cluster queue

Use this topic to understand some of the issues involved with finding messages sent to a cluster queue.

Before you can use the techniques described in these topics to find a message that did not arrive at a cluster queue, you need to determine the queue managers that host the queue to which the message was sent. You can determine this in the following ways:

- You can use the DISPLAY QUEUE command to request information about cluster queues.
- You can use the name of the queue and queue manager that is returned in the MQPMO structure.

If you specified the MQOO_BIND_ON_OPEN option for the message, these fields give the destination of the message. If the message was not bound to a particular queue and queue manager, these fields give the name of the first queue and queue manager to which the message was sent. In this case, it might not be the ultimate destination of the message.

Finding messages sent to the IBM WebSphere MQ-IMS bridge

Use this topic to understand possible causes for missing messages sent to the IBM WebSphere MQ-IMS bridge.

If you are using the IBM WebSphere MQ-IMS bridge, and your message has not arrived as expected, consider the following:

Is the IBM WebSphere MQ-IMS bridge running?

Issue the following command for the bridge queue:

```
DISPLAY QSTATUS(qname) IPPROCS CURDEPTH
```

The value of IPPROCS should be 1; if it is 0, check the following:

- Is the queue a bridge queue?
- Is IMS running?
- Has OTMA been started?
- Is IBM WebSphere MQ connected to OTMA?

Note: There are two IBM WebSphere MQ messages that you can use to establish whether you have a connection to OTMA. If message CSQ2010I is present in the job log of the task, but message CSQ2011I is not present, IBM WebSphere MQ is connected to OTMA. This message also tells you to which IBM WebSphere MQ system OTMA is connected. For more information

about the content of these messages, see  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*).

Within the queue manager there is a task processing each IMS bridge queue. This task gets from the queue, sends the request to IMS, and then does a commit. If persistent messages are used, then the commit requires disk I/O and so the process takes longer than for non-persistent messages. The time to process the get, send, and commit, limits the rate at which the task can process messages. If the task can keep up with the workload then the current depth is close to zero. If you find that the current depth is often greater than zero you might be able to increase throughput by using two queues instead of one.

Use the IMS command `/DIS OTMA` to check that OTMA is active.

If your messages are flowing to IMS, check the following:

- Use the IMS command `/DIS TMEMBER client TPIPE ALL` to display information about IMS Tpipes. From this you can determine the number of messages enqueued on, and dequeued from, each Tpipe. (Commit mode 1 messages are not usually queued on a Tpipe.)
- Use the IMS command `/DIS A` to show whether there is a dependent region available for the IMS transaction to run in.
- Use the IMS command `/DIS TRAN trancode` to show the number of messages queued for a transaction.
- Use the IMS command `/DIS PROG progname` to show if a program has been stopped.

Was the reply message sent to the correct place?

Issue the following command:

```
DISPLAY QSTATUS(*) CURDEPTH
```

Does the CURDEPTH indicate that there is a reply on a queue that you are not expecting?

Messages contain unexpected or corrupted information

Use this topic to understand some of the issues that can cause unexpected or corrupted output.

If the information contained in the message is not what your application was expecting, or has been corrupted in some way, consider the following points:

Has your application, or the application that put the message on to the queue changed?

Ensure that all changes are simultaneously reflected on all systems that need to be aware of the change.

For example, a copybook formatting the message might have been changed, in which case, both applications have to be recompiled to pick up the changes. If one application has not been recompiled, the data will appear corrupt to the other.

Check that no external source of data, such as a VSAM data set, has changed. This could also invalidate your data if any necessary recompilations have not been done. Also check that any CICS maps and TSO panels that you are using for input of message data have not changed.

Is an application sending messages to the wrong queue?

Check that the messages your application is receiving are not intended for an application servicing a different queue. If necessary, change your security definitions to prevent unauthorized applications from putting messages on to the wrong queues.

If your application has used an alias queue, check that the alias points to the correct queue.

If you altered the queue to make it a cluster queue, it might now contain messages from different application sources.

Has the trigger information been specified correctly for this queue?

Check that your application should have been started, or should a different application have been started?

Has data conversion been performed correctly?

If a message has come from a different queue manager, are the CCSIDs and encoding the same, or does data conversion need to be performed.

Check that the *Format* field of the MQMD structure corresponds with the content of the message. If not, the data conversion process might not have been able to deal with the message correctly.

If these checks do not enable you to solve the problem, check your application logic, both for the program sending the message, and for the program receiving it.

Resolving problems with channels and DQM

Use the advice given in the links provided to help you to resolve common problems that can arise when you use channels and distributed queue management.

IBM WebSphere MQ provides a utility to assist with problem determination named **amqldmpa**. During the course of problem determination, your IBM service representative might ask you to provide output from the utility.

Your IBM service representative will provide you with the parameters you require to collect the appropriate diagnostic information, and information on how you send the data you record to IBM.

Attention: You should not rely on the format of the output from this utility, as the format is subject to change without notice.

A number of validation checks are made when creating, altering, and deleting channels, and where appropriate, an error message returned. Errors might occur when the following events happen:

- A duplicate channel name is chosen when creating a channel
- Unacceptable data is entered in the channel parameter fields
- The channel to be altered is in doubt, or does not exist

Problems found during normal operation of the channels are reported to the system console and to the system log. In IBM WebSphere MQ for Windows, problems are recorded in the channel log. To diagnose a problem you must collect the relevant information from the log, and analyze this information to identify the problem. Analysis can be difficult in a network where the problem might arise at an intermediate system that is staging some of your messages. An error situation, such as a full transmission queue, followed by the dead-letter queue filling up, would result in your channel to that site closing down. In that example, the error message you receive in your error log indicates a problem originating from the remote site, but might not be able to tell you any details about the error at that site. You need to contact your counterpart at the remote site to obtain details of the problem, and to receive notification of that channel becoming available again

Use the following links to help you resolve your channel problems.

- “Using Ping to check the communication link” on page 1307

- “Channel problems and the dead-letter queue” on page 1307
-  In-doubt channels (*WebSphere MQ V7.1 Installing Guide*)
- “Channel startup negotiation errors” on page 1308
- “When a channel does not run” on page 1308
- “Retrying the link” on page 1311
- “Data structures” on page 1312
- “User exit problems” on page 1312
- “Disaster recovery” on page 1312
- “Channel switching” on page 1313
- “Connection switching” on page 1313
- “Client problems” on page 1313
- “Message monitoring” on page 1315

Related concepts:

-  Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*)
- “Troubleshooting and support” on page 1149
- “Making initial checks on Windows, UNIX and Linux systems” on page 1151
- “Making initial checks on z/OS” on page 1169
- “Making initial checks on IBM i” on page 1161
- “Reason codes” on page 1379

Related reference:

-  Communications protocol return codes (*WebSphere MQ V7.1 Reference*)

Error message from channel control

Problems found during normal operation of the channels are reported to the system console and to the system log. In WebSphere MQ for Windows, they are reported to the channel log.

Problem diagnosis starts with the collection of all relevant information from the log, and analysis of this information to identify the problem.

However, this analysis could be difficult in a network where the problem might arise at an intermediate system that is staging some of your messages. An error situation, such as transmission queue full, followed by the dead-letter queue filling up, would result in your channel to that site closing down.

In this example, the error message you receive in your error log indicates a problem originating from the remote site, but might not be able to tell you any details about the error at that site.

You need to contact your counterpart at the remote site to obtain details of the problem, and to receive notification of that channel becoming available again.

Using Ping to check the communication link

Ping is useful in determining whether the communication link and the two message channel agents that make up a message channel are functioning across all interfaces.

Ping does not use transmission queues, but it does invoke some user exit programs. If any error conditions are encountered, error messages are issued.

To use ping, you can issue the MQSC command PING CHANNEL. On z/OS and IBM i, you can also use the panel interface to select this option.

On IBM i, Windows, UNIX, and Linux platforms, you can also use the MQSC command PING QMGR to test whether the queue manager is responsive to commands. See the  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for more information about this topic.

Channel problems and the dead-letter queue

If a channel ceases to run for any reason, it might direct messages to a dead-letter queue. If the dead-letter queue cannot take the messages, a number of things occur.

In some WebSphere MQ implementations, the dead-letter queue is referred to as an *undelivered-message queue*.

A channel attribute, `USEDLQ`, can be used to determine whether the dead-letter queue is used when messages cannot be delivered by channels for what ever reason.

If a channel ceases to run for any reason, applications probably continue to place messages on transmission queues, creating a potential overflow situation. Applications can monitor transmission queues to find the number of messages waiting to be sent, but this monitoring would not be a normal function for them to carry out.

When this occurs in a message-originating node, and the local transmission queue is full, the PUT of the application fails.

When this occurs in a staging or destination node, there are three ways that the MCA copes with the situation:

1. By calling the message-retry exit, if one is defined.
2. By directing all overflow messages to a *dead-letter queue (DLQ)*, returning an exception report to applications that requested these reports.

Note: In distributed-queuing management, if the message is too large for the DLQ, the DLQ is full, or the DLQ is not available, the channel stops and the message remains on the transmission queue. Ensure your DLQ is defined, available, and sized for the largest messages you handle.

3. By closing down the channel, if neither of the previous options succeeded.
4. By returning the undelivered messages back to the sending end and returning a full report to the reply-to-queue (`MQRC_EXCEPTION_WITH_FULL_DATA` and `MQRO_DISCARD_MSG`).

If an MCA is unable to put a message on the DLQ:

- The channel stops
- Appropriate error messages are issued at the system consoles at both ends of the message channel
- The unit of work is backed out, and the messages reappear on the transmission queue at the sending channel end of the channel
- Triggering is disabled for the transmission queue

Validation checks

A number of validation checks are made when creating, altering, and deleting channels, and where appropriate, an error message returned.

Errors might occur when the following events happen:

- A duplicate channel name is chosen when creating a channel
- Unacceptable data is entered in the channel parameter fields
- The channel to be altered is in doubt, or does not exist

In-doubt relationship

If a channel is in doubt, it is typically resolved automatically on restart, so the system operator does not need to resolve a channel manually in normal circumstances.

See  In-doubt channels (*WebSphere MQ V7.1 Installing Guide*) for information about this topic.

Channel startup negotiation errors

During channel startup, the starting end has to state its position and agree channel running parameters with the corresponding channel.

It might happen that the two ends cannot agree on the parameters, in which case the channel closes down with error messages being issued to the appropriate error logs.

Shared channel recovery

The following table shows the types of shared-channel failure and how each type is handled.

Type of failure:	What happens:
Channel initiator communications subsystem failure	The channels dependent on the communications subsystem enter channel retry, and are restarted on an appropriate queue-sharing group channel initiator by a load-balanced start command.
Channel initiator failure	The channel initiator fails, but the associated queue manager remains active. The queue manager monitors the failure and initiates recovery processing.
Queue manager failure	The queue manager fails (failing the associated channel initiator). Other queue managers in the queue-sharing group monitor the event and initiate peer recovery.
Shared status failure	Channel state information is stored in Db2, so a loss of connectivity to Db2 becomes a failure when a channel state change occurs. Running channels can carry on running without access to these resources. On a failed access to Db2, the channel enters retry.

Shared channel recovery processing on behalf of a failed system requires connectivity to Db2 to be available on the system managing the recovery to retrieve the shared channel status.

When a channel does not run

If a channel persistently fails to start, there are a number of potential reasons.

Carry out the following checks:

- Check that DQM and the channels have been set up correctly. This issue is a likely problem source if the channel has never run. Reasons could be:
 - A mismatch of names between sending and receiving channels (remember that uppercase and lowercase letters are significant)
 - Incorrect channel types specified
 - The sequence number queue (if applicable) is not available, or is damaged

- The dead-letter queue is not available
- The sequence number wrap value is different on the two channel definitions
- A queue manager or communication link is not available
- A receiver channel might be in STOPPED state
- The connection might not be defined correctly
- There might be a problem with the communications software (for example, is TCP running?)
- It is possible that an in-doubt situation exists, if the automatic synchronization on startup has failed for some reason. This issue is indicated by messages on the system console, and the status panel can be used to show channels that are in doubt.

The possible responses to this situation are:

- Issue a Resolve channel request with Backout or Commit.

You need to check with your remote link supervisor to establish the number of the last message or unit of work committed. Check this number against the last number at your end of the link. If the remote end has committed a number, and that number is not yet committed at your end of the link, then issue a RESOLVE COMMIT command.

In all other cases, issue a RESOLVE BACKOUT command.

The effect of these commands is that backed out messages reappear on the transmission queue and are sent again, while committed messages are discarded.

If in doubt yourself, perhaps backing out with the probability of duplicating a sent message would be the safer decision.

- Issue a RESET CHANNEL command.

This command is for use when sequential numbering is in effect, and must be used with care. Its purpose is to reset the sequence number of messages and you must use it only after using the RESOLVE command to resolve any in-doubt situations.

- On WebSphere MQ for IBM i, z/OS, Windows, UNIX, and Linux systems, there is no need for the administrator to choose a particular sequence number to ensure that the sequence numbers are put back in step. When a sender channel starts after being reset, it informs the receiver that it has been reset and supplies the new sequence number that is to be used by both the sender and receiver.
- If the status of a receiver end of the channel is STOPPED, it can be reset by starting the receiver end.

Note: This action does not start the channel, it merely resets the status. The channel must still be started from the sender end.

Related concepts:

“Triggered channels”

“Conversion failure” on page 1310

“Network problems” on page 1310

“Dial-up problems” on page 1311



Connecting applications using distributed queuing (*WebSphere MQ V7.1 Installing Guide*)

Triggered channels

If a triggered channel refuses to run, investigate the possibility of in-doubt messages. Another possibility is that the trigger control parameter on the transmission queue has been set to NOTRIGGER by the channel.

If a triggered channel refuses to run, investigate the possibility of in-doubt messages here: “When a channel does not run” on page 1308.

Another possibility is that the trigger control parameter on the transmission queue has been set to NOTRIGGER by the channel. This issue happens when:

- There is a channel error.

- The channel was stopped because of a request from the receiver.
- The channel was stopped because of a problem on the sender that requires manual intervention.

After diagnosing and fixing the problem, start the channel manually.

An example of a situation where a triggered channel fails to start is as follows:

1. A transmission queue is defined with a trigger type of FIRST.
2. A message arrives on the transmission queue, and a trigger message is produced.
3. The channel is started, but stops immediately because the communications to the remote system are not available.
4. The remote system is made available.
5. Another message arrives on the transmission queue.
6. The second message does not increase the queue depth from zero to one, so no trigger message is produced (unless the channel is in RETRY state). If this issue happens, restart the channel manually.

On WebSphere MQ for z/OS, if the queue manager is stopped using MODE(FORCE) during channel initiator shutdown, it might be necessary to manually restart some channels after channel initiator restart.

Conversion failure

Another reason for the channel refusing to run could be that neither end is able to carry out necessary conversion of message descriptor data between ASCII and EBCDIC, and integer formats.

In this instance, communication is not possible.

Network problems

There are a number of things to check if you are experiencing network problems.

When using LU 6.2, make sure that your definitions are consistent throughout the network. For example, if you have increased the RU sizes in your CICS Transaction Server for z/OS or Communications Manager definitions, but you have a controller with a small MAXDATA value in its definition, the session might fail if you attempt to send large messages across the network. A symptom of this problem might be that channel negotiation takes place successfully, but the link fails when message transfer occurs.

When using TCP, if your channels are unreliable and your connections break, you can set a KEEPALIVE value for your system or channels. You do this using the SO_KEEPAIVE option to set a system-wide value, and on WebSphere MQ for z/OS, you can also use the KeepAlive Interval channel attribute (KAINT) to set channel-specific keepalive values. On WebSphere MQ for z/OS, you can alternatively use the RCVTIME and RCVTMIN channel initiator parameters. For more information about these parameters, see  [Checking that the other end of the channel is still available \(WebSphere MQ V7.1 Installing Guide\)](#), and  [Keepalive Interval \(KAINT\) \(WebSphere MQ V7.1 Reference\)](#).

Adopting an MCA:

The Adopt MCA function enables WebSphere MQ to cancel a receiver channel and to start a new one in its place.

For more information about this function, see  [Adopting an MCA \(WebSphere MQ V7.1 Installing Guide\)](#). For details of its parameters, see  [MQSC reference \(WebSphere MQ V7.1 Reference\)](#).

Registration time for DDNS:

When a group TCP/IP listener is started, it registers with DDNS. But there may be a delay until the address is available to the network. A channel that is started in this period, and which targets the newly registered generic name, fails with an 'error in communications configuration' message. The channel then goes into retry until the name becomes available to the network. The length of the delay will be dependent on the name server configuration used.

Dial-up problems

WebSphere MQ supports connection over dial-up lines but with TCP, some protocol providers assign a new IP address each time you dial in.

This new IP address can cause channel synchronization problems because the channel cannot recognize the new IP addresses and so cannot ensure the authenticity of the partner. If you encounter this problem, you need to use a security exit program to override the connection name for the session.

This problem does not occur when a WebSphere MQ for IBM i, Windows, UNIX or Linux systems product is communicating with another product at the same level, because the queue manager name is used for synchronization instead of the IP address.

Retrying the link

An error scenario might occur that is difficult to recognize. For example, the link and channel might be functioning perfectly, but some occurrence at the receiving end causes the receiver to stop.

Another unforeseen situation could be that the receiver system has run out of memory and is unable to complete a transaction.

You must be aware that such situations can arise, often characterized by a system that appears to be busy but is not actually moving messages. You must work with your counterpart at the far end of the link to help detect the problem and correct it.

Retry considerations

If a link failure occurs during normal operation, a sender or server channel program will itself start another instance, provided that:

1. Initial data negotiation and security exchanges are complete
2. The retry count in the channel definition is greater than zero

Note: For i5/OS, UNIX systems, and Windows, to attempt a retry a channel initiator must be running. In platforms other than WebSphere MQ for i5/OS, UNIX systems, and Windows systems, this channel initiator must be monitoring the initiation queue specified in the transmission queue that the channel is using.

Shared channel recovery on z/OS:

See "Shared channel recovery" on page 1308, for a table that shows the types of shared-channel failure and how each type is handled.

Data structures

Data structures are needed for reference when checking logs and trace entries during problem diagnosis.

More information can be found in  [Channel-exit calls and data structures \(WebSphere MQ V7.1 Reference\)](#) and  [Developing applications reference \(WebSphere MQ V7.1 Reference\)](#).

User exit problems

The interaction between the channel programs and the user-exit programs has some error-checking routines, but this facility can only work successfully when the user exits obey certain rules.

These rules are described in  [Channel-exit programs for messaging channels \(WebSphere MQ V7.1 Programming Guide\)](#). When errors occur, the most likely outcome is that the channel stops and the channel program issues an error message, together with any return codes from the user exit. Any errors detected on the user exit side of the interface can be determined by scanning the messages created by the user exit itself.

You might need to use a trace facility of your host system to identify the problem.

Disaster recovery

Disaster recovery planning is the responsibility of individual installations, and the functions performed might include the provision of regular system 'snapshot' dumps that are stored safely off-site.

These dumps would be available for regenerating the system, if some disaster were to overtake it. If this occurs, you need to know what to expect of the messages, and the following description is intended to start you thinking about it.

First a recap on system restart. If a system fails for any reason, it might have a system log that allows the applications running at the time of failure to be regenerated by replaying the system software from a sync point forward to the instant of failure. If this occurs without error, the worst that can happen is that message channel sync points to the adjacent system might fail on startup, and that the last batches of messages for the various channels are sent again. Persistent messages are recovered and sent again, nonpersistent messages might be lost.

If the system has no system log for recovery, or if the system recovery fails, or where the disaster recovery procedure is invoked, the channels and transmission queues might be recovered to an earlier state, and the messages held on local queues at the sending and receiving end of channels might be inconsistent.

Messages might have been lost that were put on local queues. The consequence of this happening depends on the particular WebSphere MQ implementation, and the channel attributes. For example, where strict message sequencing is in force, the receiving channel detects a sequence number gap, and the channel closes down for manual intervention. Recovery then depends upon application design, as in the worst case the sending application might need to restart from an earlier message sequence number.

Channel switching

A possible solution to the problem of a channel ceasing to run is to have two message channels defined for the same transmission queue, but with different communication links.

One message channel is preferred, the other is a replacement for use when the preferred channel is unavailable.

If triggering is required for these message channels, the associated process definitions must exist for each sender channel end.

To switch message channels:

- If the channel is triggered, set the transmission queue attribute NOTRIGGER.
- Ensure that the current channel is inactive.
- Resolve any in-doubt messages on the current channel.
- If the channel is triggered, change the process attribute in the transmission queue to name the process associated with the replacement channel.

In this context, some implementations allow a channel to have a blank process object definition, in which case you can omit this step as the queue manager finds and start the appropriate process object.

- Restart the channel, or if the channel was triggered, set the transmission queue attribute TRIGGER.

Connection switching

Another solution to the problem of a channel ceasing to run is to switch communication connections from the transmission queues.

Use the following steps to do the connection switching:

- If the sender channel is triggered, set the transmission queue attribute NOTRIGGER.
- Ensure that the channel is inactive.
- Change the connection and profile fields to connect to the replacement communication link.
- Ensure that the corresponding channel at the remote end has been defined.
- Restart the channel, or if the sender channel was triggered, set the transmission queue attribute TRIGGER.

Client problems

A client application might receive an unexpected error return code.

For example:

- Queue manager not available
- Queue manager name error
- Connection broken

Look in the client error log for a message explaining the cause of the failure. There might also be errors logged at the server, depending on the nature of the failure.

Terminating clients

Even though a client has terminated, it is still possible for its surrogate process to be holding its queues open. Normally this will only be for a short time until the communications layer notifies that the partner has gone.

Error logs

WebSphere MQ error messages are placed in different error logs depending on the platform. There are error logs for z/OS, Windows, UNIX and Linux systems.

Error logs for Windows

WebSphere MQ for Windows uses a number of error logs to capture messages concerning the operation of WebSphere MQ itself, any queue managers that you start, and error data coming from the channels that are in use.

The location the error logs are stored in depends on whether the queue manager name is known and whether the error is associated with a client.

- If the queue manager name is known and the queue manager is available:
`<install directory>\QMGRS\QMgrName\ERRORS\AMQERR01.LOG`
- If the queue manager is not available:
`<install directory>\QMGRS\@SYSTEM\ERRORS\AMQERR01.LOG`
- If an error has occurred with a client application:
`<install directory>\ERRORS\AMQERR01.LOG`

On Windows, you should also examine the Windows application event log for relevant messages.

Error logs on UNIX and Linux systems

WebSphere MQ on UNIX and Linux systems uses a number of error logs to capture messages concerning the operation of WebSphere MQ itself, any queue managers that you start, and error data coming from the channels that are in use. The location the error logs are stored in depends on whether the queue manager name is known and whether the error is associated with a client.

- If the queue manager name is known:
`/var/mqm/qmgrs/QMgrName/errors`
- If the queue manager name is not known (for example when there are problems in the listener or SSL handshake):
`/var/mqm/errors`

When a client is installed, and there is a problem in the client application, the following log is used:

- If an error has occurred with a client application:
`/var/mqm/errors/`

Error logs on z/OS

Error messages are written to:

- The z/OS system console
- The channel-initiator job log

If you are using the z/OS message processing facility to suppress messages, the console messages may be suppressed. See the  [Planning on z/OS \(WebSphere MQ V7.1 Installing Guide\)](#) for more information.

Message monitoring

If a message does not reach its intended destination, you can use the WebSphere MQ display route application, available through the control command **dspmqrte**, to determine the route a message takes through the queue manager network and its final location.

The WebSphere MQ display route application is described in “Message monitoring” on page 885.

First Failure Support Technology (FFST)

First Failure Support Technology (FFST) for WebSphere MQ provides information that can help IBM support personnel to diagnose a problem when a serious error occurs.

First Failure Data Capture (FFDC) provides an automated snapshot of the system environment when an unexpected internal error occurs. This snapshot is used by IBM support personnel to provide a better understanding of the state of the system and IBM WebSphere MQ when the problem occurred.

An FFST file is a file containing information for use in detecting and diagnosing software problems. In WebSphere MQ, FFST files have a file type of FDC.

Use the information in the following links to find out the names, locations and contents of FFST files in different platforms.

- “FFST: WebSphere MQ for Windows”
- “FFST: WebSphere MQ for UNIX and Linux systems” on page 1318
- “First Failure Support Technology (FFST)” on page 1320

Related concepts:

“Troubleshooting and support” on page 1149

“Troubleshooting overview” on page 1149

“Using logs” on page 1226

“Using trace” on page 1234

“Problem determination on z/OS” on page 1265

Related tasks:

“Searching knowledge bases” on page 1326

“Contacting IBM Software Support” on page 1346

FFST: WebSphere MQ for Windows

Describes the name, location, and contents of the First Failure Support Technology (FFST) files for Windows systems.

In WebSphere MQ for Windows, FFST information is recorded in a file in the `c:\Program Files\IBM\WebSphere MQ\errors` directory.

An FFST file contains one or more records. Each FFST record contains information about an error that is normally severe, and possibly unrecoverable. These records typically indicate either a configuration problem with the system or a WebSphere MQ internal error.

FFST files are named *AMQnnnnn.mm.FDC*, where:

<i>nnnnn</i>	Is the ID of the process reporting the error
<i>mm</i>	Starts at 0. If the full file name already exists, this value is incremented by one until a unique FFST file name is found. An FFST file name can already exist if a process is reused.

An instance of a process will write all FFST information to the same FFST file. If multiple errors occur during a single execution of the process, an FFST file can contain many records.

When a process writes an FFST record it also sends a record to the Event Log. The record contains the name of the FFST file to assist in automatic problem tracking. The Event log entry is made at the application level.

A typical FFST log is shown in Figure 165 on page 1317.

```

+-----+
| WebSphere MQ First Failure Symptom Report
| =====
|
| Date/Time           :- Mon January 28 2008 21:59:06 GMT
| UTC Time/Zone       :- 1201539869.892015 0 GMT
| Host Name           :- 99VXY09 (Windows XP Build 2600: Service Pack 1)
| PIDS                :- 5724H7200
| LVLS                :- 7.0.0.0
| Product Long Name   :- WebSphere MQ for Windows
| Vendor              :- IBM
| Probe Id            :- HL010004
| Application Name     :- MQM
| Component           :- hlgReserveLogSpace
| SCCS Info           :- lib/logger/amqhlge0.c, 1.26
| Line Number         :- 246
| Build Date          :- Jan 25 2008
| CMVC level          :- p000-L050202
| Build Type          :- IKAP - (Production)
| UserID              :- IBM_User
| Process Name        :- C:\Program Files\IBM\WebSphere MQ\bin\amqzlaa0.exe
| Process             :- 00003456
| Thread              :- 00000030
| QueueManager        :- qmgr2
| ConnId(1) IPCC      :- 162
| ConnId(2) QM        :- 45
| Major Errorcode     :- hrcE_LOG_FULL
| Minor Errorcode     :- OK
| Probe Type          :- MSGAMQ6709
| Probe Severity      :- 2
| Probe Description   :- AMQ6709: The log for the Queue manager is full.
| FDCSequenceNumber  :- 0
|
+-----+

MQM Function Stack
zlaMainThread
zlaProcessMessage
zlaProcessMQIRequest
zlaMQPUT
zsqMQPUT
kpiMQPUT
kqiPutIt
kqiPutMsgSegments
apiPutMessage
aqmPutMessage
aqhPutMessage
aqqWriteMsg
aqqWriteMsgData
aqlReservePutSpace
almReserveSpace
hlgReserveLogSpace
xcsFFST

MQM Trace History
-----} hlgReserveLogSpace rc=hrcW_LOG_GETTING_VERY_FULL
-----{ xllLongLockRequest
-----} xllLongLockRequest rc=OK

...

```

Figure 165. Sample WebSphere MQ for Windows First Failure Symptom Report

The Function Stack and Trace History are used by IBM to assist in problem determination. In many cases there is little that the system administrator can do when an FFST record is generated, apart from raising problems through the IBM Support Center.

In certain circumstances a small dump file can be generated in addition to an FFST file and placed in the `c:\Program Files\IBM\WebSphere MQ\errors` directory. A dump file will have the same name as the FFST file, in the form `AMQnnnnn.mm.dmp`. These files can be used by IBM to assist in problem determination.

First Failure Support Technology (FFST) files and Windows clients

The files are produced already formatted and are in the errors subdirectory of the WebSphere MQ MQI client installation directory.

These are normally severe, unrecoverable errors and indicate either a configuration problem with the system or a WebSphere MQ internal error.

The files are named `AMQnnnnn.mm.FDC`, where:

- `nnnnn` is the process ID reporting the error
- `mm` is a sequence number, normally 0

When a process creates an FFST it also sends a record to the system log. The record contains the name of the FFST file to assist in automatic problem tracking.

The system log entry is made at the “user.error” level.

First Failure Support Technology is explained in detail in [First Failure Support Technology \(FFST\)](#).

FFST: WebSphere MQ for UNIX and Linux systems

Describes the name, location, and contents of the First Failure Support Technology (FFST) files for UNIX and Linux systems.

For WebSphere MQ on UNIX and Linux systems, FFST information is recorded in a file in the `/var/mqm/errors` directory.

An FFST file contains one or more records. Each FFST record contains information about an error that is normally severe, and possibly unrecoverable. These records indicate either a configuration problem with the system or a WebSphere MQ internal error.

FFST files are named `AMQnnnnn.mm.FDC`, where:

<code>nnnnn</code>	Is the ID of the process reporting the error
<code>mm</code>	Starts at 0. If the full file name already exists, this value is incremented by one until a unique FFST file name is found. An FFST file name can already exist if a process is reused.

An instance of a process will write all FFST information to the same FFST file. If multiple errors occur during a single execution of the process, an FFST file can contain many records.

In order to read the contents of a FFST file, you must be either the creator of the file, or a member of the `mqm` group.

When a process writes an FFST record, it also sends a record to `syslog`. The record contains the name of the FFST file to assist in automatic problem tracking. The `syslog` entry is made at the `user.error` level. See the operating-system documentation about `syslog.conf` for information about configuring this.

Some typical FFST data is shown in [Figure 166](#) on page 1319.

```

+-----+
| WebSphere MQ First Failure Symptom Report
| =====
|
| Date/Time      :- Mon January 28 2008 21:59:06 GMT
| UTC Time/Zone  :- 1201539869.892015 0 GMT
| Host Name     :- mqperfh2 (HP-UX B.11.23)
| PIDS          :- 5724H7202
| LVLS         :- 7.0.0.0
| Product Long Name :- WebSphere MQ for HP-UX
| Vendor        :- IBM
| Probe Id      :- XC034255
| Application Name :- MQM
| Component     :- xcsWaitEventSem
| SCCS Info     :- lib/cs/unix/amqxerrx.c, 1.204
| Line Number   :- 6262
| Build Date    :- Jan 25 2008
| CMVC level    :- p000-L050203
| Build Type    :- IKAP - (Production)
| UserID       :- 00000106 (mqperf)
| Program Name  :- amqzmuc0
| Addressing mode :- 64-bit
| Process       :- 15497
| Thread       :- 1
| QueueManager  :- CSIM
| ConnId(2) QM  :- 4
| Major Errorcode :- OK
| Minor Errorcode :- OK
| Probe Type    :- INCORROUT
| Probe Severity :- 4
| Probe Description :- AMQ6109: An internal WebSphere MQ error has occurred.
| FDCSequenceNumber :- 0
|
+-----+
|
| MQM Function Stack
| amqzmuc0
| xcsWaitEventSem
| xcsFFST
|
| MQM Trace History
| Data: 0x00003c87
| --} xcsCheckProcess rc=OK
| --{ xcsRequestMutexSem
| --} xcsRequestMutexSem rc=OK
|
| ...

```

Figure 166. FFST report for WebSphere MQ for UNIX systems

The Function Stack and Trace History are used by IBM to assist in problem determination. In many cases there is little that the system administrator can do when an FFST report is generated, apart from raising problems through the IBM Support Center.

However, there are some problems that the system administrator might be able to solve. If the FFST shows *out of resource* or *out of space on device* descriptions when calling one of the IPC functions (for example, **semop** or **shmget**), it is likely that the relevant kernel parameter limit has been exceeded.

If the FFST report shows a problem with **setitimer**, it is likely that a change to the kernel timer parameters is needed.

To resolve these problems, increase the IPC limits, rebuild the kernel, and restart the machine.

First Failure Support Technology (FFST) files and UNIX and Linux clients

FFST logs are written when a severe WebSphere MQ error occurs. They are written to the directory `/var/mqm/errors`.

These are normally severe, unrecoverable errors and indicate either a configuration problem with the system or a WebSphere MQ internal error.

The files are named `AMQnnnnn.mm.FDC`, where:

- `nnnnn` is the process id reporting the error
- `mm` is a sequence number, normally 0

When a process creates an FFST it also sends a record to the system log. The record contains the name of the FFST file to assist in automatic problem tracking.

The system log entry is made at the "user.error" level.

First Failure Support Technology is explained in detail in [First Failure Support Technology \(FFST\)](#).

First Failure Support Technology (FFST)

Describes the role of First Failure Support Technology (FFST).

For IBM i, FFST information is recorded in a stream file in the `/QIBM/UserData/mqm/errors` directory.

These errors are normally severe, unrecoverable errors, and indicate either a configuration problem with the system or a IBM WebSphere MQ internal error.

The stream files are named `AMQnnnnn.mm.FDC`, where:

`nnnnn` Is the ID of the process reporting the error
`mm` Is a sequence number, normally 0

A copy of the job log of the failing job is written to a file with the same name as the `.FDC` file. The file name ends with `.JOB`.

Some typical FFST data is shown in the following example.

```
-----  
WebSphere MQ First Failure Symptom Report  
-----  
Date/Time      :- Mon January 28 2008 21:59:06 GMT  
UTC Time/Zone  :- 1201539869.892015 0 GMT  
Host Name      :- WINAS12B.HURSLEY.IBM.COM  
PIDS          :- 5733A38  
LVLS          :- 520  
Product Long Name :- WebSphere MQ for IBM i  
Vendor         :- IBM  
Probe Id      :- XY353001  
Application Name :- MQM  
Component     :- xehAS400ConditionHandler  
Build Date    :- Feb 25 2008  
UserID        :- 00000331 (MAYFCT)  
Program Name  :- STRMQM_R MAYFCT  
Job Name      :- 020100/MAYFCT/STRMQM_R  
Activation Group :- 101 (QMOM) (QMOM/STRMQM_R)  
Process       :- 00001689  
Thread        :- 00000001  
QueueManager  :- TEST.AS400.OE.P  
Major Errorcode :- STOP
```

```

Minor Errorcode :- OK
Probe Type      :- HALT6109
Probe Severity  :- 1
Probe Description :- 0
Arith1         :- 1 1
Comment1       :- 00d0

```

```

MQM Function Stack
lpiSPIMQConnect
zstMQConnect
ziiMQCONN
ziiClearUpAgent
xcsTerminate
xlsThreadInitialization
xcsConnectSharedMem
xstConnSetInSPbyHandle
xstConnSharedMemSet
xcsFFST

```

```

MQM Trace History
<-- xcsCheckProcess rc=xecP_E_INVALID_PID
--> xcsCheckProcess
<-- xcsCheckProcess rc=xecP_E_INVALID_PID
--> xlsThreadInitialization
--> xcsConnectSharedMem
--> xcsRequestThreadMutexSem
<-- xcsRequestThreadMutexSem rc=OK
--> xihGetConnSPDetailsFromList
<-- xihGetConnSPDetailsFromList rc=OK
--> xstCreateConnExtentList
<-- xstCreateConnExtentList rc=OK
--> xstConnSetInSPbyHandle
--> xstSerialiseSPList
--> xlsSpinLockRequest
<-- xlsSpinLockRequest rc=OK
<-- xstSerialiseSPList rc=OK
--> xstGetSetDetailsFromSPbyHandle
<-- xstGetSetDetailsFromSPbyHandle rc=OK
--> xstConnSharedMemSet
--> xstConnectExtent
--> xstAddConnExtentToList
<-- xstAddConnExtentToList rc=OK
<-- xstConnectExtent rc=OK
--> xcsBuildDumpPtr
--> xcsGetMem
<-- xcsGetMem rc=OK
<-- xcsBuildDumpPtr rc=OK
--> xcsBuildDumpPtr
<-- xcsBuildDumpPtr rc=OK
--> xcsBuildDumpPtr
<-- xcsBuildDumpPtr rc=OK
--> xcsFFST

```

```

Process Control Block
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :8bba0:0:6d E7C9C8D7 000004E0 00000699 00000000 XIHP...\...r....
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :8bbb0:1:6d 00000000 00000002 00000000 00000000 .....
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :8bbc0:2:6d 80000000 00000000 EC161F7C FC002DB0 .....@...¢
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :8bbd0:3:6d 80000000 00000000 EC161F7C FC002DB0 .....@...¢
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :8bbe0:4:6d 00000000 00000000 00000000 00000000 .....

```

```

Thread Control Block
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :1db0:20:6d E7C9C8E3 00001320 00000000 00000000 XIHT.....
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :1dc0:21:6d 00000001 00000000 00000000 00000000 .....
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :1dd0:22:6d 80000000 00000000 DD13C17B 81001000 .....A#a...
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :1de0:23:6d 00000000 00000046 00000002 00000001 .....
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :1df0:24:6d 00000000 00000000 00000000 00000000 .....

```

```

RecoveryIndex
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :2064:128:6d 00000000 ....

```

Note:

1. The MQM Trace History section is a log of the 200 most recent function trace statements, and is recorded in the FFST report regardless of any TRCMQM settings.
2. The queue manager details are recorded only for jobs that are connected to a queue manager subpool.

3. When the failing component is `xehAS400ConditionHandler`, additional data is logged in the errors directory giving extracts from the job log relating to the exception condition.

The function stack and trace history are used by IBM to assist in problem determination. In most cases, there is little that the system administrator can do when an FFST report is generated, apart from raising problems through the IBM Support Center.

FFST: WebSphere MQ for HP Integrity NonStop Server

Describes the name, location, and contents of the First Failure Support Technology™ (FFST™) files for HP Integrity NonStop Server systems.

In IBM WebSphere MQ client for HP Integrity NonStop Server systems, FFST information is recorded in a file in the `<mqpath>/var/mqm/errors` directory.

An FFST file contains one or more records. Each FFST record contains information about an error that is normally severe, and possibly unrecoverable. These records indicate either a configuration problem with the system or a IBM WebSphere MQ internal error.

FFST files are named `AMQ.nnn.xx.ppp.qq.FDC`, where:

- nnn* The name of the process that is reporting the error.
- xx* The processor number on which the process is running.
- ppp* The PIN of the process that you are tracing.
- qq* A sequence that starts at 0. If the full file name exists, this value is incremented by one until a unique FFST file name is found. An FFST file name can exist if a process is reused.

Each field can contain fewer or more digits than shown in the example.

An instance of a process writes all FFST information to the same FFST file. If multiple errors occur during a single execution of the process, an FFST file can contain many records.

To read the contents of an FFST file, you must be either the creator of the file, or a member of the `mqm` group.

When a process writes an FFST record, it also creates an EMS event.

Figure 167 on page 1323 shows a typical FFST report for a IBM WebSphere MQ client on a HP Integrity NonStop Server system:

```

+-----+
| WebSphere MQ First Failure Symptom Report
| =====
|
| Date/Time      :- Mon April 29 2013 10:21:26 EDT
| UTC Time       :- 1367245286.105303
| UTC Time Offset :- -240 (EST)
| Host Name      :- MYHOST
| Operating System :- HP NonStop J06.14, NSE-AB 069194
|
| PIDS           :- 5724H7222
| LVLS           :- 7.1.0.0
| Product Long Name :- WebSphere MQ for HP NonStop Server
| Vendor         :- IBM
| Installation Path :- /home/cmarti/client/opt/mqm
| Probe Id       :- MQ000020
| Application Name :- MQM
| Component      :- Unknown
| SCCS Info      :- S:/cmd/trace/amqxdspa.c,
| Line Number    :- 3374
| Build Date     :- Apr 24 2013
| Build Level    :- D20130424-1027
| Build Type     :- ICOL - (Development)
| File Descriptor :- 6
| Effective UserID :- 11329 (MQM.CMARTI)
| Real UserID    :- 11329 (MQM.CMARTI)
| Program Name   :- dspmqtrc
| Addressing mode :- 32-bit
| LANG           :-
| Process        :- 1,656 $Y376 OSS 469762429
| Thread(n)     :- 1
| UserApp       :- FALSE
| Last HQC      :- 0.0.0-0
| Last HSHMEMB  :- 0.0.0-0
| Major Errorcode :- krcE_UNEXPECTED_ERROR
| Minor Errorcode :- OK
| Probe Type    :- INCORROUT
| Probe Severity :- 2
| Probe Description :- AMQ6125: An internal WebSphere MQ error has occurred.
| FDCSequenceNumber :- 0
| Comment1      :- AMQ.3.520.sq_tc.0.TRC
| Comment2      :- Unrecognised hookID:0x3 at file offset 0x4b84
|
+-----+
|
| MQM Function Stack
| xcsFFST
|
| MQM Trace History
| { xppInitialiseDestructorRegistrations
| } xppInitialiseDestructorRegistrations rc=OK
| { xcsGetEnvironmentInteger
| -{ xcsGetEnvironmentString
|
| ...

```

Figure 167. Sample FFST data

The Function Stack and Trace History are used by IBM to help problem determination. In many cases, there is little that the system administrator can do when an FFST report is generated, apart from raising problems through the IBM Support Center. However, there are some problems that the system administrator might be able to solve, for example, if the FFST report shows Out of resource or Out of space on device.

For more information about FFST, see “First Failure Support Technology (FFST)” on page 1315.

IBM Support Assistant (ISA)

The IBM Support Assistant (ISA) helps you to resolve questions and problems with IBM software products by providing access to support-related information and troubleshooting tools.

ISA is available at no charge to install on your computer; you then install the relevant product add-ons. ISA has a built-in user guide, and the ISA download package includes a quick start installation and configuration guide. This topic contains a brief overview of the features of the ISA Workbench Version 4; you can find more detailed information on the [IBM Support Assistant](#) web page.

From the ISA home page you can search for information, analyze problems, and collect data to send to IBM.

Find information

Click **Find Information** to search multiple information sources concurrently. You can search the following sources.

- IBM software support documents
- IBM Developer
- IBM news groups and forums
- Google Web search
- Guided troubleshooter content
- Product documentation

Analyze problem

Click **Analyze Problem** to access diagnostic tools or a guided troubleshooter, or to collect data. More tools might be made available periodically, so check for updates by clicking **Find new add-ons**.

Collect and send data

Click **Collect and Send Data** to complete the following tasks.

- Collect diagnostic data automatically from a local or remote computer.
- Send files to IBM Support for problem determination.
- Create and submit a new problem report.
- View or update an existing problem report.

For information on installing the ISA, see “Installing the IBM Support Assistant (ISA).”

Related concepts:

“Troubleshooting overview” on page 1149

Related tasks:

“Searching knowledge bases” on page 1326

“Contacting IBM Software Support” on page 1346

Installing the IBM Support Assistant (ISA)

You can install the IBM Support Assistant (ISA) from the ISA downloads Web page.

Before you begin

Before you start:

Read the concept topic about the “IBM Support Assistant (ISA).”

About this task

Follow these steps to install ISA on your computer:

Procedure

1. Go to the  IBM Support Assistant web page to download the installation package.
2. Log in by using your IBM ID and password. If you do not have an IBM ID, click **Get an IBM ID** to create one.
3. Select the version of ISA that you want and click **Continue**.
4. Click **View license** to read the license agreement in a separate window, then select **I agree** and click **I confirm**.
5. Select the relevant operating system, click **Download now**, and save the compressed file to a temporary directory.
6. Extract the files from the compressed file to a temporary directory. The files that you extract include a quick start guide that tells you how to install, upgrade, and configure ISA.
7. Follow the instructions in the quick start guide to install ISA.

Results

When you have installed ISA successfully, you can use the desktop icon to open it, or you can click **Programs > IBM Support Assistant > IBM Support Assistant**. (The exact name of the entry in the Start menu depends on the version of ISA that you have installed.)

What to do next

Now that you have installed ISA, install the WMQ add-on, as described in “Updating the IBM Support Assistant (ISA).”

Updating the IBM Support Assistant (ISA)

You can update ISA by installing product and tool add-ons.

Before you begin

Before you start:

1. Read the concept topic about the “IBM Support Assistant (ISA)” on page 1324.
2. Install the IBM Support Assistant.

About this task

To install product add-ons and tool add-ons, complete the following steps.

Procedure

1. Open the IBM Support Assistant by clicking **Programs > IBM Support Assistant > IBM Support Assistant**.
2. Click **Update > Find New** and select either **Product Add-ons** or **Tools Add-ons**.
3. Select the appropriate product add-ons to install and click **Next**. Add-ons are categorized by product family, therefore expand **WebSphere** and select IBM WebSphere MQ.
4. Select the appropriate tool add-ons and click **Next**.
5. Read and accept the license agreement and click **Next**.
6. Click **Finish** to install the selected add-ons.
7. When the installation completes, click **Finish**, then click **Yes** to restart ISA.

What to do next

When you have installed the add-ons in successfully, click **Find Information** to search various forms of information for the products that you have selected. You can also use ISA to analyze problems and collect and send data to IBM.

Searching knowledge bases

If you have a problem with your IBM software, you want it resolved quickly. Begin by searching the available knowledge bases to determine whether the resolution to your problem is already documented.

Before you begin

1. If you have not already installed the IBM Support Assistant, you can find instructions about how to do so in “Updating the IBM Support Assistant (ISA)” on page 1325.
2. Install the IBM Support Assistant WMQ plug-in by following the instructions in “Installing the IBM Support Assistant (ISA)” on page 1324.

Procedure

1. Search the product documentation

IBM provides extensive documentation in the form of online product documentation. You can also install this documentation on your local machine or on a local intranet. You can use the powerful search function of the product documentation to query conceptual and reference information, and to find detailed instructions for completing tasks.

2. Search the IBM database for similar problems

IBM keeps records of all known problems with its licensed programs on its software support database (RETAIN). IBM support center staff continually update this database as new problems are found, and they regularly search the database to see if problems they are told about are already known. You can use one of IBM's search tools to search the database, or you can contact IBM support center to perform the search for you. For more information about searching the IBM database, see “Searching the IBM database for similar problems, and solutions” on page 1327.

3. Search the Internet

If you cannot find an answer to your question in the product documentation, search the Internet for the latest, most complete information that might help you resolve your problem, including:

- IBM technotes
- IBM downloads
- IBM Redbooks
- IBM Developer
- Forums and newsgroups
- Internet search engines

You can use the IBM Support Assistant (ISA) to help in your search of knowledge bases. With ISA, you can:

- Query multiple sources of support information
- Access available diagnostic tools
- Collect diagnostic data automatically
- Send files to IBM Support for problem determination
- Create and submit a new problem report
- View or update an existing problem report

For more information, see “IBM Support Assistant (ISA)” on page 1324.

Related concepts:

“Troubleshooting and support” on page 1149

“IBM Support Assistant (ISA)” on page 1324

Related tasks:

“Contacting IBM Software Support” on page 1346

Searching the IBM database for similar problems, and solutions

IBM maintain a database of known problems, and solutions to some of those problems. Use this topic to understand how to best search the database.

IBM keeps records of all known problems with its licensed programs on its software support database (RETAIN). IBM support center staff continually update this database as new problems are found, and they regularly search the database to see if problems they are told about are already known.

If you have access to one of IBM's search tools such as INFORMATION/ACCESS OR INFORMATION/SYSTEM you can look on the RETAIN database yourself. If not, you can contact the IBM support center to perform the search for you.

You can search the database using a string of keywords to see if a similar problem already exists. This section explains how to search the database using keywords.

You can use the keyword string (also called the symptom string) that appears in a dump or SYS1.LOGREC record to search the database, or you can build your own keyword string from the procedure described in “Building a keyword string” on page 1331. Before you use the procedures in this section, make some initial checks by searching through the following appropriate product documentation section specific to your platform:

- “Making initial checks on Windows, UNIX and Linux systems” on page 1151
- “Making initial checks on IBM i” on page 1161
- “Making initial checks on z/OS” on page 1169

If the search is successful, you find a similar problem description and, usually, a fix. If the search is unsuccessful, you should use these keywords when contacting IBM for additional assistance, or when documenting a possible *authorized program analysis report* (APAR).

Searching the IBM software support database is most effective if you:

- Always spell keywords the way they are spelled in this documentation
- Include all the appropriate keywords in any discussion with your IBM support center

Use the following topics to find out more about searching the IBM database for problems:

- “The search argument process” on page 1328
- “The keyword format” on page 1329
- “Building a keyword string” on page 1331
- “SDB format symptom-to-keyword cross reference” on page 1343
- “WebSphere MQ component and resource manager identifiers” on page 1344

The search argument process

Use this topic to understand how to search the RETAIN database.

Use the following procedure when searching the IBM software support database:

1. Using INFORMATION/ACCESS or INFORMATION/SYSTEM, search the database using the keywords you have developed. Details on how to construct suitable keywords are given in “Building a keyword string” on page 1331

Note: Do *not* use both the CSECT keyword and the load module modifier keyword at the same time for the first search. Refer to “Load module modifier keyword” on page 1337 for additional information.

2. Compare each matching APAR closing description with the current failure symptoms.
3. If you find an appropriate APAR, apply the correction or PTF.
4. If you do not find an appropriate APAR, vary the search argument by following the suggestions provided under “Techniques for varying the search process.”
5. If you still cannot find a similar problem, see “Resolving a problem” on page 1354.

Techniques for varying the search process:

You can widen or narrow the scope of your search or you can alter the keywords to make the search more precise.

To vary your search, follow these guidelines:

Dropping keywords to widen your search

If you used a complete set of keywords (as described in “Building a keyword string” on page 1331) and could not find any problem descriptions to examine, drop one or more of the following keywords and try again:

- Release-level keyword
- Load Module modifier keyword
- Recovery routine modifier keyword
- CSECT keyword

Adding keywords to narrow your search

If you tried to search with an incomplete set of keywords and found too many problem descriptions to examine, add keywords to narrow your search. For example, for storage manager abends (which produce a reason code beginning with X'00E2'), you use the CSECT name recorded in the VRA to narrow or vary the search.

Making your set of keywords more precise

If you tried to search with a complete set of keywords and found too many matching descriptions and if you received a 4-byte WebSphere MQ abend reason code, you might be able to make your set of keywords more precise. Look up the 4-byte abend reason code in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) to find additional information available for this problem.

Replacing keywords to locate problems

If your type-of-failure keyword is WAIT, LOOP, or PERFM, and if you did not find a matching problem description, replace that keyword with one of the other two listed here. Sometimes a problem that appears to be a performance problem might actually be a WAIT or LOOP; likewise, a problem that seems to be a WAIT or a LOOP might actually be recorded as a performance problem.

Using message numbers in your search

If your type-of-failure keyword is MSGx and you received more than one message near the time of the problem, repeat the search replacing the message number in the keyword with the number of each related message in turn.

Using DOC as a keyword in your search

If your type-of-failure keyword is MSGx, PERFM, or INCORROUT, and if the problem occurred immediately after you performed some action that a WebSphere MQ documentation told you to perform, the problem could be recorded as a DOC type of failure. In this case, try searching with DOC as your type-of-failure keyword, rather than with MSGx, PERFM, or INCORROUT.

The keyword format

Searches can be performed using free format keywords or structured database (SDB) format keywords. Use this topic to understand how to perform searches using different keyword formats.

The keyword formats

The keywords in “Building a keyword string” on page 1331 are described in two distinct formats: the z/OS, or free format; and the structured database (SDB) format. Structured symptoms are also called RETAIN symptoms and “failure keywords”.

If your installation has a tool for performing structured searches, you can use the SDB format. Otherwise, you should use the free format. For both formats, your choice of keywords depends on the type of failure that occurred.

- Free format
- Structured database (SDB) format

Free format

A free form keyword can consist of any piece of data that is related to the problem. To help you search the database, a set of keywords has been defined, and you can use them to narrow your search. (For example, if you know the name of the CSECT in error, you can use this to search, but if you add the MSGxx or ABEND keyword, your search will be more precise.)

The following list shows keywords defined for use in a free format search:

Table 145. Keywords defined for use in a free format search

Keyword	Meaning
ABEND	Abnormal termination of a task; no error message.
ABENDxx	Abnormal termination of a task; xx is the abend code.
ABENDUxx	User abend; xx is the abend code.
DOC	Documentation discrepancy that caused a problem.
HALTxx	Halt; xx is the halt number.
INCORROUT	Any incorrect data output, except performance degradation.
INTEG	Integrity problem.
LOOP	Loop.
MSGxx	Any message; xx is the message identifier.
PERFM	Performance degradation.
PROCCHK	Processor check.
PROGCH	Program check.
WAIT	Wait condition; undocumented and no identifier.

Table 145. Keywords defined for use in a free format search (continued)

Keyword	Meaning
WAITxx	System wait condition; xx is the identifier.

Structured database (SDB) format

The structured symptoms consist of a prefix keyword, which identifies the type of symptom, followed by a slash (/) and the data portion of the symptom.

- The prefix keyword has one through eight characters.
- All characters must be alphanumeric, #, @, or \$.
- At least one character of data is required.
- The maximum length, including the prefix, is 15 characters.

For example, the following is a structured symptom string for a message identifier of CSQC223D:

MS/CSQC223D

The following list shows the structured symptom strings:

Table 146. Keywords defined for use in a structured format search

Keyword	Meaning
AB	Abend code.
FLDS	Name of a field or control block involved with the problem.
LVLS	Level of the base system or licensed program.
MS	Message identifier.
OPCS	Operation code (opcode) for software, such as an assembler-language opcode.
PCSS	Program command or other software statement, such as JCL, a parameter, or a data set name.
PIDS	Program identifier for a component involved in the problem.
PRCS	Program return code, generated by software, including reason codes and condition codes.
PTFS	Program temporary fix (PTF) for software associated with a problem.
PUBS	Identifier of a publication associated with a problem.
RECS	Record associated with a problem.
REGS	Register for a software program associated with a problem. The value can be the register/PSW difference (<i>rrddd</i>), which the STATUS FAILDATA subcommand of IPCS provides for abends. The difference (<i>ddd</i>) is a hexadecimal offset from a probable base register or branch register (<i>rr</i>).
RIDS	Routine identifier, such as the name of a CSECT or subroutine. If the RIDS/ value has no suffix, the value is a CSECT name. The following suffixes are supported: <ul style="list-style-type: none"> • #L for a load module • #R for a recovery routine
VALU	Value in a field or register. One of the following qualifiers is required as the first character of the value: <ul style="list-style-type: none"> • B for a bit value • C for a character value • H for a hexadecimal value

Table 146. Keywords defined for use in a structured format search (continued)

Keyword	Meaning
WS	Wait state code issued by the system, or device-issued wait code. One of the following qualifiers is required as the first character of the value: <ul style="list-style-type: none">• D for disabled wait (system disabled for I/O or external interrupts)• E for enabled wait

For more information about which prefix keyword to use for which type of symptom, see “SDB format symptom-to-keyword cross reference” on page 1343.

Building a keyword string

This section describes a systematic way of selecting *keywords* to describe a failure in WebSphere MQ for z/OS. Keywords are predefined words or abbreviations that identify aspects of a program failure.

To determine which WebSphere MQ for z/OS keywords to use and the procedures for selecting them, see the flowchart in Figure 168 on page 1332.

To begin selecting your keywords:

1. Follow the procedures in “The component-identifier keyword” on page 1333 and “The release-level keyword” on page 1333. Do this for all failures.
2. Follow one of the type-of-failure keyword procedures.
3. Identify the area of the failure using CSECT and modifier keywords when appropriate. The procedures in this section refer you to these steps as needed.
4. Follow “Searching the IBM database for similar problems, and solutions” on page 1327 to learn how to search the database with your set of keywords. Do this for all failures.
5. If the search is unsuccessful, turn to “Resolving a problem” on page 1354. This helps IBM product support personnel determine whether an APAR should be submitted.

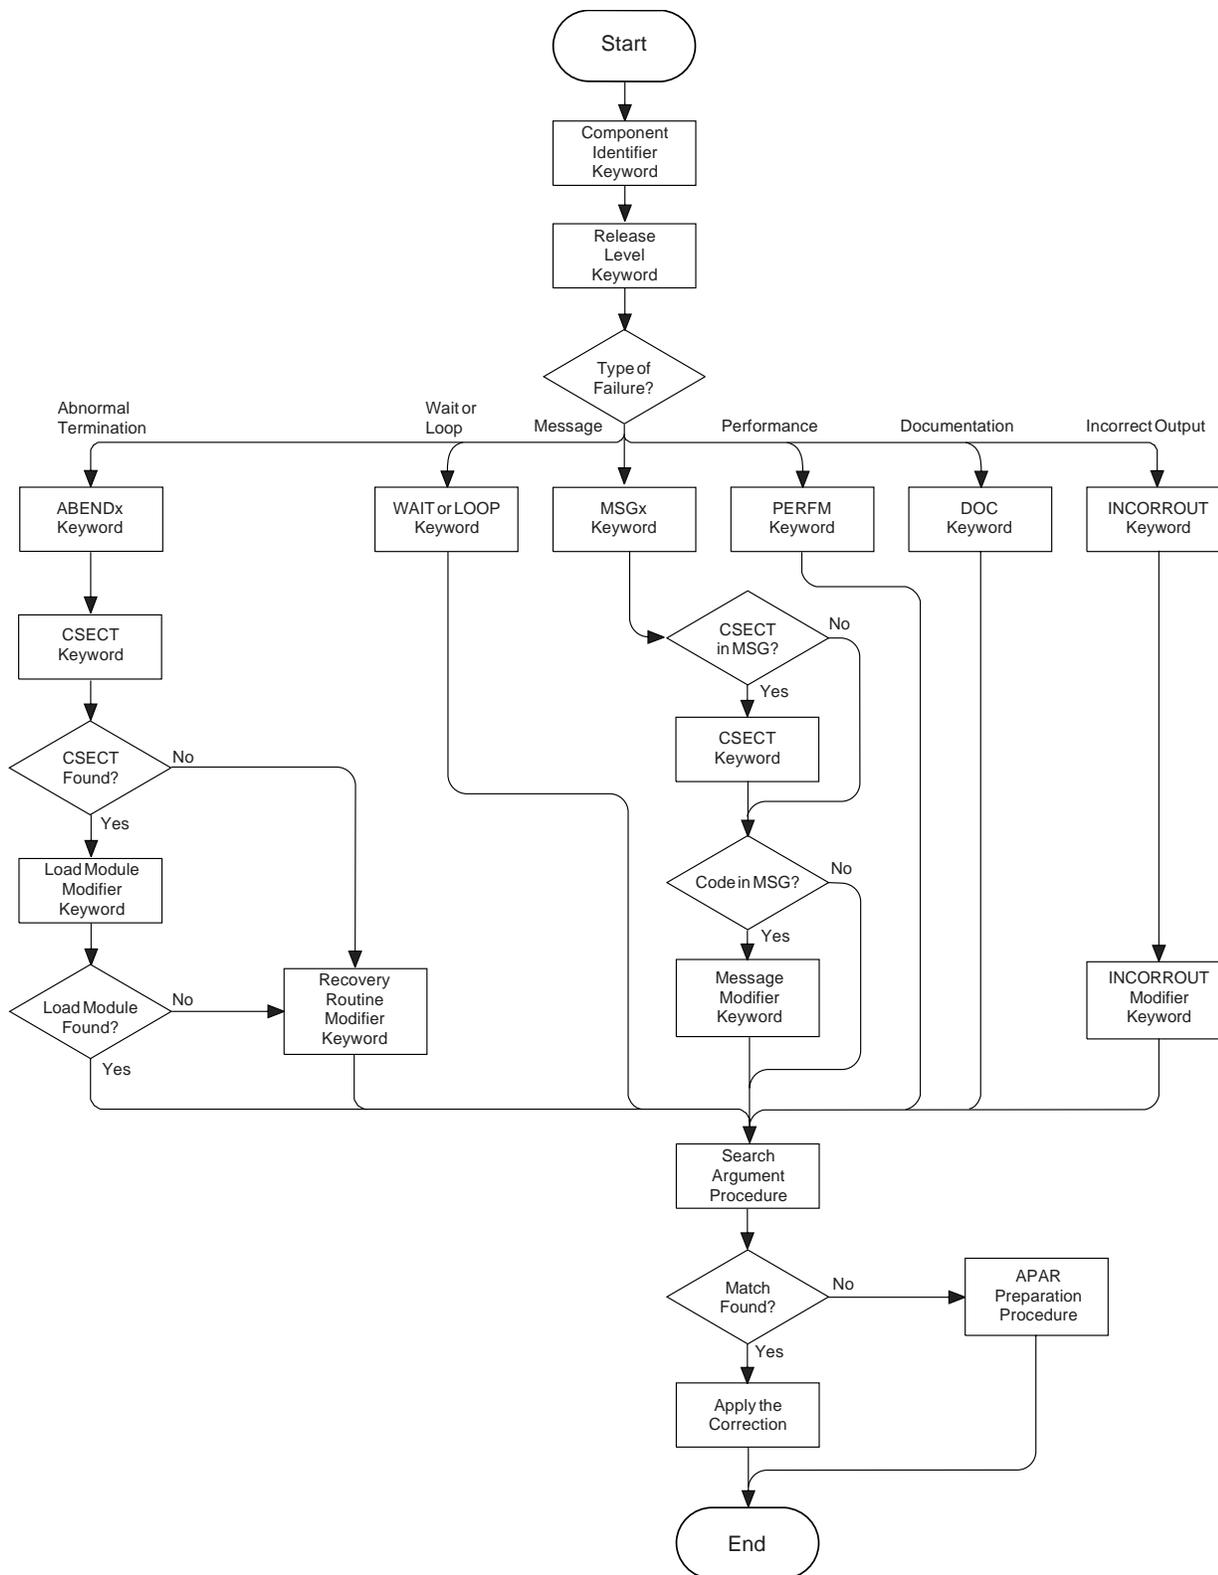


Figure 168. High-level flowchart of various sets of keywords

Use the following topics to build a keyword string for searching the IBM database for problems:

- “The component-identifier keyword” on page 1333
- “The release-level keyword” on page 1333

- “Type-of-failure keyword” on page 1334
- “The abend keyword, and its associated keywords” on page 1334
- “Wait and loop keywords” on page 1338
- “The message keyword” on page 1338
- “Performance keyword” on page 1340
- “Documentation keyword” on page 1341
- “Incorrect output keyword” on page 1342

Related concepts:

“The search argument process” on page 1328

“The keyword format” on page 1329

Related reference:

“SDB format symptom-to-keyword cross reference” on page 1343

“WebSphere MQ component and resource manager identifiers” on page 1344

The component-identifier keyword:

You can use the component-identifier as a keyword to search the IBM documentation and software support database for known problems and solutions.

The *component-identifier keyword* identifies the library within the IBM software support database that contains *authorized program analysis reports* (APARs) and *program temporary fixes* (PTFs) for the product.

The component-identifier keyword for WebSphere MQ for z/OS is **5655R3600**.

This section describes how to determine the nine-digit component identifier keyword for your failure to verify that the problem was caused by WebSphere MQ for z/OS. If the component identifier is not 5655R3600, the problem could be caused by another product.

If the problem caused a dump, display the dump title, locate the COMP= label, and note the first five characters following that label. If these characters are **R3600**, the problem was caused by WebSphere MQ for z/OS. Append those five characters to **5655** and use this as the first keyword in your search argument.

ssnm,ABN=compltn-reason,U=userid,C=compid.release.comp-function,
M=module, LOC=loadmod.csect+csect_offset

If you cannot use the dump title, display the z/OS SYMPTOM STRING in the formatted dump. Note the nine characters following the PIDS/ label.

Related concepts:

“WebSphere MQ dumps” on page 1276

The release-level keyword:

You can use the release-level as a keyword to search the IBM documentation and software support database for known problems and solutions.

The *release-level keyword* narrows the symptom search to your specific release level. Using this keyword is optional, but suggested, when searching the IBM software support database. It is required, however, when an APAR is submitted.

Locate the three-digit release identifier in the dump title. It follows COMP=R3600, for example:

COMP=R3600.710

Add the release-level to your keyword string, in one of the formats shown:

Free format

5655R3600 R710

Structured format

PIDS/5655R3600 LVLS/710

Type-of-failure keyword:

You can use the type-of-failure as a keyword to search the IBM documentation and software support database for known problems and solutions.

To narrow your search, use one or more of the type-of-failure and modifier keywords to describe an external symptom of a program failure. The various types of failures are shown in Table 147. Use this table to find the name and page number of the keyword that best matches your problem.

Table 147. Types of WebSphere MQ for z/OS failures

Problem	Procedure
Abend of the subsystem or task	"The abend keyword, and its associated keywords"
Unexpected program suspension	"Wait and loop keywords" on page 1338
Uncontrolled program looping (often signaled by repeating messages or output)	"Wait and loop keywords" on page 1338
Errors signaled by or associated with messages	"The message keyword" on page 1338
Performance degradation	"Performance keyword" on page 1340
Documentation problem	"Documentation keyword" on page 1341
Unexpected or missing output	"Incorrect output keyword" on page 1342

The abend keyword, and its associated keywords:

The abend keyword is often seen in association with other keywords. These keywords can be used together to form a keyword string.

Use the ABEND keyword when the subsystem or task terminates abnormally. This procedure describes how to locate the abend completion code and the abend reason code (if there is one), and how to use them in a set of keywords. Check the SYS1.LOGREC to determine how many abends there were. Sometimes an earlier abend causes a secondary abend that causes a dump. If no dump has been taken, try searching the database with a minimum symptom string (the component-identifier, and release-level keywords). If you cannot find any information that seems to relate to your problem, contact your IBM support center.

When a WebSphere MQ for z/OS abend occurs, you will see one of the following symptoms:

- An IEA911E message from z/OS, indicating that an SVC dump occurred. See "IEA911E message" on page 1335.
- The CSQV086E message QUEUE MANAGER ABNORMAL TERMINATION REASON=xxxxxxx. See "CSQV086E message" on page 1336.
- "CSECT keyword" on page 1336
- "Load module modifier keyword" on page 1337
- "Recovery routine modifier keyword" on page 1337

IEA911E message:

The message code IEA911E can be analyzed to give further information about abends.

You can use the information from the IEA911E message to extract further details about the abend code. Use the following steps to assist with the analysis:

1. Use the DISPLAY DUMP,TITLE command on the console to display the SVC dump title for this abend, or use one of the methods described in “WebSphere MQ dumps” on page 1276 to look at the dump title in the dump.

Note: If the first five characters of the COMP field are not R3600, or the dump title is not of the same form as Figure 162 on page 1289 or Figure 163 on page 1290, the problem was not caused by WebSphere MQ for z/OS, or you are looking at the wrong dump.

2. Locate the 3-character completion code following the word ABND.
 - If the completion code is X'071', or X'122', the operator pressed the RESTART key or canceled the job, probably to break a loop. Verify that this is the case, and turn to “Wait and loop keywords” on page 1338.
 - Otherwise, add this to your keyword string, in one of the formats shown in this topic (in this example, X'0C4' is used):

Free format

5655R3600 R710 **ABEND0C4**

Structured format

PIDS/5655R3600 LVLS/710 **AB/S00C4**

3. Some abends also have reason codes. These reason codes are usually found in message CSQV086E, and register 15 at the time of the abend. Locate the reason code for the abend either:
 - In the 4-byte reason code field in a dump title generated by WebSphere MQ for z/OS
 - In the registers at time of error in the abstract information section of the dump
 - From the value of register 15 in the error summary display
4. If the completion code is X'5C6', review the diagnostic information for the reason code in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*). Follow any procedures recommended there.

If the completion code is anything else, and you have found a reason code, check the value against the description of the abend code in the *MVS System Codes* manual to see if it is valid for the abend completion code.

5. Add the reason code to the keyword string (in this example X'00E20015' is used):

Free format

5655R3600 R710 ABEND5C6 **RC00E20015**

Structured format

PIDS/5655R3600 LVLS/710 AB/S05C6 **PRCS/00E20015**

Refer to “CSECT keyword” on page 1336.

CSQV086E message:

The text from the message code CSQV086E can be analyzed to give further information about abends.

You can use the information from the CSQV086E message to extract further details about the abend code. Use the following steps to assist with the analysis:

1. Issue the DISPLAY DUMP command to see whether any SVC dumps occurred near the time the message appeared. (See the *MVS System Commands* manual if necessary.)
2. If there was only one SVC dump for the abend, follow the procedure starting at step 1 on page 1335.
3. If there were two or more SVC dumps, follow the steps here.
 - a. Read the sections under  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*) that describe the reason code appearing in your message, and any reason codes appearing in the SVC dump titles. Reason codes appear after the completion code in the SVC dump title. For an example, see “Analyzing the dump and interpreting dump titles” on page 1288.
 - b. Compare the reason codes in the SVC dumps to determine which dump relates to the CSQV086E message.
 - c. Use that SVC dump and follow the procedure starting at step 1 on page 1335.
4. If there were two or more different abends, follow the steps here:
 - a. Determine which abend was the original cause by reviewing the time stamps in the SYS1.LOGREC entries.
 - b. Use that SVC dump and follow the procedure starting at step 1 on page 1335.
5. If there were no SVC dumps for the abend, follow the steps here.
 - a. Locate the 4-byte reason code in the message.
 - b. Review the diagnostic information in  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*). Follow any procedures recommended there.
 - c. Add this to your keyword string, in one of the formats shown in this topic (in this example, a reason code of X'00D93001' is used):

Free format

5655R3600 R710 ABEND6C6 **RC00D93001**

Structured format

PIDS/5655R3600 LVLS/710 AB/S06C6 **PRCS/00D93001**

Refer to “CSECT keyword.”

CSECT keyword:

The CSECT keyword and parameter is often associated with an abend code and can assist with identifying where the abend problem occurred.

To find the name of the failing CSECT, locate the LOC= label; the second word following it is the CSECT name. For an example, see “Analyzing the dump and interpreting dump titles” on page 1288.

Any CSECT name you locate should begin with the letters CSQ or CMQ. If you find a CSECT name with a different prefix, the problem is probably not in WebSphere MQ for z/OS.

Add the CSECT name to your keyword string:

Free format

5655R3600 R710 ABEND0C4 **CSQVATRM**

Structured format

PIDS/5655R3600 LVLS/710 AB/S00C4 **RIDS/CSQVATRM**

If required, narrow your search further by referring to “Load module modifier keyword.”

If you cannot find the CSECT, see “Recovery routine modifier keyword.”

Load module modifier keyword:

The Load module modifier keyword and parameter is often associated with an abend code and can assist with identifying where the abend problem occurred.

Use the load module modifier keyword to identify the name of the load module involved if your search using the CSECT keyword was unsuccessful, or yielded too many possible matches:

- If your search was unsuccessful, replace the CSECT name with the load module name and try again.
- If your search yielded too many possible matches, add the load module name to your string to further narrow the search.

All WebSphere MQ for z/OS load module names begin with CSQ or CMQ. If you follow these instructions and find a load module name with a different prefix, the problem is in another product.

To locate the load module name, locate the first word following the label LOC=. This is the load module name, and it precedes the CSECT name. (For an example, see “Analyzing the dump and interpreting dump titles” on page 1288.)

Add the load module name to your keyword string, or substitute it for the CSECT name as appropriate. If you are using the structured format, follow the name of the module with the characters #L to indicate that this is a load module. Search the database again using the revised keyword string. (See “Searching the IBM database for similar problems, and solutions” on page 1327.)

Free format

5655R3600 R710 ABEND5C6 RC00E50013 **CSQSLD1** CSQSVSTK (with load module name and then CSECT name)

5655R3600 R710 ABEND5C6 RC00E50013 **CSQSLD1** (with load module name only)

Structured format

PIDS/5655R3600 LVLS/710 AB/S05C6 PRCS/00E50013 **RIDS/CSQSLD1#L** RIDS/CSQSVSTK (with load module name and then CSECT name)

PIDS/5655R3600 LVLS/710 AB/S05C6 PRCS/00E50013 **RIDS/CSQSLD1#L** (with load module name only)

Recovery routine modifier keyword:

The Recovery routine modifier keyword and parameter is often associated with an abend code and can assist with identifying where the abend problem occurred.

Include the name of the recovery routine only when you could not determine the names of the CSECT and load module involved at the time of failure, after looking in both the SVC dump and the SYS1.LOGREC entry.

To obtain the recovery routine name, locate the area of the dump title containing the symbol M=. The word following this identifies the functional recovery routine (FRR) or the extended specify task abnormal exit (ESTAE). For an example, see “Analyzing the dump and interpreting dump titles” on page 1288.

Add this word to your keyword string. If you are using the structured format, follow the name of the module with the characters #R to indicate that this is a recovery routine. Search the database (see “Searching the IBM database for similar problems, and solutions” on page 1327).

Free format

5655R3600 R710 ABEND5C6 RC00E20015 **CSQTFRCV**

Structured format

PIDS/5655R3600 LVLS/710 AB/S05C6 PRCS/00E20015 **RIDS/CSQTFRCV#R**

Wait and loop keywords:

The keywords wait and loop can be used as part of a string for searching the IBM documentation and IBM software support database. This can assist in identifying known problems and resolutions.

If the problem occurred immediately after you did something a WebSphere MQ manual told you to do, the problem might be related to the documentation. If you think that this is the case, see “Documentation keyword” on page 1341.

If you have verified that the wait or loop problem cannot be resolved through other means, use the following procedure:

1. Add WAIT or LOOP to your keyword string, in one of the formats shown in this topic (in this example **WAIT** is used).

Free format

5655R3600 R710 **WAIT**

Structured format

PIDS/5655R3600 LVLS/710 **WAIT**

2. See “Searching the IBM database for similar problems, and solutions” on page 1327.

The message keyword:

The message keyword can be used to search the IBM documentation and software support database for known problems and solutions

Use the MSG keyword if an error is associated with a WebSphere MQ for z/OS message. If you received multiple messages for one error, search the database using the first message issued. If unsuccessful, search the database using the next message, then the next, and so on.

To see if other messages related to your problem have been issued, check the console for WebSphere MQ for z/OS messages, as well as messages issued by other products. If any message is prefixed with “IEC”, indicating it was issued by data management services, check the SYSLOG for messages that identify associated data set problems. SYSLOG can also help to diagnose user errors.

If your message was issued immediately after you did something that a WebSphere MQ manual told you to do, the problem might be related to the documentation rather than to the message. If this is the case, refer to “Documentation keyword” on page 1341. Otherwise, compare the message prefix with those shown in the table Message Prefixes to determine the appropriate procedure to follow.

Table 148. Message prefixes

Prefix	Component	Procedure
AMQ	WebSphere MQ (not z/OS)	Consult  WebSphere MQ messages (<i>WebSphere MQ V7.1 Programming Guide</i>)
AMT	WebSphere MQ	For details of the MA0F SupportPac see  SupportPacs for IBM MQ and other project areas
ATB	APPC	Consult <i>MVS System Messages</i>
ATR	Resource recovery services	Consult <i>MVS System Messages</i>
CBC	C/C++	Consult <i>C/MVS™ User's Guide</i>
CEE	Language Environment®	Consult <i>Language Environment Debugging Guide and Run-Time Messages</i>
CSQ	WebSphere MQ for z/OS	Follow "Procedure for WebSphere MQ for z/OS messages"
CSV	Contents supervision	Consult <i>MVS System Messages</i>
DFH	CICS	Consult <i>CICS Messages and Codes</i>
DFS	IMS	Consult <i>IMS/ESA Messages and Codes</i>
DSN	Db2	Consult <i>Db2 Messages and Codes</i>
EDC	Language Environment	Consult <i>Language Environment Debugging Guide and Run-Time Messages</i>
EZA, EZB, EZY	TCP/IP	Consult <i>z/OS V2R6.0 eNetwork CS IP Messages: Volumes 1, 2, and 3</i>
IBM	Language Environment	Consult <i>Language Environment Debugging Guide and Run-Time Messages</i>
ICH	RACF	Consult <i>z/OS Security Server (RACF) Messages and Codes</i>
IDC	Access method services	Consult <i>MVS System Messages</i>
IEA	z/OS system services	Consult <i>MVS System Messages</i>
IEC	Data management services	Consult <i>MVS System Messages</i>
IEE, IEF	z/OS system services	Consult <i>MVS System Messages</i>
IKJ	TSO	Consult <i>MVS System Messages</i>
IST	VTAM	Consult <i>VTAM Messages</i>
IWM	MVS workload management services	Consult <i>MVS System Messages</i>
IXC	Cross-system coupling facility (XCF)	Consult <i>MVS System Messages</i>
IXL	Cross-system Extended Services (XES)	Consult <i>MVS System Messages</i>

Related concepts:

"Procedure for WebSphere MQ for z/OS messages"

Procedure for WebSphere MQ for z/OS messages:

The WebSphere MQ messages can be analyzed for the possible cause of an error. To do this you must analyze the text of the message and use the components of the text to search the IBM software support database.

Analyze the text of the any WebSphere MQ messages and use the parts of the text, for example CSECT names, variables and message codes, to search IBM software support database for known problems and resolutions.

1. Check whether the name of the CSECT issuing the message appears. This name follows the message number. If no CSECT name appears, only one CSECT can issue this message.
2. Determine whether the message contains any variables, such as return or reason codes.
3. If no CSECT name appears, add the message number to your keyword string, in one of the formats shown in this topic (in this example, message CSQJ006I is used):

Free format

5655R3600 R710 **MSGCSQJ006I**

Structured format

PIDS/5655R3600 LVLS/R710 **MS/CSQJ006I**

4. If a CSECT name does appear, add both the message number and the CSECT name to your keyword string, in one of the formats shown here (in this example, a message number of CSQJ311E and a CSECT name of CSQJC005 are used):

Free format

5655R3600 R710 **MSGCSQJ311E CSQJC005**

Structured format

PIDS/5655R3600 LVLS/710 **MS/CSQJ311E RIDS/CSQJC005**

5. If the message contains return or reason codes, add these to your keyword string, in one of the formats shown in this topic:

Free format

5655R3600 R710 **MSGCSQM002I RCE**

Structured format

PIDS/5655R3600 LVLS/R710 **MS/CSQM002I PRCs/0000000E**

6. If the message contains any other types of variables, append them to your keyword string.

Free format

5655R3600 R710 **MSGCSQJ104I OPEN**

Structured format

PIDS/5655R3600 LVLS/R710 **MS/CSQJ1041 MS/OPEN**

7. See "Searching the IBM database for similar problems, and solutions" on page 1327.

Performance keyword:

The performance keyword can be used to search the IBM documentation and software support database for known problems and solutions.

You can resolve most performance problems through system tuning, which should be handled by the WebSphere MQ for z/OS system administrator. Before following the procedure show in this topic, use this checklist to verify that the performance problem cannot be resolved through other means:

- See "Dealing with performance problems on z/OS" on page 1293 to determine if you can change the way you have designed your WebSphere MQ for z/OS subsystem and applications to improve their performance.
- Verify that the performance problem is not related to a WAIT or LOOP. See "Dealing with applications that are running slowly or have stopped on z/OS" on page 1294.
- If the problem occurred immediately after you did something a WebSphere MQ manual told you to do, the problem might be related to the manual. See "Documentation keyword" on page 1341.
- If performance degraded after someone tuned WebSphere MQ for z/OS, verify that the tuning options selected were appropriate. Perhaps you can resolve the problem by choosing other options.

If you have verified that the performance problem cannot be resolved through other means, use the following procedure:

1. Record the actual performance, expected performance, and source of expected performance criteria.
2. Add PERFM to your keyword string, as shown in the following example, and see “Searching the IBM database for similar problems, and solutions” on page 1327.

Free format

5655R3600 R710 **PERFM**

Structured format

PIDS/5655R3600 LVLS/710 **PERFM**

3. If required, you can narrow your search by adding free-format keywords that describe what you were doing when you experienced the performance problem.

Documentation keyword:

The documentation keyword can be used to search the IBM software support database for known problems and solutions.

The DOC keyword identifies problems caused by incorrect or missing information in a WebSphere MQ manual. It is possible that a documentation problem could be detected when trying to resolve problems with messages, incorrect output, and performance.

Use the following procedure if you need to use the DOC keyword in your keyword string:

1. Locate the incomplete or erroneous information. Note the page, or topic number, and describe the error and the resulting problem.
2. Add the document number, hyphens omitted, to your keyword string, in one of the formats shown here (in this example, the document number for this manual (GC34-6600-01) is used):

Free format

5655R3600 R710 **DOC GC34660001**

Structured format

PIDS/5655R3600 LVLS/710 **PUBS/GC34660001**

See “Searching the IBM database for similar problems, and solutions” on page 1327.

If your search is unsuccessful, follow Step 3.

3. Broaden your search by replacing the last two digits with two asterisks (**). This searches for all problems on that document, rather than on a specific release of the document.

Free format

5655R3600 R710 **DOC GC346600****

Structured format

PIDS/5655R3600 LVLS/710 **PUBS/GC346600****

If your search is unsuccessful, follow Step 5.

4. If the problem is not too severe, send us a comment using the “Add Comment” option at the bottom of the knowledge center main pane.
5. If the problem is severe, consider initiating a DOC APAR. Use the information gathered in Step 1, and see “Resolving a problem” on page 1354.

Corrections resulting from readers' comments are included in future editions of the manual but are not included in the software support database.

Incorrect output keyword:

You can use the incorrect output keywords to search the IBM software support database for known problems and solutions.

Use the INCORROUT keyword when output was expected but not received, or when output was different from expected. However, if this problem occurred after you did something that WebSphere MQ documentation told you to do, the documentation could be in error. If this is the case, see “Documentation keyword” on page 1341.

1. Add **INCORROUT** to your existing keyword string.

Free format

5655R3600 R710 **INCORROUT**

Structured format

PIDS/5655R3600 LVLS/710 **INCORROUT**

2. Determine the function and secondary modifier keywords for your problem from Table 149 and Table 150.
3. Add the modifier keywords to your string and use it to search the database. See “Searching the IBM database for similar problems, and solutions” on page 1327.

Free format

5655R3600 R710 INCORROUT **RECOVERY BACKOUT**

Structured format

PIDS/5655R3600 LVLS/710 INCORROUT **RECOVERY BACKOUT**

Table 149. INCORROUT modifier keywords: RECOVERY

Secondary keywords	Problem occurrence
none	During recovery
BACKOUT	At backout time
CHECKPOINT	At checkpoint time
COMMIT	At commit time
LOGGING	During logging
RECOVER	During attempt to recover in-doubt
RESTART	During restart process

Table 150. INCORROUT modifier keywords: UTILITY

Secondary keywords	Problem occurrence
none	While running a utility
CSQ1LOGP	While running CSQ1LOGP
CHANGE LOG	While using the Change Log Inventory utility or CSQJUFMT
PRINT LOGMAP	While using the Print Log Map utility
COMMAND	While running the COMMAND function or SDEFS
COPY	While running the COPY function or SCOPY
EMPTY	While running the EMPTY function
FORMAT	While running the FORMAT function, COPYPAGE, RESETPAGE or PAGEINFO
LOAD	While running the LOAD function
CSQUDLQH	While running the dead.letter.queue handler

Table 150. INCORROUT modifier keywords: UTILITY (continued)

Secondary keywords	Problem occurrence
CSQ5PQSG	While running CSQ5PQSG

SDB format symptom-to-keyword cross reference

You can use structured database (SDB) formats to search the RETAIN database.

Structured database (SDB) format is one of the formats that you can use for searching the RETAIN database. Structured symptoms are also called RETAIN symptoms and “failure keywords”. Table 151 lists which prefix keyword to use for which symptom.

“Searching the IBM database for similar problems, and solutions” on page 1327 provides details about searching RETAIN.

Table 151. SDB format symptom-to-keyword cross-reference

Symptom	Keyword	Symptom	Keyword
abend	AB/	access method	RIDS/
address	ADRS/	APAR	PTFS/
assembler macro	RIDS/	assembler message	MS/
CLIST	RIDS/	command	PCSS/
compiler message	MS/	completion code	PRCS/
component	PIDS/	condition code	PRCS/
control block	FLDS/	control block offset	ADRS/
control register	REGS/	CSECT	RIDS/
data set name	PCSS/	dependent component	PIDS/
device error code	PRCS/	disabled wait (coded)	WS/
displacement	ADRS/	display	DEVS/
document	PUBS/	DSECT	FLDS/
enabled wait (coded)	WS/	error code	PRCS/
EXEC	RIDS/	feedback code	PRCS/
field	FLDS/	field value	VALU/
file mode	PCSS/	file name	PCSS/
file type	PCSS/	flag	FLDS/
floating-point register	REGS/	full-screen mode	PCSS/
function key	PCSS/	general purpose register	REGS/
hang	WS/	hung user or task	WS/
I/O operator codes	OPCS/	incorrect output	INCORROUT*
JCL card	PCSS/	JCL parameter	PCSS/
job step code	PRCS/	key	PCSS/
label, code	FLDS/	language statement	PCSS/
level	LVLS/	library name	PCSS/
line command	PCSS/	loop	LOOP*
low core address	ADRS/	machine check	SIG/
macro as a routine	RIDS/	macro as a statement	PCSS/

Table 151. SDB format symptom-to-keyword cross-reference (continued)

Symptom	Keyword	Symptom	Keyword
maintenance level	PTFS/	message	MS/
module	RIDS/	offset	ADRS/
opcode	OPCS/	operator command	PCSS/
operator key	PCSS/	operator message	MS/
option	PCSS/	overlay	OVS/
PA key	PCSS/	panel	RIDS/
parameter	PCSS/	performance	PERFM*
PF key	PCSS/	procedure name	PCSS/
process name	PCSS/	profile option	PCSS/
program check	AB/	program id	RIDS/
program key	PCSS/	program statement	PCSS/
PSW	FLDS/	PTF, PE or otherwise	PTFS/
publication	PUBS/	PUT level	PTFS/
reason code	PRCS/	register value	VALU/
register	REGS/	release level	LVLS/
reply to message	PCSS/	reply to prompt	PCSS/
request code	OPCS/	response to message	PCSS/
response to prompt	PCSS/	return code	PRCS/
routine	RIDS/	service level	PTFS/
special character	PCSS/	SRL	PUBS/
statement	PCSS/	status code	PRCS/
step code	PRCS/	structure word	FLDS/
subroutines	RIDS/	SVC	OPCS/
SYSGEN parameter	PCSS/	system check	PRCS/
table	FLDS/	terminal key	PCSS/
value	VALU/	variable	FLDS/
wait (coded)	WS/	wait (uncoded)	WAIT*

Note: An asterisk (*) indicates that there is no prefix keyword for this type of problem. Use the type-of-failure keyword shown for searches of the software support database.

WebSphere MQ component and resource manager identifiers

Use this topic as a reference for component-identifiers, and resource manager identifiers.

The component-identifier keyword identifies the library within the IBM software support database that contains *authorized program analysis reports* (APARs) and *program temporary fixes* (PTFs) for the product. Resource manager identifiers (RMIDs) are used to limit the volume of data collected in a trace.

Table 152. WebSphere MQ component and resource manager identifiers

ID	Prefix	Hex ID	Component name	RMID
—	AMT	—	AMI	—
—	CMQ	—	Application header files	—
m	CSQm	X'94'	Connection manager	148
t	CSQt	X'A3'	Topic manager	163
A	CSQA	X'C1'	Application interface	—
B	CSQB	X'C2'	Batch adapter	—
C	CSQC	X'C3'	CICS adapter	—
E	CSQE	X'C5'	coupling facility manager	197
F	CSQF	X'C6'	Message generator	24
G	CSQG	X'C7'	Functional recovery manager	199
H	CSQH	X'C8'	Security manager interface	200
I	CSQI	X'C9'	Data manager	201
J	CSQJ	X'D1'	Recovery log manager	4
L	CSQL	X'D3'	Lock manager	211
M	CSQM	X'D4'	Message manager	212
N	CSQN	X'D5'	Command server	213
O	CSQO	X'D6'	Operations and control	—
P	CSQP	X'D7'	Buffer manager	215
Q	CSQQ	X'D8'	IMS adapter	—
R	CSQR	X'D9'	Recovery manager	3
S	CSQS	X'E2'	Storage manager	6
T	CSQT	X'E3'	Timer services	227
U	CSQU	X'E4'	Utilities	—
V	CSQV	X'E5'	Agent services	2
W	CSQW	X'E6'	Instrumentation facilities	16, 26
X	CSQX	X'E7'	Distributed queuing	231
Y	CSQY	X'E8'	Initialization procedures and general services	1
Z	CSQZ	X'E9'	System parameter manager	12
1	CSQ1	X'F1'	Service facilities	—
2	CSQ2	X'F2'	WebSphere MQ-IMS bridge	242
3	CSQ3	X'F3'	Subsystem support	7, 8
4	CSQ4	X'F4'	Sample programs	—
5	CSQ5	X'F5'	Db2 manager	245
6	CSQ6	X'F6'	Customization	—
7	CSQ7	X'F7'	Dump formatting	—
8	CSQ8	X'F8'	Installation	—
9	CSQ9	X'F9'	Generalized command preprocessor	23
—	IMQ	—	C++ bindings	—

Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

Before you begin

Before you start:

1. If you have not already installed the IBM Support Assistant, you can find instructions about how to do so in “Installing the IBM Support Assistant (ISA)” on page 1324.
2. Install the IBM Support Assistant WMQ plug-in by following the instructions in “Updating the IBM Support Assistant (ISA)” on page 1325.

About this task

Before you contact IBM Software Support, you must ensure that your company has an active IBM software subscription and support contract, and that you are authorized to submit problems to IBM:

- If you use IBM distributed software products (including, but not limited to, Tivoli, Lotus, and Rational products, as well as Db2 and WebSphere products that run on Windows or UNIX and Linux operating systems), enroll in Passport Advantage in one of the following ways:
 - Online: Go to the [Passport Advantage](#)[®] and [Passport Advantage Express](#)[®] web site, then click the **Subscription and Support** tab.
 - By telephone: Go to the [Software Support Handbook](#), click **Contacts**, then **Worldwide contacts**.
- If you have a Subscription and Support (S & S) contract, go to the [Open service request Web page](#).
- You might also have an IBMLink, CATIA, Linux, S/390, iSeries, pSeries, zSeries, or other support agreement.
- For IBM eServer software products (including, but not limited to, Db2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software subscription and support agreement by working directly with an IBM marketing representative or an IBM Business Partner. For more information about support for eServer software products, go to the [IBM Support Portal](#).

If you are not sure what type of software subscription and support contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States or, from other countries, go to the contacts page of the [Software Support Handbook](#) and click the name of your geographic region for telephone numbers of people who provide support for your location.

Follow the steps in this topic to contact IBM Software Support:

Procedure

1. “Determine the effect of the problem on your business” on page 1347
2. “Describe your problem and gather background information” on page 1347
3. “Submit your problem to IBM Software Support” on page 1347

Determine the effect of the problem on your business

About this task

When you report a problem to IBM, you will be asked to supply a severity level. Therefore, you need to understand and assess the effect on your business of the problem that you are reporting. Use the following criteria:

Severity	Effect on business
Severity 1	Critical effect on business: You are unable to use the program, resulting in a critical effect on operations. This condition requires an immediate solution.
Severity 2	Significant effect on business: The program is usable but is severely limited.
Severity 3	Some effect on business: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal effect on business: The problem has little effect on operations, or a reasonable workaround to the problem has been implemented.

Describe your problem and gather background information

About this task

When you are explaining a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you to solve the problem efficiently. To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? (For example, hardware, operating system, networking software, and so on.)
- Are you currently using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submit your problem to IBM Software Support

About this task

You can submit your problem in one of three ways:

- By using the IBM Support Assistant: To collect data automatically and submit a request to IBM Software Support, open the IBM Support Assistant and click **Service**. (For more information, see “IBM Support Assistant (ISA)” on page 1324.)
- Online: Go to the  Software Support Handbook and enter your information into the appropriate problem submission tool.
- By telephone: For the telephone number to call in your country, go to the contacts page of the  Software Support Handbook and click the name of your geographic region for telephone numbers of people who provide support for your location.

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support might create an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround for you to implement until the APAR is resolved and a fix is delivered.

IBM publishes resolved APARs on the IBM product support Web pages daily, so that other users who experience the same problem can benefit from the same resolutions.

Related concepts:

“Troubleshooting and support” on page 1149

“Dealing with the support center”

“What the support center needs to know” on page 1350

“What happens next” on page 1351

“Collecting documentation for the problem” on page 1352

“Sending the documentation to the change team” on page 1353

“Resolving a problem” on page 1354

Related tasks:

“Getting product fixes” on page 1355

Dealing with the support center

Use this topic to understand how to report problems, and correspond with the IBM support center.

If you choose to contact the support center by telephone, your first contact at the support center is the call receipt operator, who takes initial details and puts your problem on a queue. You are then contacted by a support center representative, and your problem is taken from there.

Alternatively, you might have access to an electronic system for reporting problems to the support center. In this case, a support center representative will respond to your communication.

The support center needs to know as much as possible about your problem, so have the information ready before contacting them. If contacting them by telephone, it is a good idea to write the information about a problem reporting sheet such as the one shown in Figure 169 on page 1349.

There are two advantages of using a problem reporting sheet when contacting the IBM support center:

- In a telephone conversation, you will be better prepared to respond to the questions that you might be asked if you have all your findings before you on a sheet of paper.
- You can use the information for planning, organizing, and establishing priorities for controlling and resolving these problems.

Related concepts:

“What the support center needs to know”

“What happens next” on page 1351

What the support center needs to know

It is important that you understand what information the support center requires.

When you contact the support center, whether by telephone or electronically, you need to state the name of your organization and your *access code*. Your access code is a unique code authorizing you to use IBM PSS Software Support, and you provide it every time you contact the center. This information is used to access your customer profile, which contains information about your address, relevant contact names, telephone numbers, and details of the IBM products at your installation.

The support center needs to know whether this is a new problem, or a further communication regarding an existing one. If it is new, it is assigned a unique *incident number*. A *problem management record* (PMR) is opened on the RETAIN system, where all activity associated with your problem is recorded. The problem remains ‘open’ until you are in agreement with the support center that it has been resolved and can now be closed.

Make a note of the incident number on your own problem reporting sheet. The support center expects you to quote the incident number in all future communications connected with this problem.

If the problem is new to you, you need to state the source of the problem within your system software—that is, the program that seems to be the cause of the problem. Because you are reading this manual, it is likely that you have already identified WebSphere MQ for z/OS as the problem source. You also have to give the version and release number.

You need to give the *severity level* for the problem. Severity levels can be 1, 2, 3, or 4 and they have the following meanings:

Level 1

This indicates that you cannot use the system, and have a critical condition that needs immediate attention.

Level 2

This indicates that you can use the system, but that operation is severely restricted.

Level 3

This indicates that you can use the program, with limited functions, but the problem is not critical to your overall operation.

Level 4

This indicates that you have found a way to work around the problem; however, further action is required to correct the problem.

When deciding the severity of the problem, take care not to understate it, or to over state it. The support center procedures depend on the severity level so that the most appropriate use can be made of the center's skills and resources. A severity level 1 problem is normally dealt with immediately.

Next, you need to state a brief description of the problem. You might also be asked to quote the WebSphere MQ for z/OS symptom string, or to give any keywords associated with the problem. The primary keywords are ABEND, WAIT, LOOP, PERFM, INCORROUT, MSG, and DOC, corresponding exactly with the problem classification types used in “Building a keyword string” on page 1331. Strings containing other keywords are also useful. These are not predefined, and might include such items as a message or message number, an abend code, any parameters known to be associated with the problem, or, for example, STARTUP or INITIALIZATION.

The keywords are then used as search arguments on the RETAIN database, to see if your problem is a known one that has already been the subject of an *authorized program analysis report* (APAR).

Finally, let the support center know if any of the following events occurred before the problem appeared:

- Changes in the level of z/OS or licensed programs
- Regenerations
- PTFs applied
- Additional features used
- Application programs changed
- Unusual operator action

You might be asked to give values from a formatted dump or trace table, or to carry out some special activity, for example to set a trap, or to use trace with a specific type of selectivity, and then to report the results.

Note: You will be given guidance by the support center on how to obtain this information.

How your problem is then progressed depends on its nature. The representative who handles the problem gives you guidance on what is required from you. The possibilities are described in the next section.

Related concepts:

“What happens next”

What happens next

Use this topic to understand what happens after you contact the IBM support center.

Details of your problem are passed to the appropriate support group using the RETAIN problem management system. Your problem, assuming that it is one associated with WebSphere MQ for z/OS, is put on the WebSphere MQ for z/OS queue. The problems are dealt with in order of receipt and severity level.

At first, an IBM support center representative uses the keywords that you have provided to search the RETAIN database. If your problem is found to be one already known to IBM, and a fix has been devised for it, a *program temporary fix* (PTF) can be dispatched to you quickly. Alternatively, you might be asked to try running your installation using different settings.

If the RETAIN search is unsuccessful, you are asked to provide more information about your problem. Guidance on collecting and sending documentation to IBM is given in “Collecting documentation for the problem” on page 1352 and “Sending the documentation to the change team” on page 1353.

If the problem requires a change to the WebSphere MQ for z/OS code or documentation, an *authorized program analysis report* (APAR) is submitted. This is dealt with by the WebSphere MQ for z/OS Development Support Group or *change team* and provides a means of tracking the change.

It might be necessary to have several follow-up communications, depending on the complexity of the symptoms and your system environment. In every case, the actions taken by you and the support center are entered in the original PMR. The representative can then be acquainted with the full history of the problem before the next communication.

Collecting documentation for the problem

Use this topic to understand what documentation to collect for the support center.

As a rule, the documentation you need to submit for a problem includes all the material you need yourself to do problem determination. Some of the documentation is common to all WebSphere MQ for z/OS problems, and some is specific to particular types of problem.

Make sure that the problem you have described can be seen in the documentation you send. If the problem has ambiguous symptoms, you need to reveal the sequence of events leading to the failure. Tracing is valuable in this respect but you need to provide details that trace cannot give. You are encouraged to annotate your documentation, if your annotation is legible and does not cover up vital information. You can highlight any data in hardcopy you send, using transparent highlighting markers. You can also write notes in the margins, preferably using a red pen so that the notes are not overlooked.

Finally, note that if you send too little documentation, or if it is unreadable, the change team will have to return your problem marked “insufficient documentation”. It is, therefore, worthwhile preparing your documentation carefully and sending everything relevant to the problem.

The general documentation is described in this topic. However, these are only guidelines, you must find out from the IBM support center representative precisely what documentation you need to send for your specific problem.

General documentation needed for all problems with WebSphere MQ for z/OS

Here is a list of the general documentation you might be asked to submit for a PMR:

- Any hardcopy or softcopy illustrating the symptoms of the problem.
- The dump of the problem, see “Getting a dump” on page 1277.
- The appropriate SYS1.LOGREC records, see “SYS1.LOGREC information” on page 1292.
- The system log.
- A portion of the WebSphere MQ for z/OS job log.
- Trace records, see “Using trace for problem determination on z/OS” on page 1244.
- Trace information produced by the CICS or IMS adapter, see “Other types of trace” on page 1250.
- Buffer pool statistics, see the sections on using SMF and type 115 records in  *Configuring z/OS (WebSphere MQ V7.1 Installing Guide)*.
- Listings of relevant application programs.
- A list of PTFs and APARs applied.

Use the following sample JCL produce a list of all the PTFs, APARs, user modifications, and product code that has been installed in the global zone specified by SMPCSI:

```
//SMPELIST EXEC PGM=GIMSMP,REGION=4096K
//SMPCSI DD DSN=shlqua1.global.csi,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SMPCTL DD *
        SET BOUNDARY(GLOBAL) .
        LIST
            SYSMODS
            APARS
            FUNCTIONS
            PTFS
            USERMODS
        ALLZONES .
/*
```

- Dump of coupling facility administration structure, see Figure 159 on page 1279.
- Dump of coupling facility application structure, see Figure 159 on page 1279.
- Dump of other queue managers in queue-sharing group, see Figure 155 on page 1278.
- Definitions of the objects in your system.

You can find these by using the DISPLAY command for each type of object, for example:

```
DISPLAY QUEUE (*) ALL
```

Or, if you regularly make a backup of your object using the MAKEDEF feature of CSQUTIL COMMAND function, the output from that backup job.

- Your WebSphere MQ Db2 tables.

You can get these by using the sample job CSQ45STB in thlqual.SCSQPROCS to produce a report of all the Db2 tables used by WebSphere MQ.

Because of the size of SVC dumps in the cross memory environment, transfer the SYS1.DUMPxx data set to a tape or like device. You can use the PRDMP service aid program to transfer the SYS1.DUMPxx data set contents to another data set for archiving until the problem is resolved. Alternatively, your support center representative might give you the address of an FTP site where you can send your dump electronically.

Depending on the nature of the problem, the IBM support center might ask you to send the entire dump on tape. This allows the support center to extract any additional data needed to resolve the problem; for example, CSA, SQA, or the private storage area.

Sending the documentation to the change team

Use this topic to ensure that you have the correct documentation prepared for the support center.

When submitting documentation for your problem, your IBM support center will advise you on the most appropriate method to use. Contact the support center for details.

Each item submitted must have the following information attached and visible:

- The PMR number assigned by IBM
- A list of data sets on the tape (application source program, JCL, or data)
- A list of how the tape was made, including:
 - The exact JCL listing or the commands used
 - The recording mode and density
 - Tape labeling
 - The record format, logical record length, and block size used for each data set

When the change team receives the package, this is noted on your PMR record on the RETAIN system. The team then investigates the problem. Sometimes, they will ask for more documentation, perhaps specifying some trap you must apply to get it.

When you are satisfied that the problem is solved, a code is entered on RETAIN to close the PMR, and if necessary, you are provided with a fix.

You can inquire any time at your support center on how your PMR is progressing, particularly if it is a problem of high severity.

Resolving a problem

After a problem is confirmed an APAR can be raised and a PTF released. Use this topic to understand the APAR and PTF process.

An APAR

An *authorized program analysis report* (APAR) is the means by which a problem with an IBM program is documented, tracked, and corrected. It is also used to track problems with IBM documents.

An APAR is raised by the IBM change team when a new problem is reported for which a program or documentation change is required. It is separate to the PMR that is raised when you report first report the problem.

When the change team solves the problem, they might produce a local fix enabling you to get your system running properly again. Finally, a *program temporary fix* (PTF) is produced to replace the module in error, and the APAR is closed.

The APAR process

The first step in the APAR process is that an IBM support center representative enters your APAR into the RETAIN system. The APAR text contains a description of your problem. If you have found a means of getting round the problem, details of this are entered as well. Your name is also entered, so that the support center knows whom to contact if the change team needs to ask anything further about the APAR documentation.

When the APAR has been entered, you are given an APAR number. You must write this number on all the documentation you submit to the change team. This number is always associated with the APAR and its resolution and, if a code change is required, with the fix as well.

During the APAR process, the change team might ask you to test the fix on your system.

Lastly, you need to apply the PTF resulting from the APAR when it becomes available.

Applying the fix

When the change team have created a fix for your problem, they might want you to test it on your system.

When the team is confident that the fix is satisfactory, the APAR is certified and the APAR is closed.

Occasionally, the solution to the APAR requires a change to the documentation only. In some circumstances, the APAR might be closed with a classification code of **FIN**, which means that if there is a subsequent release of IBM WebSphere MQ for z/OS, a fix for this problem can be provided at this time.

The APAR becomes a PTF

If the solution requires a code change to the current release, when the APAR is closed the change is distributed as a PTF.

If you want a PTF to resolve a specific problem, you can order it explicitly by its PTF number through the IBM support center. Otherwise, you can wait for the PTF to be sent out on the standard distribution tape. For more information on migration PTF see the [🔗 WebSphere MQ Support, Migration PTFs web page](#).

Getting product fixes

A product fix might be available to resolve your problem. You can determine what fixes are available by launching a query from the IBM Support Assistant.

Before you begin

1. If you have not already installed the IBM Support Assistant, you can find instructions about how to do so in “Installing the IBM Support Assistant (ISA)” on page 1324.
2. Install the IBM Support Assistant WMQ plug-in by following the instructions in “Updating the IBM Support Assistant (ISA)” on page 1325.

About this task

To launch a query from the IBM Support Assistant:

1. Open the IBM Support Assistant from the Start menu by clicking **Programs > IBM Support Assistant > IBM Support Assistant**.
2. Click **Product Information, WMQ, Support page, Download**, then **Recommended fixes**. This Web page provides links to the latest available maintenance for the in-service products.

To receive weekly e-mail notifications about fixes and other news about IBM products, follow these steps.

Procedure

1. From the support site ( WebSphere MQ support Web page), click **Subscribe to this product** in the Notifications box on the page.
2. If you have already registered for "Notifications", skip to the next step. If you have not registered, click **register now** on the sign-in page and follow the on-screen instructions.
3. Sign in to "My notifications".
4. Click the **Subscribe** tab. A list of products families is shown.
5. In the Software column, click **WebSphere**. A list of products is shown.
6. Select the product for which you want to receive notifications (for example, **WebSphere MQ**), then click **Continue**.
7. Set options to determine what notifications you receive, how often you receive them, and to which folder they are saved, then click **Submit**.

Related concepts:

“Troubleshooting and support” on page 1149

“Troubleshooting overview” on page 1149

 Applying and removing maintenance (*WebSphere MQ V7.1 Installing Guide*)

Related tasks:

“Contacting IBM Software Support” on page 1346

“Searching knowledge bases” on page 1326

Recovering after failure

Follow a set of procedures to recover after a serious problem.

About this task

Use the recovery methods described here if you cannot resolve the underlying problem by using the diagnostic techniques described throughout the Troubleshooting and support section of the product

documentation. If your problem cannot be resolved by using these recovery techniques, contact your IBM Support Center.

Procedure

See the following links for instructions on how to recover from different types of failures:

- “Disk drive failures”
- “Damaged queue manager object” on page 1357
- “Damaged single object” on page 1358
- “Automatic media recovery failure” on page 1358

See the following links for instructions on how to recover from different types of failures on IBM WebSphere MQ for z/OS:

- “Shared queue problems” on page 1359
- “Active log problems” on page 1359
- “Archive log problems” on page 1365
- “BSDS problems” on page 1368
- “Page set problems” on page 1372
- “Coupling facility and Db2 problems” on page 1374
- “Problems with long-running units of work” on page 1377
- “IMS-related problems” on page 1377
- “Hardware problems” on page 1379

Related concepts:

“Troubleshooting and support” on page 1149

“Troubleshooting overview” on page 1149

“Making initial checks on Windows, UNIX and Linux systems” on page 1151



Backing up and restoring WebSphere MQ (*WebSphere MQ V7.1 Installing Guide*)



Planning for backup and recovery on z/OS (*WebSphere MQ V7.1 Installing Guide*)

Disk drive failures

You might have problems with a disk drive containing either the queue manager data, the log, or both. Problems can include data loss or corruption. The three cases differ only in the part of the data that survives, if any.

In *all* cases first check the directory structure for any damage and, if necessary, repair such damage. If you lose queue manager data, the queue manager directory structure might have been damaged. If so, re-create the directory tree manually before you restart the queue manager.

If damage has occurred to the queue manager data files, but not to the queue manager log files, then the queue manager will normally be able to restart. If any damage has occurred to the queue manager log files, then it is likely that the queue manager will not be able to restart.

Having checked for structural damage, there are a number of things you can do, depending on the type of logging that you use.

- **Where there is major damage to the directory structure or any damage to the log**, remove all the old files back to the QMgrName level, including the configuration files, the log, and the queue manager directory, restore the last backup, and restart the queue manager.

- **For linear logging with media recovery**, ensure that the directory structure is intact and restart the queue manager. If the queue manager restarts, check, using MQSC commands such as DISPLAY QUEUE, whether any other objects have been damaged. Recover those you find, using the `rcrmqobj` command. For example:

```
rcrmqobj -m QMgrName -t all *
```

where QMgrName is the queue manager being recovered. `-t all *` indicates that all damaged objects of any type are to be recovered. If only one or two objects have been reported as damaged, you can specify those objects by name and type here.

- **For linear logging with media recovery and with an undamaged log**, you might be able to restore a backup of the queue manager data leaving the existing log files and log control file unchanged. Starting the queue manager applies the changes from the log to bring the queue manager back to its state when the failure occurred.

This method relies on two things:

1. You must restore the checkpoint file as part of the queue manager data. This file contains the information determining how much of the data in the log must be applied to give a consistent queue manager.
2. You must have the oldest log file required to start the queue manager at the time of the backup, and all subsequent log files, available in the log file directory.

If this is not possible, restore a backup of both the queue manager data and the log, both of which were taken at the same time. This causes message integrity to be lost.

- **For circular logging**, if the queue manager log files are damaged, restore the queue manager from the latest backup that you have. Once you have restored the backup, restart the queue manager and check for damaged objects. However, because you do not have media recovery, you must find other ways of re-creating the damaged objects.

If the queue manager log files are not damaged, the queue manager will normally be able to restart. Following the restart you must identify all damaged objects, then delete and redefine them.

Damaged queue manager object

What to do if the queue manager reports a damaged object during normal operation.

There are two ways of recovering in these circumstances, depending on the type of logging you use:

- **For linear logging**, manually delete the file containing the damaged object and restart the queue manager. (You can use the `dspmqls` command to determine the real, file-system name of the damaged object.) Media recovery of the damaged object is automatic.
- **For circular logging**, restore the last backup of the queue manager data and log, and restart the queue manager.

There is a further option if you are using circular logging. For a damaged queue, or other object, delete the object and define the object again. In the case of a queue, this option does not allow you to recover any data on the queue.

Note: Restoring from backup is likely to be out of date, due to the fact that you must have your queue manager shutdown in order to get a clean backup of the queue files.

Damaged single object

If a single object is reported as damaged during normal operation, for linear logging you can re-create the object from its media image. However, for circular logging you cannot re-create a single object.

Automatic media recovery failure

If a local queue required for queue manager startup with a linear log is damaged, and the automatic media recovery fails, restore the last backup of the queue manager data and log and restart the queue manager.

Example recovery scenarios

Use this topic as a reference for various recovery scenario procedures.

This topic describes procedures for recovering WebSphere MQ after various error conditions. These error conditions are grouped in the following categories:

Table 153. Example recovery scenarios

Problem category	Problem	Where to look next
Shared queue problems	Conflicting definitions for both private and shared queues.	"Shared queue problems" on page 1359
Active log problems	<ul style="list-style-type: none">• Dual logging is lost.• Active log has stopped.• One or both copies of the active log data set are damaged.• Write errors on active log data set.• Active log is becoming full or is full.• Read errors on active log data set.	"Active log problems" on page 1359
Archive log problems	<ul style="list-style-type: none">• Insufficient DASD space to complete offloading active log data sets.• Offload task has terminated abnormally.• Archive data set allocation problem.  1• Read I/O errors on the archive data set during restart.	"Archive log problems" on page 1365
BSDS problems	<ul style="list-style-type: none">• Error opening BSDS.• Log content does not correspond with BSDS information.• Both copies of the BSDS are damaged.• Unequal time stamps.• Dual BSDS data sets are out of synchronization.• I/O error on BSDS.	"BSDS problems" on page 1368
Page set problems	<ul style="list-style-type: none">• Page set full.• A page set has an I/O error.	"Page set problems" on page 1372
coupling facility and Db2 problems	<ul style="list-style-type: none">• Storage medium full.• Db2 system fails.• Db2 data-sharing group fails.• Db2 and the coupling facility fail.	"Coupling facility and Db2 problems" on page 1374
Unit of work problems	A long-running unit of work is encountered.	"Problems with long-running units of work" on page 1377

Table 153. Example recovery scenarios (continued)

Problem category	Problem	Where to look next
IMS problems	<ul style="list-style-type: none"> • An IMS application terminates abnormally. • The IMS adapter cannot connect to WebSphere MQ. • IMS not operational. 	"IMS-related problems" on page 1377
Hardware problems	Media recovery procedures	"Hardware problems" on page 1379

Shared queue problems

Problems occur if WebSphere MQ discovers that a page set based queue, and a shared queue of the same name are defined.

Symptoms

WebSphere MQ issues the following message:

```
CSQI063E +CSQ1 QUEUE queue-name IS BOTH PRIVATE AND SHARED
```

During queue manager restart, WebSphere MQ discovered that a page set based queue and a shared queue of the same name coexist.

System action

Once restart processing has completed, any **MQOPEN** request to that queue name fails, indicating the coexistence problem.

System programmer action

None.

Operator action

Delete one version of the queue to allow processing of that queue name. If there are messages on the queue that must be kept, you can use the **MOVE QLOCAL** command to move them to the other queue.

Active log problems

Use this topic to resolve different problems with the active logs.

This topic covers the following active log problems:

- "Dual logging is lost"
- "Active log stopped" on page 1360
- "One or both copies of the active log data set are damaged" on page 1361
- "Write I/O errors on an active log data set" on page 1361
- "I/O errors occur while reading the active log" on page 1362
- "Active log is becoming full" on page 1364
- Active log is full

Dual logging is lost

Symptoms

WebSphere MQ issues the following message:

```
CSQJ004I +CSQ1 ACTIVE LOG COPY n INACTIVE, LOG IN SINGLE MODE,  
ENDRBA=...
```

Having completed one active log data set, WebSphere MQ found that the subsequent (COPY *n*) data sets were not offloaded or were marked stopped.

System action

WebSphere MQ continues in single mode until offloading has been completed, then returns to dual mode.

System programmer action

None.

Operator action

Check that the offload process is proceeding and is not waiting for a tape mount. You might need to run the print log map utility to determine the state of all data sets. You might also need to define additional data sets.

Active log stopped**Symptoms**

WebSphere MQ issues the following message:

```
CSQJ030E +CSQ1 RBA RANGE startriba TO endriba NOT AVAILABLE IN ACTIVE  
LOG DATA SETS
```

System action

The active log data sets that contain the RBA range reported in message CSQJ030E are unavailable to WebSphere MQ. The status of these logs is STOPPED in the BSDS. The queue manager terminates with a dump.

System programmer action

You must resolve this problem before restarting the queue manager. The log RBA range must be available for WebSphere MQ to be recoverable. An active log that is marked as STOPPED in the BSDS will never be reused or archived and this creates a hole in the log.

Look for messages that indicate why the log data set has stopped, and follow the instructions for those messages.

Modify the BSDS active log inventory to reset the STOPPED status. To do this, follow this procedure after the queue manager has terminated:

1. Use the print log utility (CSQJU004) to obtain a copy of the BSDS log inventory. This shows the status of the log data sets.
2. Use the DELETE function of the change log inventory utility (CSQJU003) to delete the active log data sets that are marked as STOPPED.
3. Use the NEWLOG function of CSQJU003 to add the active logs back into the BSDS inventory. The starting and ending RBA for each active log data set must be specified on the NEWLOG statement. (The correct values to use can be found from the print log utility report obtained in Step 1.)
4. Rerun CSQJU004. The active log data sets that were marked as STOPPED are now shown as NEW and NOT REUSABLE. These active logs will be archived in due course.
5. Restart the queue manager.

Note: If your queue manager is running in dual BSDS mode, you must update both BSDS inventories.

One or both copies of the active log data set are damaged

Symptoms

WebSphere MQ issues the following messages:

```
CSQJ102E +CSQ1 LOG RBA CONTENT OF LOG DATA SET DSNAME=...,
          STARTRBA=..., ENDRBA=...,
          DOES NOT AGREE WITH BSDS INFORMATION
CSQJ232E +CSQ1 OUTPUT DATA SET CONTROL INITIALIZATION PROCESS FAILED
```

System action

Queue manager startup processing is terminated.

System programmer action

If one copy of the data set is damaged, carry out these steps:

1. Rename the damaged active log data set and define a replacement data set.
2. Copy the undamaged data set to the replacement data set.
3. Use the change log inventory utility to:
 - Remove information relating to the damaged data set from the BSDS.
 - Add information relating to the replacement data set to the BSDS.
4. Restart the queue manager.

If both copies of the active log data sets are damaged, the current page sets are available, **and the queue manager shut down cleanly**, carry out these steps:

1. Rename the damaged active log data sets and define replacement data sets.
2. Use the change log records utility to:
 - Remove information relating to the damaged data set from the BSDS.
 - Add information relating to the replacement data set to the BSDS.
3. Rename the current page sets and define replacement page sets.
4. Use CSQUTIL (FORMAT and RESETPAGE) to format the replacement page sets and copy the renamed page sets to them. The RESETPAGE function also resets the log information in the replacement page sets.

If the queue manager did not shut down cleanly, you must either restore your system from a previous known point of consistency, or perform a cold start (described in “Reinitializing a queue manager” on page 325).

Operator action

None.

Write I/O errors on an active log data set

Symptoms

WebSphere MQ issues the following message:

```
CSQJ105E +CSQ1 csect-name LOG WRITE ERROR DSNAME=...,
LOGRBA=..., ERROR STATUS=ccccffss
```

System action

WebSphere MQ carries out these steps:

1. Marks the log data set that has the error as TRUNCATED in the BSDS.
2. Goes on to the next available data set.
3. If dual active logging is used, truncates the other copy at the same point.

The data in the truncated data set is offloaded later, as usual.

The data set will be reused on the next cycle.

System programmer action

None.

Operator action

If errors on this data set still exist, shut down the queue manager after the next offload process. Then use Access Method Services (AMS) and the change log inventory utility to add a replacement. (For instructions, see "Changing the BSDS" on page 294.)

I/O errors occur while reading the active log

Symptoms

WebSphere MQ issues the following message:

```
CSQJ106E +CSQ1 LOG READ ERROR DSNAME=..., LOGRBA=...,
ERROR STATUS=ccccffss
```

System action

This depends on when the error occurred:

- If the error occurs during the offload process, the process tries to read the RBA range from a second copy.
 - If no second copy exists, the active log data set is stopped.
 - If the second copy also has an error, only the original data set that triggered the offload process is stopped. The archive log data set is then terminated, leaving a gap in the archived log RBA range.
 - This message is issued:

```
CSQJ124E +CSQ1 OFFLOAD OF ACTIVE LOG SUSPENDED FROM
RBA xxxxxx TO RBA xxxxxx DUE TO I/O ERROR
```

- If the second copy is satisfactory, the first copy is not stopped.
- If the error occurs during recovery, WebSphere MQ provides data from specific log RBAs requested from another copy or archive. If this is unsuccessful, recovery does not succeed, and the queue manager terminates abnormally.
- If the error occurs during restart, if dual logging is used, WebSphere MQ continues with the alternative log data set, otherwise the queue manager ends abnormally.

System programmer action

Look for system messages, such as IEC prefixed messages, and try to resolve the problem using the recommended actions for these messages.

If the active log data set has been stopped, it is not used for logging. The data set is not deallocated; it is still used for reading. Even if the data set is not stopped, an active log data set that gives persistent errors should be replaced.

Operator action

None.

Replacing the data set

How you replace the data set depends on whether you are using single or dual active logging.

If you are using dual active logging:

1. Ensure that the data has been saved.
The data is saved on the other active log and this can be copied to a replacement active log.
2. Stop the queue manager and delete the data set with the error using Access Method Services.
3. Redefine a new log data set using Access Method Services DEFINE so that you can write to it. Use DFDSS or Access Method Services REPRO to copy the good log in to the redefined data set so that you have two consistent, correct logs again.
4. Use the change log inventory utility, CSQJU003, to update the information in the BSDS about the corrupted data set as follows:
 - a. Use the DELETE function to remove information about the corrupted data set.
 - b. Use the NEWLOG function to name the new data set as the new active log data set and give it the RBA range that was successfully copied.
You can run the DELETE and NEWLOG functions in the same job step. Put the DELETE statement before NEWLOG statement in the SYSIN input data set.
5. Restart the queue manager.

If you are using single active logging:

1. Ensure that the data has been saved.
2. Stop the queue manager.
3. Determine whether the data set with the error has been offloaded:
 - a. Use the CSQJU003 utility to list information about the archive log data sets from the BSDS.
 - b. Search the list for a data set with an RBA range that includes the RBA of the corrupted data set.
4. If the corrupted data set has been offloaded, copy its backup in the archive log to a new data set. Then, skip to step 6.
5. If an active log data set is stopped, an RBA is not offloaded. Use DFDSS or Access Method Services REPRO to copy the data from the corrupted data set to a new data set.
If further I/O errors prevent you from copying the entire data set, a gap occurs in the log.

Note: Queue manager restart will not be successful if a gap in the log is detected.

6. Use the change log inventory utility, CSQJU003, to update the information in the BSDS about the corrupted data set as follows:
 - a. Use the DELETE function to remove information about the corrupted data set.
 - b. Use the NEWLOG function to name the new data set as the new active log data set and to give it the RBA range that was successfully copied.
The DELETE and NEWLOG functions can be run in the same job step. Put the DELETE statement before NEWLOG statement in the SYSIN input data set.
7. Restart the queue manager.

Active log is becoming full

The active log can fill up for several reasons, for example, delays in offloading and excessive logging. If an active log runs out of space, this has serious consequences. When the active log becomes full, the queue manager halts processing until an offload process has been completed. If the offload processing stops when the active log is full, the queue manager can end abnormally. Corrective action is required before the queue manager can be restarted.

Symptoms

Because of the serious implications of an active log becoming full, the queue manager issues the following warning message when the last available active log data set is 5% full:

```
CSQJ110E +CSQ1 LAST COPYn ACTIVE LOG DATA SET IS nnn PERCENT FULL
```

and reissues the message after each additional 5% of the data set space is filled. Each time the message is issued, the offload process is started.

System action

Messages are issued and offload processing started. If the active log becomes full, further actions are taken. See "Active log is full"

System programmer action

Use the DEFINE LOG command to dynamically add further active log data sets. This permits WebSphere MQ to continue its normal operation while the error causing the offload problems is corrected. For more information about the DEFINE LOG command, see  DEFINE LOG (*WebSphere MQ V7.1 Reference*).

Active log is full

Symptoms

When the active log becomes full, the queue manager halts processing until an offload process has been completed. If the offload processing stops when the active log is full, the queue manager can end abnormally. Corrective action is required before the queue manager can be restarted.

WebSphere MQ issues the following message:

```
CSQJ111A +CSQ1 OUT OF SPACE IN ACTIVE LOG DATA SETS
```

and an offload process is started. The queue manager then halts processing until the offload process has been completed.

System action

WebSphere MQ waits for an available active log data set before resuming normal WebSphere MQ processing. Normal shut down, with either QUIESCE or FORCE, is not possible because the shutdown sequence requires log space to record system events related to shut down (for example, checkpoint records). If the offload processing stops when the active log is full, the queue manager stops with an X'6C6' abend; restart in this case requires special attention. For more details, see "Problem determination on z/OS" on page 1265.

System programmer action

You can provide additional active log data sets before restarting the queue manager. This permits WebSphere MQ to continue its normal operation while the error causing the offload process

problems is corrected. To add new active log data sets, use the change log inventory utility (CSQJU003) when the queue manager is not active. For more details about adding new active log data sets, see “Changing the BSDS” on page 294.

Consider increasing the number of logs by:

1. Making sure that the queue manager is stopped, then using the Access Method Services DEFINE command to define a new active log data set.
2. Defining the new active log data set in the BSDS using the change log inventory utility (CSQJU003).

When you restart the queue manager, offloading starts automatically during startup, and any work that was in progress when WebSphere MQ was forced to stop is recovered.

Operator action

Check whether the offload process is waiting for a tape drive. If it is, mount the tape. If you cannot mount the tape, force WebSphere MQ to stop by using the z/OS CANCEL command.

Archive log problems

Use this topic to investigate, and resolve problems with the archive logs.

This topic covers the following archive log problems:

- “Allocation problems”
- “Offload task terminated abnormally” on page 1366
- “Insufficient DASD space to complete offload processing” on page 1366
- “Read I/O errors on the archive data set while WebSphere MQ is restarting” on page 1368

Allocation problems

Symptoms

WebSphere MQ issues the following message:

```
CSQJ103E +CSQ1 LOG ALLOCATION ERROR DSNAME=dsname,  
        ERROR STATUS=eeeeiii, SMS REASON CODE=sss
```

z/OS dynamic allocation provides the ERROR STATUS. If the allocation was for offload processing, the following message is also displayed:

```
CSQJ115E +CSQ1 OFFLOAD FAILED, COULD NOT ALLOCATE AN ARCHIVE  
        DATA SET
```

System action

The following actions take place:

- If the input is needed for recovery, and recovery is not successful, and the queue manager ends abnormally.
- If the active log had become full and an offload task was scheduled but not completed, the offload task tries again the next time it is triggered. The active log does not reuse a data set that has not yet been archived.

System programmer action

None.

Operator action

Check the allocation error code for the cause of the problem, and correct it. Ensure that drives are available, and either restart or wait for the offload task to be retried. Be careful if a DFP/DFSMS

ACS user-exit filter has been written for an archive log data set, because this can cause a device allocation error when the queue manager tries to read the archive log data set.

Offload task terminated abnormally

Symptoms

No specific WebSphere MQ message is issued for write I/O errors.

Only a z/OS error recovery program message appears. If you get WebSphere MQ message CSQJ128E, the offload task has ended abnormally; see  CSQJ128E: LOG OFFLOAD TASK FAILED FOR ACTIVE LOG dsname (*WebSphere MQ V7.1 Reference*) for more information.

System action

The following actions take place:

- The offload task abandons the output data set; no entry is made in the BSDS.
- The offload task dynamically allocates a new archive and restarts offloading from the point at which it was previously triggered.
- If an error occurs on the new data set:
 - In dual archive mode, the following message is generated and the offload processing changes to single mode:

```
CSQJ114I +CSQ1 ERROR ON ARCHIVE DATA SET, OFFLOAD
CONTINUING WITH ONLY ONE ARCHIVE DATA SET BEING
GENERATED
```

- In single archive mode, the output data set is abandoned. Another attempt to process this RBA range is made the next time offload processing is triggered.
- The active log does not wrap around; if there are no more active logs, data is not lost.

System programmer action

None.

Operator action

Ensure that offload task is allocated on a reliable drive and control unit.

Insufficient DASD space to complete offload processing

Symptoms

While offloading the active log data sets to DASD, the process terminates unexpectedly. WebSphere MQ issues the following message:

```
CSQJ128E +CSQ1 LOG OFF-LOAD TASK FAILED FOR ACTIVE LOG nnnnn
```

The error is preceded by z/OS messages IEC030I, IEC031I, or IEC032I.

System action

WebSphere MQ de-allocates the data set on which the error occurred. If WebSphere MQ is running in dual archive mode, WebSphere MQ changes to single archive mode and continues the offload task. If the offload task cannot be completed in single archive mode, the active log data sets cannot be offloaded, and the state of the active log data sets remains NOT REUSABLE. Another attempt to process the RBA range of the abandoned active log data sets is made the next time the offload task is triggered.

System programmer action

Quiesce the queue manager (using the WebSphere MQ command STOP QMGR MODE(QUIESCE)) to restrict logging activity until the z/OS abend is resolved.

The most likely causes of these symptoms are:

- The size of the archive log data set is too small to contain the data from the active log data sets during offload processing. All the secondary space allocations have been used. This condition is normally accompanied by z/OS message IEC030I.

The most likely causes of these symptoms are:

- The size of the archive log data set is too small to contain the data from the active log data sets during offload processing. All the secondary space allocations have been used. This condition is normally accompanied by z/OS message IEC030I.

To solve the problem

1. Issue the command CANCEL <queue-manager name> to cancel the queue manager job
2. Increase the primary or secondary allocations (or both) for the archive log data set (in the CSQ6ARVP system parameters).

If the data to be offloaded is large, you can mount another online storage volume or make one available to WebSphere MQ.

3. Restart the queue manager.

- All available space on the DASD volumes to which the archive data set is being written has been exhausted. This condition is normally accompanied by z/OS message IEC032I.

To solve the problem, make more space available on the DASD volumes, or make another online storage volume available for WebSphere MQ.

- The primary space allocation for the archive log data set (as specified in the CSQ6ARVP system parameters) is too large to allocate to any available online DASD device. This condition is normally accompanied by z/OS message IEC032I.

To solve the problem, make more space available on the DASD volumes, or make another online storage volume available for WebSphere MQ. If this is not possible, you must adjust the value of PRIQTY in the CSQ6ARVP system parameters to reduce the primary allocation.

(For details, see  Using CSQ6ARVP (*WebSphere MQ V7.1 Installing Guide*.)

Note: If you reduce the primary allocation, you might have to increase the size of the secondary space allocation to avoid future abends.

- All available space on the DASD volumes to which the archive data set is being written has been exhausted. This condition is normally accompanied by z/OS message IEC032I.

To solve the problem, make more space available on the DASD volumes, or make another online storage volume available for WebSphere MQ.

- The primary space allocation for the archive log data set (as specified in the CSQ6ARVP system parameters) is too large to allocate to any available online DASD device. This condition is normally accompanied by z/OS message IEC032I.

To solve the problem, make more space available on the DASD volumes, or make another online storage volume available for WebSphere MQ. If this is not possible, you must adjust the value of PRIQTY in the CSQ6ARVP system parameters to reduce the primary allocation. (For

details, see  Using CSQ6ARVP (*WebSphere MQ V7.1 Installing Guide*.)

Note: If you reduce the primary allocation, you might have to increase the size of the secondary space allocation to avoid future abends.

Operator action

None.

Read I/O errors on the archive data set while WebSphere MQ is restarting

Symptoms

No specific WebSphere MQ message is issued; only the z/OS error recovery program message appears.

System action

This depends on whether a second copy exists:

- If a second copy exists, it is allocated and used.
- If a second copy does not exist, restart is not successful.

System programmer action

None.

Operator action

Try to restart, using a different drive.

BSDS problems

Use this topic to investigate, and resolve problems with BSDS.

For background information about the bootstrap data set (BSDS), see the  Planning your IBM WebSphere MQ environment on z/OS (*WebSphere MQ V7.1 Installing Guide*).

This topic describes the following BSDS problems:

- “Error occurs while opening the BSDS”
- “Log content does not agree with the BSDS information” on page 1369
- “Both copies of the BSDS are damaged” on page 1369
- “Unequal time stamps” on page 1370
- “Out of synchronization” on page 1371
- “I/O error” on page 1371

Normally, there are two copies of the BSDS, but if one is damaged, WebSphere MQ immediately changes to single BSDS mode. However, the damaged copy of the BSDS must be recovered before restart. If you are in single mode and damage the only copy of the BSDS, or if you are in dual mode and damage both copies, use the procedure described in “Recovering the BSDS” on page 298.

This section covers some of the BSDS problems that can occur at startup. Problems *not* covered here include:

- RECOVER BSDS command errors (messages CSQJ301E - CSQJ307I)
- Change log inventory utility errors (message CSQJ123E)
- Errors in the BSDS backup being dumped by offload processing (message CSQJ125E)

For more information about the problems, see  z/OS Messages and Codes (*WebSphere MQ V7.1 Product Overview Guide*).

Error occurs while opening the BSDS

Symptoms

WebSphere MQ issues the following message:

```
CSQJ100E +CSQ1 ERROR OPENING BSDSn DSNAME=..., ERROR STATUS=eei
```

where *eei* is the VSAM return code. For information about VSAM codes, see the *DFSMS/MVS Macro Instructions for Data Sets* documentation. For an explanation of this message, see  *z/OS Messages and Codes (WebSphere MQ V7.1 Product Overview Guide)*.

System action

During system initialization, the startup is terminated.

During a RECOVER BSDS command, the system continues in single BSDS mode.

System programmer action

None.

Operator action

Carry out these steps:

1. Run the print log map utility on both copies of the BSDS, and compare the lists to determine which copy is accurate or current.
2. Rename the data set that had the problem, and define a replacement for it.
3. Copy the accurate data set to the replacement data set, using Access Method Services.
4. Restart the queue manager.

Log content does not agree with the BSDS information**Symptoms**

WebSphere MQ issues the following message:

```
CSQJ102E +CSQ1 LOG RBA CONTENT OF LOG DATA SET DSNAME=...,  
STARTRBA=..., ENDRBA=...,  
DOES NOT AGREE WITH BSDS INFORMATION
```

This message indicates that the change log inventory utility was used incorrectly or that a down-level data set is being used.

System action

Queue manager startup processing is terminated.

System programmer action

None.

Operator action

Run the print log map utility and the change log inventory utility to print and correct the contents of the BSDS.

Both copies of the BSDS are damaged**Symptoms**

WebSphere MQ issues the following messages:

```
CSQJ107E +CSQ1 READ ERROR ON BSDS
          DSNAME=... ERROR STATUS=0874
CSQJ117E +CSQ1 REG8 INITIALIZATION ERROR READING BSDS
          DSNAME=... ERROR STATUS=0874
CSQJ119E +CSQ1 BOOTSTRAP ACCESS INITIALIZATION PROCESSING FAILED
```

System action

Queue manager startup processing is terminated.

System programmer action

Carry out these steps:

1. Rename the data set, and define a replacement for it.
2. Locate the BSDS associated with the most recent archive log data set, and copy it to the replacement data set.
3. Use the print log map utility to print the contents of the replacement BSDS.
4. Use the print log records utility to print a summary report of the active log data sets missing from the replacement BSDS, and to establish the RBA range.
5. Use the change log inventory utility to update the missing active log data set inventory in the replacement BSDS.
6. If dual BSDS data sets had been in use, copy the updated BSDS to the second copy of the BSDS.
7. Restart the queue manager.

Operator action

None.

Unequal time stamps**Symptoms**

WebSphere MQ issues the following message:

```
CSQJ120E +CSQ1 DUAL BSDS DATA SETS HAVE UNEQUAL TIME STAMPS,
          SYSTEM BSDS1=...,BSDS2=...,
          UTILITY BSDS1=...,BSDS2=...
```

The possible causes are:

- One copy of the BSDS has been restored. All information about the restored BSDS is down-level. The down-level BSDS has the lower time stamp.
- One of the volumes containing the BSDS has been restored. All information about the restored volume is down-level. If the volume contains any active log data sets or WebSphere MQ data, they are also down-level. The down-level volume has the lower time stamp.
- Dual logging has degraded to single logging, and you are trying to start without recovering the damaged log.
- The queue manager terminated abnormally after updating one copy of the BSDS but before updating the second copy.

System action

WebSphere MQ attempts to resynchronize the BSDS data sets using the more recent copy. If this fails, queue manager startup is terminated.

System programmer action

None.

Operator action

If automatic resynchronization fails, carry out these steps:

1. Run the print log map utility on both copies of the BSDS, compare the lists to determine which copy is accurate or current.
2. Rename the down-level data set and define a replacement for it.
3. Copy the good data set to the replacement data set, using Access Method Services.
4. If applicable, determine whether the volume containing the down-level BSDS has been restored. If it has been restored, all data on that volume, such as the active log data, is also down-level.

If the restored volume contains active log data and you were using dual active logs on separate volumes, you need to copy the current version of the active log to the down-level log data set. See "Recovering logs" on page 289 for details of how to do this.

Out of synchronization

Symptoms

WebSphere MQ issues the following message:

```
CSQJ122E +CSQ1 DUAL BSDS DATA SETS ARE OUT OF SYNCHRONIZATION
```

The system time stamps of the two data sets are identical. Differences can exist if operator errors occurred while the change log inventory utility was being used. (For example, the change log inventory utility was only run on one copy.) The change log inventory utility sets a private time stamp in the BSDS control record when it starts, and a close flag when it ends. WebSphere MQ checks the change log inventory utility time stamps and, if they are different, or they are the same but one close flag is not set, WebSphere MQ compares the copies of the BSDSs. If the copies are different, CSQJ122E is issued.

System action

Queue manager startup is terminated.

System programmer action

None.

Operator action

Carry out these steps:

1. Run the print log map utility on both copies of the BSDS, and compare the lists to determine which copy is accurate or current.
2. Rename the data set that had the problem, and define a replacement for it.
3. Copy the accurate data set to the replacement data set, using access method services.
4. Restart the queue manager.

I/O error

Symptoms

WebSphere MQ changes to single BSDS mode and issues the user message:

```
CSQJ126E +CSQ1 BSDS ERROR FORCED SINGLE BSDS MODE
```

This is followed by one of the following messages:

```
CSQJ107E +CSQ1 READ ERROR ON BSDS  
      DSNAME=... ERROR STATUS=...  
CSQJ108E +CSQ1 WRITE ERROR ON BSDS  
      DSNAME=... ERROR STATUS=...
```

System action

The BSDS mode changes from dual to single.

System programmer action

None.

Operator action

Carry out these steps:

1. Use Access Method Services to rename or delete the damaged BSDS and to define a new BSDS with the same name as the BSDS that had the error. Example control statements can be found in job CSQ4BREC in thlqual.SCSQPROC.
2. Issue the WebSphere MQ command RECOVER BSDS to make a copy of the good BSDS in the newly allocated data set and reinstate dual BSDS mode. See also "Recovering the BSDS" on page 298.

Page set problems

Use this topic to investigate, and resolve problems with the page sets.

This topic covers the problems that you might encounter with page sets:

- "Page set I/O errors" describes what happens if a page set is damaged.
- "Page set full" on page 1373 describes what happens if there is not enough space on the page set for any more MQI operations.

Page set I/O errors

Problem

A page set has an I/O error.

Symptoms

This message is issued:

```
CSQP004E +CSQ1 csect-name I/O ERROR STATUS ret-code  
PSID psid RBA rba
```

System action

The queue manager terminates abnormally.

System programmer action

None.

Operator action

Repair the I/O error cause.

If none of the page sets are damaged, restart the queue manager. WebSphere MQ automatically restores the page set to a consistent state from the logs.

If one or more page sets are damaged:

1. Rename the damaged page sets and define replacement page sets.
2. Copy the most recent backup page sets to the replacement page sets.
3. Restart the queue manager. WebSphere MQ automatically applies any updates that are necessary from the logs.

You cannot restart the queue manager if page set zero is not available. If one of the other page sets is not available, you can comment out the page set DD statement in the queue manager start-up JCL procedure. This lets you defer recovery of the defective page set, enabling other users to continue accessing WebSphere MQ.

When you add the page set back to the JCL procedure, system restart reads the log from the point where the page set was removed from the JCL to the end of the log. This procedure could take a long time if numerous data has been logged.

A reason code of MQRC_PAGESET_ERROR is returned to any application that tries to access a queue defined on a page set that is not available.

When you have restored the defective page set, restore its associated DD statement and restart the queue manager.

The operator actions described here are only possible if all log data sets are available. If your log data sets are lost or damaged, see Restarting if you have lost your log data sets.

Page set full

Problem

There is not enough space on a page set for one of the following:

- MQPUT or MQPUT1 calls to be completed
- Object manipulation commands to be completed (for example, DEFINE QLOCAL)
- MQOPEN calls for dynamic queues to be completed

Symptoms

The request fails with reason code MQRC_STORAGE_MEDIUM_FULL. The queue manager cannot complete the request because there is not enough space remaining on the page set.

Reason code MQRC_STORAGE_MEDIUM_FULL can occur even when the page set expand attribute is set to EXPAND(USER). Before the reason code MQRC_STORAGE_MEDIUM_FULL is returned to the application code, the queue manager will attempt to expand the page set and retry the API request. On a heavily loaded system it is possible that the expanded storage can be used by other IO operations before the retry of the API. See "Managing page sets" on page 300.

The cause of this problem could be messages accumulating on a transmission queue because they cannot be sent to another system.

System action

Further requests that use this page set are blocked until enough messages are removed or objects deleted to make room for the new incoming requests.

Operator action

Use the WebSphere MQ command DISPLAY USAGE PSID(*) to identify which page set is full.

System programmer action

You can either enlarge the page set involved or reduce the loading on that page set by moving queues to another page set. See "Managing page sets" on page 300 for more information about these tasks. If the cause of the problem is messages accumulating on the transmission queue, consider starting distributed queuing to transmit the messages.

Coupling facility and Db2 problems

Use this topic to investigate, and resolve problems with the coupling facility, and Db2.

This section covers the problems that you might encounter with the coupling facility and Db2:

- “Storage medium full”
- “A Db2 system fails”
- “A Db2 data-sharing group fails” on page 1375
- “Db2 and the coupling facility fail” on page 1376

Storage medium full

Problem

A coupling facility structure is full.

Symptoms

If a queue structure becomes full, return code MQRC_STORAGE_MEDIUM_FULL is returned to the application.

If the administration structure becomes full, the exact symptoms depend on which processes experience the error, they might range from no responses to CMDSCOPE(GROUP) commands, to queue manager failure as a result of problems during commit processing.

System programmer action

You can use WebSphere MQ to inhibit **MQPUT** operations to some of the queues in the structure to prevent applications from writing more messages, start more applications to get messages from the queues, or quiesce some of the applications that are putting messages to the queue.

Alternatively you can use XES facilities to alter the structure size in place. The following z/OS command alters the size of the structure:

```
SETXCF START,ALTER,STRNAME=structure-name,SIZE=newsiz
```

where *newsiz* is a value that is less than the value of MAXSIZE specified on the CFRM policy for the structure, but greater than the current coupling facility size.

You can monitor the utilization of a coupling facility structure with the DISPLAY CFSTATUS command.

A Db2 system fails

If a Db2 subsystem that WebSphere MQ is connected to fails, WebSphere MQ attempts to reconnect to the subsystem, and continue working. If you specified a Db2 group attach name in the QSGDATA parameter of the CSQ6SYSP system parameter module, WebSphere MQ reconnects to another active Db2 that is a member of the same data-sharing group as the failed Db2, if one is available on the same z/OS image.

There are some queue manager operations that do not work while WebSphere MQ is not connected to Db2. These are:

- Deleting a shared queue or group object definition.
- Altering, or issuing **MQSET** on, a shared queue or group object definition. The restriction of **MQSET** on shared queues means that operations such as triggering or the generation of performance events do not work correctly.
- Defining new shared queues or group objects.
- Displaying shared queues or group objects.
- Starting, stopping, or other actions for shared channels.

- Reading the shared queue definition from Db2 the first time that the shared queue is open by issuing an MQOPEN.

Other WebSphere MQ API operations continue to function as normal for shared queues, and all WebSphere MQ operations can be performed against the queue manager private versions (COPY objects) built from GROUP objects. Similarly, any shared channels that are running continue normally until they end or have an error, when they go into retry state.

When WebSphere MQ reconnects to Db2, resynchronization is performed between the queue manager and Db2. This involves notifying the queue manager of new objects that have been defined in Db2 while it was disconnected (other queue managers might have been able to continue working as normal on other z/OS images through other Db2 subsystems), and updating object attributes of shared queues that have changed in Db2. Any shared channels in retry state are recovered.

If a Db2 fails, it might have owned locks on Db2 resources at the time of failure. In some cases, this might make certain WebSphere MQ objects unavailable to other queue managers that are not otherwise affected. To resolve this, restart the failed Db2 so that it can perform recovery processing and release the locks.

A Db2 data-sharing group fails

If an entire Db2 data-sharing group fails, recovery might be to the time of failure, or to a previous point in time.

In the case of recovery to the point of failure, WebSphere MQ reconnects when Db2 has been recovered, the resynchronization process takes places, and normal queue manager function is resumed.

However, if Db2 is recovered to a previous point in time, there might be inconsistencies between the actual queues in the coupling facility structures and the Db2 view of those queues. For example, at the point in time Db2 is recovered to, a queue existed that has since been deleted and its location in the coupling facility structure reused by the definition of a new queue that now contains messages.

If you find yourself in this situation, you must stop all the queue managers in the queue-sharing group, clear out the coupling facility structures, and restart the queue managers. You must then use WebSphere MQ commands to define any missing objects. To do this, use the following procedure:

1. Prevent WebSphere MQ from reconnecting to Db2 by starting Db2 in utility mode, or by altering security profiles.
2. If you have any important messages on shared queues, you might be able to offload them using the COPY function of the CSQUTIL utility program, but this might not work.
3. Terminate all queue managers.
4. Use the following z/OS command to clear all structures:

```
SETXCF FORCE,STRUCTURE,STRNAME=
```

5. Restore Db2 to a historical point in time.
6. Reestablish queue manager access to Db2.
7. Restart the queue managers.
8. Recover the WebSphere MQ definitions from backup copies.
9. Reload any offloaded messages to the shared queues.

When the queue managers restart, they attempt to resynchronize local COPY objects with the Db2 GROUP objects. This might cause WebSphere MQ to attempt to do the following:

- Create COPY objects for old GROUP objects that existed at the point in time Db2 has recovered to.
- Delete COPY objects for GROUP objects that were created since the point in time Db2 has recovered to and so do not exist in the database.

The DELETE of COPY objects is attempted with the NOPURGE option, so it fails for queue managers that still have messages on these COPY queues.

Db2 and the coupling facility fail

If the coupling facility fails, the queue manager might fail, and Db2 will also fail if it is using this coupling facility.

Recover Db2 using Db2 recovery procedures. When Db2 has been restarted, you can restart the queue managers. The CF administration structure will also have failed, but this is rebuilt by restarting all the queue managers within the queue-sharing group.

If a single application structure within the coupling facility suffers a failure, the effect on the queue manager depends on the level of the queue manager and the CFLEVEL of the failed CF structure:

- If the CF application structure is CFLEVEL(3) or higher and RECOVER is set to YES, it will not be usable until you recover the CF structure by issuing an MQSC RECOVER CFSTRUCT command to the queue manager that will do the recovery. You can specify a single CF structure to be recovered, or you can recover several CF structures simultaneously. The queue manager performing the recovery locates the relevant backups on all the other queue managers' logs using the data in Db2 and the bootstrap data sets. The queue manager replays these backups in the correct time sequence across the queue sharing group, from just before the last backup through to the point of failure. If a recoverable application structure has failed, any further application activity is prevented until the structure has been recovered. If the administration structure has also failed, all the queue managers in the queue-sharing group must be started before the RECOVER CFSTRUCT command can be issued. All queue managers can continue working with local queues and queues in other CF structures during recovery of a failed CF structure.
- If the CF application structure is CFLEVEL(3) or higher and RECOVER is set to NO, the structure is automatically reallocated by the next **MQOPEN** request performed on a queue defined in the structure. All messages are lost, as the structure can only contain non-persistent messages.
- If the CF application structure has a CFLEVEL less than 3, the queue manager fails. On queue manager restart, peer recovery attempts to connect to the structure, detect that the structure has failed and allocate a new version of the structure. All messages on shared queues that were in CF structures affected by the coupling facility failure are lost.

Since IBM WebSphere MQ Version 7.1, queue managers in queue-sharing-groups have been able to tolerate loss of connectivity to coupling facility structures without failing. If the structure has experienced a connection failure, attempts are made to rebuild the structure in another coupling facility with better connectivity in order to regain access to shared queues as soon as possible.

See  WebSphere MQ Script (MQSC) Command Reference (*WebSphere MQ V7.1 Reference*) for details of the RECOVER CFSTRUCT command.

Problems with long-running units of work

Use this topic to investigate, and resolve problems with long-running units of work.

This topic explains what to do if you encounter a long-running unit of work during restart. In this context, this means a unit of work that has been active for a long time (possibly days or even weeks) so that the origin RBA of the unit of work is outside the scope of the current active logs. This means that restart could take a long time, because all the log records relating to the unit of work have to be read, which might involve reading archive logs.

Old unit of work found during restart

Problem

A unit of work with an origin RBA that predates the oldest active log has been detected during restart.

Symptoms

WebSphere MQ issues the following message:

```
CSQR020I +CSQ1 OLD UOW FOUND
```

System action

Information about the unit of work is displayed, and message CSQR021D is issued, requesting a response from the operator.

System programmer action

None.

Operator action

Decide whether to commit the unit of work or not. If you choose not to commit the unit of work, it is handled by normal restart recovery processing. Because the unit of work is old, this is likely to involve using the archive log, and so takes longer to complete.

IMS-related problems

Use this topic to investigate, and resolve problems with IMS and WebSphere MQ..

This topic includes plans for the following problems that you might encounter in the IMS environment:

- “IMS cannot connect to WebSphere MQ”
- “IMS application problem” on page 1378
- “IMS is not operational” on page 1378

IMS cannot connect to WebSphere MQ

Problem

The IMS adapter cannot connect to WebSphere MQ.

Symptoms

IMS remains operative. The IMS adapter issues these messages for control region connect:

- CSQQ001I
- CSQQ002E
- CSQQ003E
- CSQQ004E
- CSQQ005E
- CSQQ007E

For details, see the  z/OS Messages and Codes (*WebSphere MQ V7.1 Reference*) documentation.

If an IMS application program tries to access WebSphere MQ while the IMS adapter cannot connect, it can either receive a completion code and reason code, or terminate abnormally. This depends on the value of the REO option in the SSM member of IMS PROCLIB.

System action

All connection errors are also reported in the IMS message DFS3611.

System programmer action

None.

Operator action

Analyze and correct the problem, then restart the connection with the IMS command:

```
/START SUBSYS subsysname
```

IMS requests the adapter to resolve in-doubt units of recovery.

IMS application problem**Problem**

An IMS application terminates abnormally.

Symptoms

The following message is sent to the user's terminal:

```
DFS555I TRANSACTION tran-id ABEND abcode  
MSG IN PROCESS: message data:
```

where *tran-id* represents any IMS transaction that is terminating abnormally and *abcode* is the abend code.

System action

IMS requests the adapter to resolve the unit of recovery. IMS remains connected to WebSphere MQ.

System programmer action

None.

Operator action

As indicated in message DFS554A on the IMS master terminal.

IMS is not operational**Problem**

IMS is not operational.

Symptoms

More than one symptom is possible:

- IMS waits or loops
WebSphere MQ cannot detect a wait or loop in IMS, so you must find the origin of the wait or loop. This can be IMS, IMS applications, or the IMS adapter.
- IMS terminates abnormally.
 - See the manuals *IMS/ESA Messages and Codes* and *IMS/ESA Failure Analysis Structure Tables* for more information.

- If threads are connected to WebSphere MQ when IMS terminates, WebSphere MQ issues message CSQ3201E. This message indicates that WebSphere MQ end-of-task (EOT) routines have been run to clean up and disconnect any connected threads.

System action

WebSphere MQ detects the IMS error and:

- Backs out in-flight work.
- Saves in-doubt units of recovery to be resolved when IMS is reconnected.

System programmer action

None.

Operator action

Resolve and correct the problem that caused IMS to terminate abnormally, then carry out an emergency restart of IMS. The emergency restart:

- Backs out in-flight transactions that changed IMS resources.
- Remembers the transactions with access to WebSphere MQ that might be in doubt.

You might need to restart the connection to WebSphere MQ with the IMS command:

```
/START SUBSYS subsysname
```

During startup, IMS requests the adapter to resolve in-doubt units of recovery.

Hardware problems

Use this topic as a starting point to investigate hardware problems.

If a hardware error causes data to be unreadable, WebSphere MQ can still be recovered by using the *media recovery* technique:

1. To recover the data, you need a backup copy of the data. Use DFDSS or Access Method Services REPRO regularly to make a copy of your data.
2. Reinstall the most recent backup copy.
3. Restart the queue manager.

The more recent your backup copy, the more quickly your subsystem can be made available again.

When the queue manager restarts, it uses the archive logs to reinstate changes made since the backup copy was taken. You must keep sufficient archive logs to enable WebSphere MQ to reinstate the changes fully. Do not delete archive logs until there is a backup copy that includes all the changes in the log.

Reason codes

You can use the following messages and reason codes to help you solve problems with your IBM WebSphere MQ components or applications.

-  Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)
- “API completion and reason codes” on page 1380
- “PCF reason codes” on page 1590
- “Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes” on page 1672
- “WCF custom channel exceptions” on page 1676
-  IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

API completion and reason codes

For each call, a completion code and a reason code are returned by the queue manager or by an exit routine, to indicate the success or failure of the call.

For more information about the WebSphere MQ API, see  Developing applications (*WebSphere MQ V7.1 Programming Guide*), and the reference information in  Developing applications reference (*WebSphere MQ V7.1 Reference*).

For a full list and explanation of the API reason codes, see “API reason codes.”

API completion codes

The following is a list of the completion codes (MQCC) returned by WebSphere MQ

0: Successful completion (MQCC_OK)

The call completed fully; all output parameters have been set.

The *Reason* parameter always has the value MQRC_NONE in this case.

1: Warning (partial completion) (MQCC_WARNING)

The call completed partially. Some output parameters might have been set in addition to the *CompCode* and *Reason* output parameters.

The *Reason* parameter gives additional information.

2: Call failed (MQCC_FAILED)

The processing of the call did not complete, and the state of the queue manager is normally unchanged; exceptions are specifically noted. Only the *CompCode* and *Reason* output parameters have been set; all other parameters are unchanged.

The reason might be a fault in the application program, or it might be a result of some situation external to the program, for example the application's authority might have been revoked. The *Reason* parameter gives additional information.

Related concepts:

 IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

Related reference:

 Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)

“PCF reason codes” on page 1590

“Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes” on page 1672

“WCF custom channel exceptions” on page 1676

API reason codes

The reason code parameter (*Reason*) is a qualification to the completion code parameter (*CompCode*).

If there is no special reason to report, MQRC_NONE is returned. A successful call returns MQCC_OK and MQRC_NONE.

If the completion code is either MQCC_WARNING or MQCC_FAILED, the queue manager always reports a qualifying reason; details are given under each call description.

Where user exit routines set completion codes and reasons, they should adhere to these rules. In addition, any special reason values defined by user exits should be less than zero, to ensure that they do not conflict with values defined by the queue manager. Exits can set reasons already defined by the queue manager, where these are appropriate.

Reason codes also occur in:

- The *Reason* field of the MQDLH structure
- The *Feedback* field of the MQMD structure

In addition to the following list of reason codes, in numeric order, providing detailed information to help you understand them, including:

- An explanation of the circumstances that have caused the code to be raised
- The associated completion code
- Suggested programmer actions in response to the code

0 (0000) (RC0): MQRC_NONE:

Explanation

The call completed normally. The completion code (*CompCode*) is MQCC_OK.

Completion Code

MQCC_OK

Programmer response

None.

900 (0384) (RC900): MQRC_APPL_FIRST:

Explanation

This is the lowest value for an application-defined reason code returned by a data-conversion exit. Data-conversion exits can return reason codes in the range MQRC_APPL_FIRST through MQRC_APPL_LAST to indicate particular conditions that the exit has detected.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

As defined by the writer of the data-conversion exit.

999 (03E7) (RC999): MQRC_APPL_LAST:

Explanation

This is the highest value for an application-defined reason code returned by a data-conversion exit. Data-conversion exits can return reason codes in the range MQRC_APPL_FIRST through MQRC_APPL_LAST to indicate particular conditions that the exit has detected.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

As defined by the writer of the data-conversion exit.

2001 (07D1) (RC2001): MQRC_ALIAS_BASE_Q_TYPE_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued specifying an alias queue as the destination, but the *BaseQName* in the alias queue definition resolves to a queue that is not a local queue, a local definition of a remote queue, or a cluster queue,

or,

a queue in a distribution list contains an alias queue that is pointing to a topic object.

Completion Code

MQCC_FAILED

Programmer response

Correct the queue definitions.

2002 (07D2) (RC2002): MQRC_ALREADY_CONNECTED:

Explanation

An MQCONN or MQCONNX call was issued, but the application is already connected to the queue manager.

- On z/OS, this reason code occurs for batch and IMS applications only; it does not occur for CICS applications.
- On UNIX, IBM i, Linux and Windows, this reason code occurs if the application attempts to create a nonshared handle when a nonshared handle exists for the thread. A thread can have no more than one nonshared handle.
- On UNIX, IBM i, Linux and Windows, this reason code occurs if an MQCONN call is issued from within an MQ channel exit, API Crossing Exit, or Async Consume Callback function, and a shared hConn is bound to this thread.
- On UNIX, IBM i, Linux and Windows, this reason code occurs if an MQCONNX call that does not specify one of the MQCNO_HANDLE_SHARE_* options is issued from within an MQ channel exit, API Crossing Exit, or Async Consume Callback function, and a shared hConn is bound to this thread.
- On Windows, MTS objects do not receive this reason code, as additional connections to the queue manager are permitted.

Completion Code

MQCC_WARNING

Programmer response

None. The *hConn* parameter returned has the same value as was returned for the previous MQCONN or MQCONNX call.

An MQCONN or MQCONNX call that returns this reason code does *not* mean that an additional MQDISC call must be issued to disconnect from the queue manager. If this reason code is returned

because the application has been called in a situation where the MQCONN has already been done, do *not* issue a corresponding MQDISC, because this causes the application that issued the original MQCONN or MQCONNX call to be disconnected as well.

2003 (07D3) (RC2003): MQRC_BACKED_OUT:

Explanation

The current unit of work encountered an unrecoverable error or was backed out. This reason code is issued in the following cases:

- On an MQCMIT or MQDISC call, when the commit operation fails and the unit of work is backed out. All resources that participated in the unit of work are returned to their state at the start of the unit of work. The MQCMIT or MQDISC call completes with MQCC_WARNING in this case.
 - On z/OS, this reason code occurs only for batch applications.
- On an MQGET, MQPUT, or MQPUT1 call that is operating within a unit of work, when the unit of work already encountered an error that prevents the unit of work from being committed (for example, when the log space is exhausted). The application must issue the appropriate call to back out the unit of work. (For a unit of work that is coordinated by the queue manager, this call is the MQBACK call, although the MQCMIT call has the same effect in these circumstances.) The MQGET, MQPUT, or MQPUT1 call completes with MQCC_FAILED in this case.
 - On z/OS, this case does not occur.
- On an asynchronous consumption callback (registered by an MQCB call), the unit of work is backed out and the asynchronous consumer should call MQBACK.
 - On z/OS, this case does not occur.
- For the IBM WebSphere MQ client on HP Integrity NonStop Server using TMF, this return code can occur:
 - For MQGET, MQPUT, and MQPUT1 calls, if you have an active transaction that is being coordinated by TMF, but the IBM WebSphere MQ part of the transaction is rolled back because of inactivity on the transaction.
 - If the TMF/Gateway detects that TMF is rolling back the current transaction before the application finishes with it.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Check the returns from previous calls to the queue manager. For example, a previous MQPUT call might have failed.

2004 (07D4) (RC2004): MQRC_BUFFER_ERROR:

Explanation

The *Buffer* parameter is not valid for one of the following reasons:

- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The parameter pointer points to storage that cannot be accessed for the entire length specified by *BufferLength*.
- For calls where *Buffer* is an output parameter: the parameter pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2005 (07D5) (RC2005): MQRC_BUFFER_LENGTH_ERROR:

Explanation

The *BufferLength* parameter is not valid, or the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason can also be returned to an MQ MQI client program on the MQCONN or MQCONNX call if the negotiated maximum message size for the channel is smaller than the fixed part of any call structure.

This reason should also be returned by the MQZ_ENUMERATE_AUTHORITY_DATA installable service component when the *AuthorityBuffer* parameter is too small to accommodate the data to be returned to the invoker of the service component.

This reason code can also be returned when a zero length multicast message has been supplied where a positive length is required.

Completion Code

MQCC_FAILED

Programmer response

Specify a value that is zero or greater. For the mqAddString and mqSetString calls, the special value MQBL_NULL_TERMINATED is also valid.

2006 (07D6) (RC2006): MQRC_CHAR_ATTR_LENGTH_ERROR:

Explanation

CharAttrLength is negative (for MQINQ or MQSET calls), or is not large enough to hold all selected attributes (MQSET calls only). This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Specify a value large enough to hold the concatenated strings for all selected attributes.

2007 (07D7) (RC2007): MQRC_CHAR_ATTRS_ERROR:

Explanation

CharAttrs is not valid. The parameter pointer is not valid, or points to read-only storage for MQINQ calls or to storage that is not as long as implied by *CharAttrLength*. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2008 (07D8) (RC2008): MQRC_CHAR_ATTRS_TOO_SHORT:

Explanation

For MQINQ calls, *CharAttrLength* is not large enough to contain all of the character attributes for which MQCA_* selectors are specified in the *Selectors* parameter.

The call still completes, with the *CharAttrs* parameter string filled in with as many character attributes as there is room for. Only complete attribute strings are returned: if there is insufficient space remaining to accommodate an attribute in its entirety, that attribute and subsequent character attributes are omitted. Any space at the end of the string not used to hold an attribute is unchanged.

An attribute that represents a set of values (for example, the namelist *Names* attribute) is treated as a single entity—either all of its values are returned, or none.

Completion Code

MQCC_WARNING

Programmer response

Specify a large enough value, unless only a subset of the values is needed.

2009 (07D9) (RC2009): MQRC_CONNECTION_BROKEN:

Explanation

Connection to the queue manager has been lost. This can occur because the queue manager has ended. If the call is an MQGET call with the MQGMO_WAIT option, the wait has been canceled. All connection and object handles are now invalid.

For MQ MQI client applications, it is possible that the call did complete successfully, even though this reason code is returned with a *CompCode* of MQCC_FAILED.

Completion Code

MQCC_FAILED

Programmer response

Applications can attempt to reconnect to the queue manager by issuing the MQCONN or MQCONNX call. It might be necessary to poll until a successful response is received.

- On z/OS for CICS applications, it is not necessary to issue the MQCONN or MQCONNX call, because CICS applications are connected automatically.

Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2010 (07DA) (RC2010): MQRC_DATA_LENGTH_ERROR:

Explanation

The *DataLength* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason can also be returned to an MQ MQI client program on the MQGET, MQPUT, or MQPUT1 call, if the *BufferLength* parameter exceeds the maximum message size that was negotiated for the client channel.

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

If the error occurs for an MQ MQI client program, also check that the maximum message size for the channel is big enough to accommodate the message being sent; if it is not big enough, increase the maximum message size for the channel.

2011 (07DB) (RC2011): MQRC_DYNAMIC_Q_NAME_ERROR:

Explanation

On the MQOPEN call, a model queue is specified in the *ObjectName* field of the *ObjDesc* parameter, but the *DynamicQName* field is not valid, for one of the following reasons:

- *DynamicQName* is completely blank (or blank up to the first null character in the field).
- Characters are present that are not valid for a queue name.
- An asterisk is present beyond the 33rd position (and before any null character).
- An asterisk is present followed by characters that are not null and not blank.

This reason code can also sometimes occur when a server application opens the reply queue specified by the *ReplyToQ* and *ReplyToQMgr* fields in the MQMD of a message that the server has just received. In this case the reason code indicates that the application that sent the original message placed incorrect values into the *ReplyToQ* and *ReplyToQMgr* fields in the MQMD of the original message.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid name.

2012 (07DC) (RC2012): MQRC_ENVIRONMENT_ERROR:

Explanation

The call is not valid for the current environment.

- On z/OS, one of the following reasons apply:
 - An MQCONN or MQCONNX call was issued, but the application has been linked with an adapter that is not supported in the environment in which the application is running. For example, when the application is linked with the MQ RRS adapter, but the application is running in a Db2 Stored Procedure address space. RRS is not supported in this environment. Stored Procedures that use the MQ RRS adapter must run in a Db2 WLM-managed Stored Procedure address space.
 - An MQCMIT or MQBACK call was issued, but the application has been linked with the RRS batch adapter CSQBRSTB. This adapter does not support the MQCMIT and MQBACK calls.
 - An MQCMIT or MQBACK call was issued in the CICS or IMS environment.
 - The RRS subsystem is not operational on the z/OS system that ran the application.
 - An MQCTL call with MQOP_START or an MQCB call registering an Event Listener was issued, but the application is not allowed to create a POSIX thread.
 - A WebSphere MQ classes for Java application has instantiated an MQQueueManager object using the CLIENT transport. The z/OS environment only supports the use of the BINDINGS transport.
- On HP OpenVMS, IBM i, HP Integrity NonStop Server, UNIX systems, and Windows, one of the following applies:
 - The application is linked to the wrong libraries (threaded or nonthreaded).
 - An MQBEGIN, MQCMIT, or MQBACK call was issued, but an external unit-of-work manager is in use. For example, this reason code occurs on Windows when an MTS object is running as a DTC transaction. This reason code also occurs if the queue manager does not support units of work.
 - The MQBEGIN call was issued in an MQ MQI client environment.
 - An MQXCLWLN call was issued, but the call did not originate from a cluster workload exit.
 - An MQCONNX call was issued specifying the option MQCNO_HANDLE_SHARE_NONE on an MQ channel exit, an API Exit, or a Callback function. The reason code occurs only if a shared *hConn* is bound to the application thread.
 - A IBM WebSphere MQ Object is unable to connect fastpath.
 - A WebSphere MQ classes for Java application has created an MQQueueManager object that uses the CLIENT transport, and then called MQQueueManager.begin(). This method can only be called on MQQueueManager objects that use the BINDINGS transport.
- On Windows, when using the managed .NET client, an attempt was made to use one of the unsupported features:
 - Unmanaged channel exits
 - Secure Sockets Layer (SSL)
 - XA Transactions
 - Communications other than TCP/IP
 - Channel compression
- On Solaris, if you install IBM WebSphere MQ V7.5 to a non-default location and then make it a primary installation, an error message is displayed. The error message shows that linking with libraries, libmqmcs, and libmqmzse has been deprecated, and that you must re-link your applications to avoid using the libmqmcs and libmqmzse libraries. You can set the environment variable `AMQ_NO_MQMCS_MSG` to ensure that IBM WebSphere MQ does not display this error message in the error logs.

The MQCONN or MQCONNX call can succeed only if connecting to a queue manager associated with the same installation owning the library that contains the MQCONN or MQCONNX call.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following (as appropriate):

- On z/OS:
 - Link the application with the correct adapter.
 - Modify the application to use the SRRCMIT and SRRBACK calls in place of the MQCMIT and MQBACK calls. Alternatively, link the application with the RRS batch adapter CSQBRRSI. This adapter supports MQCMIT and MQBACK in addition to SRRCMIT and SRRBACK.
 - For a CICS or IMS application, issue the appropriate CICS or IMS call to commit or backout the unit of work.
 - Start the RRS subsystem on the z/OS system that is running the application.
 - If your application uses Language Environment (LE) ensure that it uses the DLL interface and it runs with POSIX(ON).
 - Ensure that your application is allowed to use Unix System Services (USS).
 - Ensure that your Connection Factory definitions for local z/OS applications and WebSphere Application Server applications use Transport Type with bindings mode connections.
- In the other environments:
 - Link the application with the correct libraries (threaded or nonthreaded).
 - Remove from the application the call or feature that is not supported.
 - Change your application to run setuid, if you want to run fastpath.

2013 (07DD) (RC2013): MQRC_EXPIRY_ERROR:

Explanation

On an MQPUT or MQPUT1 call, the value specified for the *Expiry* field in the message descriptor MQMD is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a value that is greater than zero, or the special value MQEI_UNLIMITED.

2014 (07DE) (RC2014): MQRC_FEEDBACK_ERROR:

Explanation

On an MQPUT or MQPUT1 call, the value specified for the *Feedback* field in the message descriptor MQMD is not valid. The value is not MQFB_NONE, and is outside both the range defined for system feedback codes and the range defined for application feedback codes.

Completion Code

MQCC_FAILED

Programmer response

Specify MQFB_NONE, or a value in the range MQFB_SYSTEM_FIRST through MQFB_SYSTEM_LAST, or MQFB_APPL_FIRST through MQFB_APPL_LAST.

2016 (07E0) (RC2016): MQRC_GET_INHIBITED:

Explanation

MQGET calls are currently inhibited for the queue, or for the queue to which this queue resolves.

Completion Code

MQCC_FAILED

Programmer response

If the system design allows get requests to be inhibited for short periods, retry the operation later.

2017 (07E1) (RC2017): MQRC_HANDLE_NOT_AVAILABLE:

Explanation

An MQOPEN, MQPUT1 or MQSUB call was issued, but the maximum number of open handles allowed for the current task has already been reached. Be aware that when a distribution list is specified on the MQOPEN or MQPUT1 call, each queue in the distribution list uses one handle.

- On z/OS, "task" means a CICS task, a z/OS task, or an IMS-dependent region.

In addition, the MQSUB call allocates two handles when you do not provide an object handle on input.

Completion Code

MQCC_FAILED

Programmer response

Check whether the application is issuing MQOPEN calls without corresponding MQCLOSE calls. If it is, modify the application to issue the MQCLOSE call for each open object as soon as that object is no longer needed.

Also check whether the application is specifying a distribution list containing a large number of queues that are consuming all of the available handles. If it is, increase the maximum number of handles that the task can use, or reduce the size of the distribution list. The maximum number of open handles that a task can use is given by the *MaxHandles* queue manager attribute.

2018 (07E2) (RC2018): MQRC_HCONN_ERROR:

Explanation

The connection handle *Hconn* is not valid, for one of the following reasons:

- The parameter pointer is not valid, or (for the MQCONN or MQCONNX call) points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The value specified was not returned by a preceding MQCONN or MQCONNX call.
- The value specified has been made invalid by a preceding MQDISC call.
- The handle is a shared handle that has been made invalid by another thread issuing the MQDISC call.
- The handle is a shared handle that is being used on the MQBEGIN call (only nonshared handles are valid on MQBEGIN).
- The handle is a nonshared handle that is being used a thread that did not create the handle.
- The call was issued in the MTS environment in a situation where the handle is not valid (for example, passing the handle between processes or packages; note that passing the handle between library packages *is* supported).
- The conversion program is not defined as OPENAPI, when the MQXCNVC call is invoked by running a character conversion exit program with CICS TS 3.2 or higher. When the conversion process runs, the TCB is switched to the Quasi Reentrant (QR) TCB, making the connection incorrect.

Completion Code

MQCC_FAILED

Programmer response

Ensure that a successful MQCONN or MQCONNX call is performed for the queue manager, and that an MQDISC call has not already been performed for it. Ensure that the handle is being used within its valid scope (see the description of MQCONN in  MQCONN – Connect queue manager (*WebSphere MQ V7.1 Reference*) for more information about MQCONN).

- On z/OS, also check that the application has been linked with the correct stub; this is CSQCSTUB for CICS applications, CSQBSTUB for batch applications, and CSQQSTUB for IMS applications. Also, the stub used must not belong to a release of the queue manager that is more recent than the release on which the application will run.

Ensure the character conversion exit program run by your CICS TS 3.2 or higher application, which invokes the MQXCNVC call, is defined as OPENAPI. This definition prevents the 2018 MQRC_HCONN_ERROR error caused by from an incorrect connection, and allows the MQGET to complete.

2019 (07E3) (RC2019): MQRC_HOBJ_ERROR:

Explanation

The object handle *Hobj* is not valid, for one of the following reasons:

- The parameter pointer is not valid, or (for the MQOPEN call) points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The value specified was not returned by a preceding MQOPEN call.
- The value specified has been made invalid by a preceding MQCLOSE call.

- The handle is a shared handle that has been made invalid by another thread issuing the MQCLOSE call.
- The handle is a nonshared handle that is being used by a thread that did not create the handle.
- The call is MQGET or MQPUT, but the object represented by the handle is not a queue.

Completion Code

MQCC_FAILED

Programmer response

Ensure that a successful MQOPEN call is performed for this object, and that an MQCLOSE call has not already been performed for it. Ensure that the handle is being used within its valid scope (see the description of MQOPEN in  MQOPEN – Open object (*WebSphere MQ V7.1 Reference*) for more information).

2020 (07E4) (RC2020): MQRC_INHIBIT_VALUE_ERROR:

Explanation

On an MQSET call, the value specified for either the MQIA_INHIBIT_GET attribute or the MQIA_INHIBIT_PUT attribute is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value for the *InhibitGet* or *InhibitPut* queue attribute.

2021 (07E5) (RC2021): MQRC_INT_ATTR_COUNT_ERROR:

Explanation

On an MQINQ or MQSET call, the *IntAttrCount* parameter is negative (MQINQ or MQSET), or smaller than the number of integer attribute selectors (MQIA_*) specified in the *Selectors* parameter (MQSET only). This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Specify a value large enough for all selected integer attributes.

2022 (07E6) (RC2022): MQRC_INT_ATTR_COUNT_TOO_SMALL:

Explanation

On an MQINQ call, the *IntAttrCount* parameter is smaller than the number of integer attribute selectors (MQIA_*) specified in the *Selectors* parameter.

The call completes with MQCC_WARNING, with the *IntAttrs* array filled in with as many integer attributes as there is room for.

Completion Code

MQCC_WARNING

Programmer response

Specify a large enough value, unless only a subset of the values is needed.

2023 (07E7) (RC2023): MQRC_INT_ATTRS_ARRAY_ERROR:

Explanation

On an MQINQ or MQSET call, the *IntAttrs* parameter is not valid. The parameter pointer is not valid (MQINQ and MQSET), or points to read-only storage or to storage that is not as long as indicated by the *IntAttrCount* parameter (MQINQ only). (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2024 (07E8) (RC2024): MQRC_SYNCPOINT_LIMIT_REACHED:

Explanation

An MQGET, MQPUT, or MQPUT1 call failed because it would have caused the number of uncommitted messages in the current unit of work to exceed the limit defined for the queue manager (see the *MaxUncommittedMsgs* queue-manager attribute). The number of uncommitted messages is the sum of the following since the start of the current unit of work:

- Messages put by the application with the MQPMO_SYNCPOINT option
- Messages retrieved by the application with the MQGMO_SYNCPOINT option
- Trigger messages and COA report messages generated by the queue manager for messages put with the MQPMO_SYNCPOINT option
- COD report messages generated by the queue manager for messages retrieved with the MQGMO_SYNCPOINT option
- On HP Integrity NonStop Server, this reason code occurs when the maximum number of I/O operations in a single TM/MP transaction has been exceeded.

When publishing messages out of syncpoint to topics it is possible to receive this reason code; see

 Publications under syncpoint for more information.

Completion Code

MQCC_FAILED

Programmer response

Check whether the application is looping. If it is not, consider reducing the complexity of the application. Alternatively, increase the queue-manager limit for the maximum number of uncommitted messages within a unit of work.

- On z/OS, the limit for the maximum number of uncommitted messages can be changed by using the ALTER QMGR command.
- On IBM i, the limit for the maximum number of uncommitted messages can be changed by using the CHGMQM command.
- On HP Integrity NonStop Server, the application should cancel the transaction and retry with a smaller number of operations in the unit of work. See the *MQSeries for Tandem NonStop Kernel System Management Guide* for more details.

2025 (07E9) (RC2025): MQRC_MAX_CONNS_LIMIT_REACHED:

Explanation

The MQCONN or MQCONNX call was rejected because the maximum number of concurrent connections has been exceeded.

- On z/OS, the connection limits are 32767 for both TSO and Batch.
- On HP OpenVMS, IBM i, HP Integrity NonStop Server, UNIX systems, and Windows, this reason code can also occur on the MQOPEN call.
- When using Java applications, the connection manager might define a limit to the number of concurrent connections.

Completion Code

MQCC_FAILED

Programmer response

Either increase the size of the appropriate parameter value, or reduce the number of concurrent connections.

2026 (07EA) (RC2026): MQRC_MD_ERROR:

Explanation

The MQMD structure is not valid, for one of the following reasons:

- The *StrucId* field is not MQMD_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQMD structure are set correctly.

2027 (07EB) (RC2027): MQRC_MISSING_REPLY_TO_Q:

Explanation

On an MQPUT or MQPUT1 call, the *ReplyToQ* field in the message descriptor MQMD is blank, but one or both of the following is true:

- A reply was requested (that is, MQMT_REQUEST was specified in the *MsgType* field of the message descriptor).
- A report message was requested in the *Report* field of the message descriptor.

Completion Code

MQCC_FAILED

Programmer response

Specify the name of the queue to which the reply message or report message is to be sent.

2029 (07ED) (RC2029): MQRC_MSG_TYPE_ERROR:

Explanation

Either:

- On an MQPUT or MQPUT1 call, the value specified for the *MsgType* field in the message descriptor (MQMD) is not valid.
- A message processing program received a message that does not have the expected message type. For example, if the WebSphere MQ command server receives a message which is not a request message (MQMT_REQUEST) then it rejects the request with this reason code.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value for the *MsgType* field. In the case where a request is rejected by a message processing program, refer to the documentation for that program for details of the message types that it supports.

2030 (07EE) (RC2030): MQRC_MSG_TOO_BIG_FOR_Q:

Explanation

An MQPUT or MQPUT1 call was issued to put a message on a queue, but the message was too long for the queue and MQMF_SEGMENTATION_ALLOWED was not specified in the *MsgFlags* field in MQMD. If segmentation is not allowed, the length of the message cannot exceed the lesser of the queue *MaxMsgLength* attribute and queue-manager *MaxMsgLength* attribute.

- On z/OS, the queue manager does not support the segmentation of messages; if MQMF_SEGMENTATION_ALLOWED is specified, it is accepted but ignored.

This reason code can also occur when `MQMF_SEGMENTATION_ALLOWED` is specified, but the nature of the data present in the message prevents the queue manager splitting it into segments that are small enough to place on the queue:

- For a user-defined format, the smallest segment that the queue manager can create is 16 bytes.
- For a built-in format, the smallest segment that the queue manager can create depends on the particular format, but is greater than 16 bytes in all cases other than `MQFMT_STRING` (for `MQFMT_STRING` the minimum segment size is 16 bytes).

`MQRC_MSG_TOO_BIG_FOR_Q` can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

Completion Code

`MQCC_FAILED`

Programmer response

Check whether the *BufferLength* parameter is specified correctly; if it is, do one of the following:

- Increase the value of the queue's *MaxMsgLength* attribute; the queue-manager's *MaxMsgLength* attribute may also need increasing.
- Break the message into several smaller messages.
- Specify `MQMF_SEGMENTATION_ALLOWED` in the *MsgFlags* field in `MQMD`; this will allow the queue manager to break the message into segments.

2031 (07EF) (RC2031): MQRC_MSG_TOO_BIG_FOR_Q_MGR:

Explanation

An `MQPUT` or `MQPUT1` call was issued to put a message on a queue, but the message was too long for the queue manager and `MQMF_SEGMENTATION_ALLOWED` was not specified in the *MsgFlags* field in `MQMD`. If segmentation is not allowed, the length of the message cannot exceed the lesser of the queue-manager *MaxMsgLength* attribute and queue *MaxMsgLength* attribute.

This reason code can also occur when `MQMF_SEGMENTATION_ALLOWED` is specified, but the nature of the data present in the message prevents the queue manager splitting it into segments that are small enough for the queue-manager limit:

- For a user-defined format, the smallest segment that the queue manager can create is 16 bytes.
- For a built-in format, the smallest segment that the queue manager can create depends on the particular format, but is greater than 16 bytes in all cases other than `MQFMT_STRING` (for `MQFMT_STRING` the minimum segment size is 16 bytes).

`MQRC_MSG_TOO_BIG_FOR_Q_MGR` can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

This reason also occurs if a channel, through which the message is to pass, has restricted the maximum message length to a value that is actually less than that supported by the queue manager, and the message length is greater than this value.

- On z/OS, this return code is issued only if you are using CICS for distributed queuing. Otherwise, `MQRC_MSG_TOO_BIG_FOR_CHANNEL` is issued.

Completion Code

MQCC_FAILED

Programmer response

Check whether the *BufferLength* parameter is specified correctly; if it is, do one of the following:

- Increase the value of the queue-manager's *MaxMsgLength* attribute; the queue's *MaxMsgLength* attribute may also need increasing.
- Break the message into several smaller messages.
- Specify MQMF_SEGMENTATION_ALLOWED in the *MsgFlags* field in MQMD; this will allow the queue manager to break the message into segments.
- Check the channel definitions.

2033 (07F1) (RC2033): MQRC_NO_MSG_AVAILABLE:

Explanation

An MQGET call was issued, but there is no message on the queue satisfying the selection criteria specified in MQMD (the *MsgId* and *CorrelId* fields), and in MQGMO (the *Options* and *MatchOptions* fields). Either the MQGMO_WAIT option was not specified, or the time interval specified by the *WaitInterval* field in MQGMO has expired. This reason is also returned for an MQGET call for browse, when the end of the queue has been reached.

This reason code can also be returned by the mqGetBag and mqExecute calls. mqGetBag is similar to MQGET. For the mqExecute call, the completion code can be either MQCC_WARNING or MQCC_FAILED:

- If the completion code is MQCC_WARNING, some response messages were received during the specified wait interval, but not all. The response bag contains system-generated nested bags for the messages that were received.
- If the completion code is MQCC_FAILED, no response messages were received during the specified wait interval.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

If this is an expected condition, no corrective action is required.

If this is an unexpected condition, check that:

- The message was put on the queue successfully.
- The unit of work (if any) used for the MQPUT or MQPUT1 call was committed successfully.
- The options controlling the selection criteria are specified correctly. All of the following can affect the eligibility of a message for return on the MQGET call:
 - MQGMO_LOGICAL_ORDER
 - MQGMO_ALL_MSGS_AVAILABLE
 - MQGMO_ALL_SEGMENTS_AVAILABLE
 - MQGMO_COMPLETE_MSG
 - MQMO_MATCH_MSG_ID
 - MQMO_MATCH_CORREL_ID

- MQMO_MATCH_GROUP_ID
- MQMO_MATCH_MSG_SEQ_NUMBER
- MQMO_MATCH_OFFSET
- Value of *MsgId* field in MQMD
- Value of *CorrelId* field in MQMD

Consider waiting longer for the message.

2034 (07F2) (RC2034): MQRC_NO_MSG_UNDER_CURSOR:

Explanation

An MQGET call was issued with either the MQGMO_MSG_UNDER_CURSOR or the MQGMO_BROWSE_MSG_UNDER_CURSOR option. However, the browse cursor is not positioned at a retrievable message. This is caused by one of the following:

- The cursor is positioned logically before the first message (as it is before the first MQGET call with a browse option has been successfully performed).
- The message the browse cursor was positioned on has been locked or removed from the queue (probably by some other application) since the browse operation was performed.
- The message the browse cursor was positioned on has expired.

Completion Code

MQCC_FAILED

Programmer response

Check the application logic. This may be an expected reason if the application design allows multiple servers to compete for messages after browsing. Consider also using the MQGMO_LOCK option with the preceding browse MQGET call.

2035 (07F3) (RC2035): MQRC_NOT_AUTHORIZED:

General explanation

Explanation

The user of the application or channel that produced the error, is not authorized to perform the operation attempted:

- On an MQCONN or MQCONNX call, the user is not authorized to connect to the queue manager.
 - On z/OS, for CICS applications, MQRC_CONNECTION_NOT_AUTHORIZED is issued instead.
- On an MQOPEN or MQPUT1 call, the user is not authorized to open the object for the option(s) specified.
 - On z/OS, if the object being opened is a model queue, this reason also arises if the user is not authorized to create a dynamic queue with the required name.
- On an MQCLOSE call, the user is not authorized to delete the object, which is a permanent dynamic queue, and the *Hobj* parameter specified on the MQCLOSE call is not the handle returned by the MQOPEN call that created the queue.
- On a command, the user is not authorized to issue the command, or to access the object it specifies.

This reason code can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct queue manager or object was specified, and that appropriate authority exists.

Specific problems generating RC2035

JMSWMQ2013 invalid security authentication

See [🔗](#) Invalid security authentication for information your IBM WebSphere MQ JMS application fails with security authentication errors

MQRC_NOT_AUTHORIZED on a queue or channel

See [🔗](#) MQRC_NOT_AUTHORIZED on a queue for information when MQRC 2035 (MQRC_NOT_AUTHORIZED) is returned where a user is not authorized to perform the function. Determine which object the user cannot access and provide the user access to the object.

MQRC_NOT_AUTHORIZED (AMQ4036 on a client) as an administrator

See [🔗](#) MQRC_NOT_AUTHORIZED as an administrator for information when MQRC 2035 (MQRC_NOT_AUTHORIZED) is returned where you try to use a user ID that is a IBM WebSphere MQ Administrator, to remotely access the queue manager through a client connection.

MQS_REPORT_NOAUTH

See [🔗](#) MQS_REPORT_NOAUTH for information on using this environment variable to better diagnose return code 2035 (MQRC_NOT_AUTHORIZED). The use of this environment variable generates errors in the queue manager error log, but does not generate a Failure Data Capture (FDC).

MQSAUTHERRORS

See [🔗](#) MQSAUTHERRORS for information on using this environment variable to generate FDC files related to return code 2035 (MQRC_NOT_AUTHORIZED). The use of this environment variable generates an FDC, but does not generate errors in the queue manager error log.

2036 (07F4) (RC2036): MQRC_NOT_OPEN_FOR_BROWSE:

Explanation

An MQGET call was issued with one of the following options:

- MQGMO_BROWSE_FIRST
- MQGMO_BROWSE_NEXT
- MQGMO_BROWSE_MSG_UNDER_CURSOR
- MQGMO_MSG_UNDER_CURSOR

but either the queue had not been opened for browse, or you are using WebSphere MQ Multicast messaging.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_BROWSE when the queue is opened.

If you are using WebSphere MQ Multicast messaging, you cannot specify browse options with an MQGET call.

2037 (07F5) (RC2037): MQRC_NOT_OPEN_FOR_INPUT:**Explanation**

An MQGET call was issued to retrieve a message from a queue, but the queue had not been opened for input.

Completion Code

MQCC_FAILED

Programmer response

Specify one of the following when the queue is opened:

- MQOO_INPUT_SHARED
- MQOO_INPUT_EXCLUSIVE
- MQOO_INPUT_AS_Q_DEF

2038 (07F6) (RC2038): MQRC_NOT_OPEN_FOR_INQUIRE:**Explanation**

An MQINQ call was issued to inquire object attributes, but the object had not been opened for inquire.

An MQINQ call was issued for a topic handle in WebSphere MQ Multicast.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_INQUIRE when the object is opened.

MQINQ is not supported for topic handles in WebSphere MQ Multicast.

2039 (07F7) (RC2039): MQRC_NOT_OPEN_FOR_OUTPUT:

Explanation

An MQPUT call was issued to put a message on a queue, but the queue had not been opened for output.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_OUTPUT when the queue is opened.

2040 (07F8) (RC2040): MQRC_NOT_OPEN_FOR_SET:

Explanation

An MQSET call was issued to set queue attributes, but the queue had not been opened for set.

An MQSET call was issued for a topic handle in WebSphere MQ Multicast.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_SET when the object is opened.

MQSET is not supported for topic handles in WebSphere MQ Multicast.

2041 (07F9) (RC2041): MQRC_OBJECT_CHANGED:

Explanation

Object definitions that affect this object have been changed since the *Hobj* handle used on this call was returned by the MQOPEN call. For more information about the MQOPEN call, see  MQOPEN – Open object (*WebSphere MQ V7.1 Reference*).

This reason does not occur if the object handle is specified in the *Context* field of the *PutMsgOpts* parameter on the MQPUT or MQPUT1 call.

Completion Code

MQCC_FAILED

Programmer response

Issue an MQCLOSE call to return the handle to the system. It is then usually sufficient to reopen the object and retry the operation. However, if the object definitions are critical to the application logic, an MQINQ call can be used after reopening the object, to obtain the new values of the object attributes.

2042 (07FA) (RC2042): MQRC_OBJECT_IN_USE:

Explanation

An MQOPEN call was issued, but the object in question has already been opened by this or another application with options that conflict with those specified in the *Options* parameter. This arises if the request is for shared input, but the object is already open for exclusive input; it also arises if the request is for exclusive input, but the object is already open for input (of any sort).

MCAs for receiver channels, or the intra-group queuing agent (IGQ agent), may keep the destination queues open even when messages are not being transmitted; this results in the queues appearing to be "in use". Use the MQSC command DISPLAY QSTATUS to find out who is keeping the queue open.

- On z/OS, this reason can also occur for an MQOPEN or MQPUT1 call, if the object to be opened (which can be a queue, or for MQOPEN a namelist or process object) is in the process of being deleted.

Completion Code

MQCC_FAILED

Programmer response

System design should specify whether an application is to wait and retry, or take other action.

2043 (07FB) (RC2043): MQRC_OBJECT_TYPE_ERROR:

Explanation

On the MQOPEN or MQPUT1 call, the *ObjectType* field in the object descriptor MQOD specifies a value that is not valid. For the MQPUT1 call, the object type must be MQOT_Q.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid object type.

2044 (07FC) (RC2044): MQRC_OD_ERROR:

Explanation

On the MQOPEN or MQPUT1 call, the object descriptor MQOD is not valid, for one of the following reasons:

- The *StrucId* field is not MQOD_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQOD structure are set correctly.

2045 (07FD) (RC2045): MQRC_OPTION_NOT_VALID_FOR_TYPE:

Explanation

On an MQOPEN or MQCLOSE call, an option is specified that is not valid for the type of object or queue being opened or closed.

For the MQOPEN call, this includes the following cases:

- An option that is inappropriate for the object type (for example, MQOO_OUTPUT for an MQOT_PROCESS object).
- An option that is unsupported for the queue type (for example, MQOO_INQUIRE for a remote queue that has no local definition).
- One or more of the following options:
 - MQOO_INPUT_AS_Q_DEF
 - MQOO_INPUT_SHARED
 - MQOO_INPUT_EXCLUSIVE
 - MQOO_BROWSE
 - MQOO_INQUIRE
 - MQOO_SET

when either:

- the queue name is resolved through a cell directory, or
- *ObjectQMgrName* in the object descriptor specifies the name of a local definition of a remote queue (to specify a queue-manager alias), and the queue named in the *RemoteQMgrName* attribute of the definition is the name of the local queue manager.

For the MQCLOSE call, this includes the following case:

- The MQCO_DELETE or MQCO_DELETE_PURGE option when the queue is not a dynamic queue.

This reason code can also occur on the MQOPEN call when the object being opened is of type MQOT_NAMELIST, MQOT_PROCESS, or MQOT_Q_MGR, but the *ObjectQMgrName* field in MQOD is neither blank nor the name of the local queue manager.

Completion Code

MQCC_FAILED

Programmer response

Specify the correct option. For the MQOPEN call, ensure that the *ObjectQMgrName* field is set correctly. For the MQCLOSE call, either correct the option or change the definition type of the model queue that is used to create the new queue.

2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR:

Explanation

The *Options* parameter or field contains options that are not valid, or a combination of options that is not valid.

- For the MQOPEN, MQCLOSE, MQXCNVC, mqBagToBuffer, mqBufferToBag, mqCreateBag, and mqExecute calls, *Options* is a separate parameter on the call.

This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

- For the MQBEGIN, MQCONNX, MQGET, MQPUT, and MQPUT1 calls, *Options* is a field in the relevant options structure (MQBO, MQCNO, MQGMO, or MQPMO).
- For more information about option errors for WebSphere MQ Multicast see:  MQI concepts and how they relate to multicast.

Completion Code

MQCC_FAILED

Programmer response

Specify valid options. Check the description of the *Options* parameter or field to determine which options and combinations of options are valid. If multiple options are being set by adding the individual options together, ensure that the same option is not added twice. For more information, see  Rules for validating MQI options (*WebSphere MQ V7.1 Reference*).

2047 (07FF) (RC2047): MQRC_PERSISTENCE_ERROR:

Explanation

On an MQPUT or MQPUT1 call, the value specified for the *Persistence* field in the message descriptor MQMD is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify one of the following values:

- MQPER_PERSISTENT
- MQPER_NOT_PERSISTENT
- MQPER_PERSISTENCE_AS_Q_DEF

2048 (0800) (RC2048): MQRC_PERSISTENT_NOT_ALLOWED:

Explanation

On an MQPUT or MQPUT1 call, the value specified for the *Persistence* field in MQMD (or obtained from the *DefPersistence* queue attribute) specifies MQPER_PERSISTENT, but the queue on which the message is being placed does not support persistent messages. Persistent messages cannot be placed on temporary dynamic queues.

This reason code can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

Completion Code

MQCC_FAILED

Programmer response

Specify MQPER_NOT_PERSISTENT if the message is to be placed on a temporary dynamic queue. If persistence is required, use a permanent dynamic queue or predefined queue in place of a temporary dynamic queue.

Be aware that server applications are recommended to send reply messages (message type MQMT_REPLY) with the same persistence as the original request message (message type MQMT_REQUEST). If the request message is persistent, the reply queue specified in the *ReplyToQ* field in the message descriptor MQMD cannot be a temporary dynamic queue. Use a permanent dynamic queue or predefined queue as the reply queue in this situation.

On z/OS, you cannot put persistent messages to a shared queue if the CFSTRUCT that the queue uses is defined with RECOVER(NO). Either put only non-persistent messages to this queue or change the queue definition to RECOVER(YES). If you put a persistent message to a queue that uses a CFSTRUCT with RECOVER(NO) the put will fail with MQRC_PERSISTENT_NOT_ALLOWED.

2049 (0801) (RC2049): MQRC_PRIORITY_EXCEEDS_MAXIMUM:

Explanation

An MQPUT or MQPUT1 call was issued, but the value of the *Priority* field in the message descriptor MQMD exceeds the maximum priority supported by the local queue manager, as shown by the *MaxPriority* queue-manager attribute. The message is accepted by the queue manager, but is placed on the queue at the queue manager's maximum priority. The *Priority* field in the message descriptor retains the value specified by the application that put the message.

Completion Code

MQCC_WARNING

Programmer response

None required, unless this reason code was not expected by the application that put the message.

2050 (0802) (RC2050): MQRC_PRIORITY_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the value of the *Priority* field in the message descriptor MQMD is not valid. The maximum priority supported by the queue manager is given by the *MaxPriority* queue-manager attribute.

Completion Code

MQCC_FAILED

Programmer response

Specify a value in the range zero through *MaxPriority*, or the special value MQPRI_PRIORITY_AS_Q_DEF.

2051 (0803) (RC2051): MQRC_PUT_INHIBITED:

Explanation

MQPUT and MQPUT1 calls are currently inhibited for the queue, or for the queue to which this queue resolves.

This reason code can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

Completion Code

MQCC_FAILED

Programmer response

If the system design allows put requests to be inhibited for short periods, retry the operation later.

2052 (0804) (RC2052): MQRC_Q_DELETED:

Explanation

An *Hobj* queue handle specified on a call refers to a dynamic queue that has been deleted since the queue was opened. For more information about the deletion of dynamic queues, see the description of

MQCLOSE in  MQCLOSE – Close object (*WebSphere MQ V7.1 Reference*).

- On z/OS, this can also occur with the MQOPEN and MQPUT1 calls if a dynamic queue is being opened, but the queue is in a logically-deleted state. See MQCLOSE for more information about this.

Completion Code

MQCC_FAILED

Programmer response

Issue an MQCLOSE call to return the handle and associated resources to the system (the MQCLOSE call will succeed in this case). Check the design of the application that caused the error.

2053 (0805) (RC2053): MQRC_Q_FULL:

Explanation

An MQPUT or MQPUT1 call, or a command, failed because the queue is full, that is, it already contains the maximum number of messages possible, as specified by the *MaxQDepth* queue attribute.

This reason code can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

Completion Code

MQCC_FAILED

Programmer response

Retry the operation later. Consider increasing the maximum depth for this queue, or arranging for more instances of the application to service the queue.

2055 (0807) (RC2055): MQRC_Q_NOT_EMPTY:

Explanation

An MQCLOSE call was issued for a permanent dynamic queue, but the call failed because the queue is not empty or still in use. One of the following applies:

- The MQCO_DELETE option was specified, but there are messages on the queue.
- The MQCO_DELETE or MQCO_DELETE_PURGE option was specified, but there are uncommitted get or put calls outstanding against the queue.

See the usage notes pertaining to dynamic queues for the MQCLOSE call for more information.

This reason code is also returned from a command to clear or delete or move a queue, if the queue contains uncommitted messages (or committed messages in the case of delete queue without the purge option).

Completion Code

MQCC_FAILED

Programmer response

Check why there might be messages on the queue. Be aware that the *CurrentQDepth* queue attribute might be zero even though there are one or more messages on the queue; this can happen if the messages have been retrieved as part of a unit of work that has not yet been committed. If the messages can be discarded, try using the MQCLOSE call with the MQCO_DELETE_PURGE option. Consider retrying the call later.

2056 (0808) (RC2056): MQRC_Q_SPACE_NOT_AVAILABLE:

Explanation

An MQPUT or MQPUT1 call was issued, but there is no space available for the queue on disk or other storage device.

This reason code can also occur in the *Feedback* field in the message descriptor of a report message; in this case it indicates that the error was encountered by a message channel agent when it attempted to put the message on a remote queue.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Check whether an application is putting messages in an infinite loop. If not, make more disk space available for the queue.

2057 (0809) (RC2057): MQRC_Q_TYPE_ERROR:

Explanation

One of the following occurred:

- On an MQOPEN call, the *ObjectQMgrName* field in the object descriptor MQOD or object record MQOR specifies the name of a local definition of a remote queue (to specify a queue-manager alias), and in that local definition the *RemoteQMgrName* attribute is the name of the local queue manager. However, the *ObjectName* field in MQOD or MQOR specifies the name of a model queue on the local queue manager; this is not allowed. For more information, see  MQOPEN – Open object (*WebSphere MQ V7.1 Reference*).
- On an MQPUT1 call, the object descriptor MQOD or object record MQOR specifies the name of a model queue.
- On a previous MQPUT or MQPUT1 call, the *ReplyToQ* field in the message descriptor specified the name of a model queue, but a model queue cannot be specified as the destination for reply or report messages. Only the name of a predefined queue, or the name of the *dynamic* queue created from the model queue, can be specified as the destination. In this situation the reason code MQRC_Q_TYPE_ERROR is returned in the *Reason* field of the MQDLH structure when the reply message or report message is placed on the dead-letter queue.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid queue.

2058 (080A) (RC2058): MQRC_Q_MGR_NAME_ERROR:

Explanation

On an MQCONN or MQCONNX call, the value specified for the *QMgrName* parameter is not valid or not known. This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

- On z/OS for CICS applications, this reason can occur on *any* call if the original connect specified an incorrect or unrecognized name.

This reason code can also occur if an MQ MQI client application attempts to connect to a queue manager within an MQ-client queue-manager group (see the *QMgrName* parameter of MQCONN), and either:

- Queue-manager groups are not supported.
- There is no queue-manager group with the specified name.

Completion Code

MQCC_FAILED

Programmer response

Use an all-blank name if possible, or verify that the name used is valid.

If you are using CICS Resyncmember(Groupresync), use the queue-sharing group (QSG) name in the MQNAME rather than the queue manager name.

2059 (080B) (RC2059): MQRC_Q_MGR_NOT_AVAILABLE:

Explanation

This error occurs:

1. On an MQCONN or MQCONNX call, the queue manager identified by the *QMgrName* parameter is not available for connection.
 - On z/OS:
 - For batch applications, this reason can be returned to applications running in LPARs that do not have a queue manager installed.
 - For CICS applications, this reason can occur on any call if the original connect specified a queue manager with a name that was recognized, but which is not available.
 - On IBM i, this reason can also be returned by the MQOPEN and MQPUT1 calls, when MQHC_DEF_HCONN is specified for the *Hconn* parameter by an application running in compatibility mode.
2. On an MQCONN or MQCONNX call from a IBM WebSphere MQ MQI client application:
 - Attempting to connect to a queue manager within an MQ-client queue-manager group when none of the queue managers in the group is available for connection (see the *QMgrName* parameter of the MQCONN call).
 - If the client channel fails to connect, perhaps because of an error with the client-connection or the corresponding server-connection channel definitions.
 - The z/OS Client Attachment feature has not been installed.
3. If a command uses the *CommandScope* parameter specifying a queue manager that is not active in the queue-sharing group.

4. In a multiple installation environment, where an application attempts to connect to a queue manager associated with an installation of IBM WebSphere MQ Version 7.1, or later, but has loaded libraries from IBM WebSphere MQ Version 7.0.1. IBM WebSphere MQ Version 7.0.1 cannot load libraries from other versions of IBM WebSphere MQ.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the queue manager has been started. If the connection is from a client application, check the channel definitions, channel status, and error logs.

In a multiple installation environment, ensure that IBM WebSphere MQ Version 7.1, or later, libraries are loaded by the operating system. For more information, see  Connecting applications in a multiple installation environment (*WebSphere MQ V7.1 Installing Guide*).

2061 (080D) (RC2061): MQRC_REPORT_OPTIONS_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the *Report* field in the message descriptor MQMD contains one or more options that are not recognized by the local queue manager. The options that cause this reason code to be returned depend on the destination of the message; see the description of REPORT in

 Report options and message flags (*WebSphere MQ V7.1 Reference*) for more details.

This reason code can also occur in the *Feedback* field in the MQMD of a report message, or in the *Reason* field in the MQDLH structure of a message on the dead-letter queue; in both cases it indicates that the destination queue manager does not support one or more of the report options specified by the sender of the message.

Completion Code

MQCC_FAILED

Programmer response

Do the following:

- Ensure that the *Report* field in the message descriptor is initialized with a value when the message descriptor is declared, or is assigned a value prior to the MQPUT or MQPUT1 call. Specify MQRO_NONE if no report options are required.
- Ensure that the report options specified are valid; see the *Report* field described in the description of MQMD in  Report options and message flags (*WebSphere MQ V7.1 Reference*) for valid report options.
- If multiple report options are being set by adding the individual report options together, ensure that the same report option is not added twice.
- Check that conflicting report options are not specified. For example, do not add both MQRO_EXCEPTION and MQRO_EXCEPTION_WITH_DATA to the *Report* field; only one of these can be specified.

2062 (080E) (RC2062): MQRC_SECOND_MARK_NOT_ALLOWED:

Explanation

An MQGET call was issued specifying the MQGMO_MARK_SKIP_BACKOUT option in the *Options* field of MQGMO, but a message has already been marked within the current unit of work. Only one marked message is allowed within each unit of work.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Modify the application so that no more than one message is marked within each unit of work.

2063 (080F) (RC2063): MQRC_SECURITY_ERROR:

Explanation

An MQCONN, MQCONNX, MQOPEN, MQPUT1, or MQCLOSE call was issued, but it failed because a security error occurred.

- On z/OS, the security error was returned by the External Security Manager.

Completion Code

MQCC_FAILED

Programmer response

Note the error from the security manager, and contact your system programmer or security administrator.

- On IBM i, the FFST log will contain the error information.

2065 (0811) (RC2065): MQRC_SELECTOR_COUNT_ERROR:

Explanation

On an MQINQ or MQSET call, the *SelectorCount* parameter specifies a value that is not valid. This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Specify a value in the range 0 through 256.

2066 (0812) (RC2066): MQRC_SELECTOR_LIMIT_EXCEEDED:

Explanation

On an MQINQ or MQSET call, the *SelectorCount* parameter specifies a value that is larger than the maximum supported (256).

Completion Code

MQCC_FAILED

Programmer response

Reduce the number of selectors specified on the call; the valid range is 0 through 256.

2067 (0813) (RC2067): MQRC_SELECTOR_ERROR:

Explanation

An MQINQ or MQSET call was issued, but the *Selectors* array contains a selector that is not valid for one of the following reasons:

- The selector is not supported or out of range.
- The selector is not applicable to the type of object with attributes that are being inquired upon or set.
- The selector is for an attribute that cannot be set.

This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

An MQINQ call was issued for a managed handle in WebSphere MQ Multicast, inquiring a value other than *Current Depth*.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the value specified for the selector is valid for the object type represented by *Hobj*. For the MQSET call, also ensure that the selector represents an integer attribute that can be set.

MQINQ for managed handles in WebSphere MQ Multicast can only inquire on *Current Depth*.

2068 (0814) (RC2068): MQRC_SELECTOR_NOT_FOR_TYPE:

Explanation

On the MQINQ call, one or more selectors in the *Selectors* array is not applicable to the type of the queue with attributes that are being inquired upon.

This reason also occurs when the queue is a cluster queue that resolved to a remote instance of the queue. In this case only a subset of the attributes that are valid for local queues can be inquired. See the

usage notes in the description of MQINQ in  MQINQ – Inquire object attributes (*WebSphere MQ V7.1 Reference*) for more information about MQINQ.

The call completes with MQCC_WARNING, with the attribute values for the inapplicable selectors set as follows:

- For integer attributes, the corresponding elements of *IntAttrs* are set to MQIAV_NOT_APPLICABLE.
- For character attributes, the appropriate parts of the *CharAttrs* string are set to a character string consisting entirely of asterisks (*).

Completion Code

MQCC_WARNING

Programmer response

Verify that the selector specified is the one that was intended.

If the queue is a cluster queue, specifying one of the MQOO_BROWSE, MQOO_INPUT_*, or MQOO_SET options in addition to MQOO_INQUIRE forces the queue to resolve to the local instance of the queue. However, if there is no local instance of the queue the MQOPEN call fails.

2069 (0815) (RC2069): MQRC_SIGNAL_OUTSTANDING:

Explanation

An MQGET call was issued with either the MQGMO_SET_SIGNAL or MQGMO_WAIT option, but there is already a signal outstanding for the queue handle *Hobj*.

This reason code occurs only in the following environments: z/OS, Windows 95, Windows 98.

Completion Code

MQCC_FAILED

Programmer response

Check the application logic. If it is necessary to set a signal or wait when there is a signal outstanding for the same queue, a different object handle must be used.

2070 (0816) (RC2070): MQRC_SIGNAL_REQUEST_ACCEPTED:

Explanation

An MQGET call was issued specifying MQGMO_SET_SIGNAL in the *GetMsgOpts* parameter, but no suitable message was available; the call returns immediately. The application can now wait for the signal to be delivered.

- On z/OS, the application should wait on the Event Control Block pointed to by the *Signal1* field.
- On Windows 95, Windows 98, the application should wait for the signal Windows message to be delivered.

This reason code occurs only in the following environments: z/OS, Windows 95, Windows 98.

Completion Code

MQCC_WARNING

Programmer response

Wait for the signal; when it is delivered, check the signal to ensure that a message is now available. If it is, reissue the MQGET call.

- On z/OS, wait on the ECB pointed to by the *Signal1* field and, when it is posted, check it to ensure that a message is now available.
- On Windows 95, Windows 98, the application (thread) should continue executing its message loop.

2071 (0817) (RC2071): MQRC_STORAGE_NOT_AVAILABLE:

Explanation

The call failed because there is insufficient main storage available.

Completion Code

MQCC_FAILED

Programmer response

Ensure that active applications are behaving correctly, for example, that they are not looping unexpectedly. If no problems are found, make more main storage available.

- On z/OS, if no application problems are found, ask your system programmer to increase the size of the region in which the queue manager runs.

2072 (0818) (RC2072): MQRC_SYNCPOINT_NOT_AVAILABLE:

Explanation

Either the MQGMO_SYNCPOINT option was used with an MQGET call, or the MQPMO_SYNCPOINT option was used with an MQPUT or MQPUT1 call, but the local queue manager was unable to honor the request. If the queue manager does not support units of work, the *SyncPoint* queue-manager attribute has the value MQSP_NOT_AVAILABLE.

This reason code can also occur on the MQGET, MQPUT, and MQPUT1 calls when an external unit-of-work coordinator is used. If that coordinator requires an explicit call to start the unit of work, but the application has not issued that call before the MQGET, MQPUT, or MQPUT1 call, reason code MQRC_SYNCPOINT_NOT_AVAILABLE is returned.

- On IBM i, this reason code means that IBM i Commitment Control is not started, or is unavailable for use by the queue manager.
- On HP Integrity NonStop Server, this reason code means that the client has detected that the application has an active transaction that is being coordinated by the Transaction Management Facility (TMF), but that a z/OS queue manager is unable to be coordinated by TMF.

This reason code can also be returned if the MQGMO_SYNCPOINT or the MQPMO_SYNCPOINT option was used for IBM WebSphere MQ Multicast messaging. Transactions are not supported for multicast.

Completion Code

MQCC_FAILED

Programmer response

Remove the specification of MQGMO_SYNCPOINT or MQPMO_SYNCPOINT, as appropriate.

- On IBM i, ensure that Commitment Control is started. If this reason code occurs after Commitment Control is started, contact your system programmer.
- On HP Integrity NonStop Server, ensure that your z/OS queue manager has the relevant APAR applied. Check with the IBM support center for APAR details.

2075 (081B) (RC2075): MQRC_TRIGGER_CONTROL_ERROR:

Explanation

On an MQSET call, the value specified for the MQIA_TRIGGER_CONTROL attribute selector is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value.

2076 (081C) (RC2076): MQRC_TRIGGER_DEPTH_ERROR:

Explanation

On an MQSET call, the value specified for the MQIA_TRIGGER_DEPTH attribute selector is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a value that is greater than zero.

2077 (081D) (RC2077): MQRC_TRIGGER_MSG_PRIORITY_ERR:

Explanation

On an MQSET call, the value specified for the MQIA_TRIGGER_MSG_PRIORITY attribute selector is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a value in the range zero through the value of *MaxPriority* queue-manager attribute.

2078 (081E) (RC2078): MQRC_TRIGGER_TYPE_ERROR:**Explanation**

On an MQSET call, the value specified for the MQIA_TRIGGER_TYPE attribute selector is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value.

2079 (081F) (RC2079): MQRC_TRUNCATED_MSG_ACCEPTED:**Explanation**

On an MQGET call, the message length was too large to fit into the supplied buffer. The MQGMO_ACCEPT_TRUNCATED_MSG option was specified, so the call completes. The message is removed from the queue (subject to unit-of-work considerations), or, if this was a browse operation, the browse cursor is advanced to this message.

The *DataLength* parameter is set to the length of the message before truncation, the *Buffer* parameter contains as much of the message as fits, and the MQMD structure is filled in.

Completion Code

MQCC_WARNING

Programmer response

None, because the application expected this situation.

2080 (0820) (RC2080): MQRC_TRUNCATED_MSG_FAILED:**Explanation**

On an MQGET call, the message length was too large to fit into the supplied buffer. The MQGMO_ACCEPT_TRUNCATED_MSG option was *not* specified, so the message has not been removed from the queue. If this was a browse operation, the browse cursor remains where it was before this call, but if MQGMO_BROWSE_FIRST was specified, the browse cursor is positioned logically before the highest-priority message on the queue.

The *DataLength* field is set to the length of the message before truncation, the *Buffer* parameter contains as much of the message as fits, and the MQMD structure is filled in.

Completion Code

MQCC_WARNING

Programmer response

Supply a buffer that is at least as large as *DataLength*, or specify MQGMO_ACCEPT_TRUNCATED_MSG if not all of the message data is required.

2082 (0822) (RC2082): MQRC_UNKNOWN_ALIAS_BASE_Q:

Explanation

An MQOPEN or MQPUT1 call was issued specifying an alias queue as the target, but the *BaseQName* in the alias queue attributes is not recognized as a queue name.

This reason code can also occur when *BaseQName* is the name of a cluster queue that cannot be resolved successfully.

MQRC_UNKNOWN_ALIAS_BASE_Q might indicate that the application is specifying the **ObjectQmgrName** of the queue manager that it is connecting to, and the queue manager that is hosting the alias queue. This means that the queue manager looks for the alias target queue on the specified queue manager and fails because the alias target queue is not on the local queue manager. The **ObjectQmgrName** parameter is typically left blank so that the clustering decides which queue manager to route to.

Completion Code

MQCC_FAILED

Programmer response

Correct the queue definitions.

2085 (0825) (RC2085): MQRC_UNKNOWN_OBJECT_NAME:

Explanation

An MQOPEN or MQPUT1 call was issued, but the object identified by the *ObjectName* and *ObjectQMgrName* fields in the object descriptor MQOD cannot be found. One of the following applies:

- The *ObjectQMgrName* field is one of the following:
 - Blank
 - The name of the local queue manager
 - The name of a local definition of a remote queue (a queue-manager alias) in which the *RemoteQMgrName* attribute is the name of the local queue managerbut no object with the specified *ObjectName* and *ObjectType* exists on the local queue manager.
- The object being opened is a cluster queue that is hosted on a remote queue manager, but the local queue manager does not have a defined route to the remote queue manager.
- The object being opened is a queue definition that has QSGDISP(GROUP). Such definitions cannot be used with the MQOPEN and MQPUT1 calls.
- The MQOD in the failing application specifies the name of the local queue manager in *ObjectQMgrName*. The local queue manager does not host the particular cluster queue specified in *ObjectName*.
The solution in this environment is to leave *ObjectQMgrName* of the MQOD blank.

This can also occur in response to a command that specifies the name of an object or other item that does not exist.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid object name. Ensure that the name is padded with blanks at the end, if necessary. If this is correct, check the object definitions.

2086 (0826) (RC2086): MQRC_UNKNOWN_OBJECT_Q_MGR:

Explanation

On an MQOPEN or MQPUT1 call, the *ObjectQMgrName* field in the object descriptor MQOD does not satisfy the naming rules for objects. For more information, see  *ObjectQMgrName (MQCHAR48)* (*WebSphere MQ V7.1 Reference*).

This reason also occurs if the *ObjectType* field in the object descriptor has the value MQOT_Q_MGR, and the *ObjectQMgrName* field is not blank, but the name specified is not the name of the local queue manager.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid queue manager name. To refer to the local queue manager, a name consisting entirely of blanks or beginning with a null character can be used. Ensure that the name is padded with blanks at the end, or terminated with a null character if necessary.

2087 (0827) (RC2087): MQRC_UNKNOWN_REMOTE_Q_MGR:

Explanation

On an MQOPEN or MQPUT1 call, an error occurred with the queue-name resolution, for one of the following reasons:

- *ObjectQMgrName* is blank or the name of the local queue manager, *ObjectName* is the name of a local definition of a remote queue (or an alias to one), and one of the following is true:
 - *RemoteQMgrName* is blank or the name of the local queue manager. Note that this error occurs even if *XmitQName* is not blank.
 - *XmitQName* is blank, but there is no transmission queue defined with the name of *RemoteQMgrName*, and the *DefXmitQName* queue-manager attribute is blank.
 - *RemoteQMgrName* and *RemoteQName* specify a cluster queue that cannot be resolved successfully, and the *DefXmitQName* queue-manager attribute is blank.
 - On z/OS only, the *RemoteQMgrName* is the name of a queue manager in the Queue Sharing group but intra-group queuing is disabled.
- *ObjectQMgrName* is the name of a local definition of a remote queue (containing a queue-manager alias definition), and one of the following is true:
 - *RemoteQName* is not blank.
 - *XmitQName* is blank, but there is no transmission queue defined with the name of *RemoteQMgrName*, and the *DefXmitQName* queue-manager attribute is blank.
- *ObjectQMgrName* is not:
 - Blank
 - The name of the local queue manager
 - The name of a transmission queue

- The name of a queue-manager alias definition (that is, a local definition of a remote queue with a blank *RemoteQName*)

but the *DefXmitQName* queue-manager attribute is blank and the queue manager is not part of a queue-sharing group with intra-group queuing enabled.

- *ObjectQMgrName* is the name of a model queue.
- The queue name is resolved through a cell directory. However, there is no queue defined with the same name as the remote queue manager name obtained from the cell directory, and the *DefXmitQName* queue-manager attribute is blank.

Completion Code

MQCC_FAILED

Programmer response

Check the values specified for *ObjectQMgrName* and *ObjectName*. If these are correct, check the queue definitions.

2090 (082A) (RC2090): MQRC_WAIT_INTERVAL_ERROR:

Explanation

On the MQGET call, the value specified for the *WaitInterval* field in the *GetMsgOpts* parameter is not valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a value greater than or equal to zero, or the special value MQWI_UNLIMITED if an indefinite wait is required.

2091 (082B) (RC2091): MQRC_XMIT_Q_TYPE_ERROR:

Explanation

On an MQOPEN or MQPUT1 call, a message is to be sent to a remote queue manager. The *ObjectName* or *ObjectQMgrName* field in the object descriptor specifies the name of a local definition of a remote queue but one of the following applies to the *XmitQName* attribute of the definition:

- *XmitQName* is not blank, but specifies a queue that is not a local queue
- *XmitQName* is blank, but *RemoteQMgrName* specifies a queue that is not a local queue

This reason also occurs if the queue name is resolved through a cell directory, and the remote queue manager name obtained from the cell directory is the name of a queue, but this is not a local queue.

Completion Code

MQCC_FAILED

Programmer response

Check the values specified for *ObjectName* and *ObjectQMgrName*. If these are correct, check the queue definitions.

2092 (082C) (RC2092): MQRC_XMIT_Q_USAGE_ERROR:

Explanation

On an MQOPEN or MQPUT1 call, a message is to be sent to a remote queue manager, but one of the following occurred:

- *ObjectQMgrName* specifies the name of a local queue, but it does not have a *Usage* attribute of MQUS_TRANSMISSION.
- The *ObjectName* or *ObjectQMgrName* field in the object descriptor specifies the name of a local definition of a remote queue but one of the following applies to the *XmitQName* attribute of the definition:
 - *XmitQName* is not blank, but specifies a queue that does not have a *Usage* attribute of MQUS_TRANSMISSION
 - *XmitQName* is blank, but *RemoteQMgrName* specifies a queue that does not have a *Usage* attribute of MQUS_TRANSMISSION
 - *XmitQName* specifies the queue SYSTEM.QSG.TRANSMIT.QUEUE the IGQ queue manager attribute indicates that IGQ is DISABLED.
- The queue name is resolved through a cell directory, and the remote queue manager name obtained from the cell directory is the name of a local queue, but it does not have a *Usage* attribute of MQUS_TRANSMISSION.

Completion Code

MQCC_FAILED

Programmer response

Check the values specified for *ObjectName* and *ObjectQMgrName*. If these are correct, check the queue definitions.

2093 (082D) (RC2093): MQRC_NOT_OPEN_FOR_PASS_ALL:

Explanation

An MQPUT call was issued with the MQPMO_PASS_ALL_CONTEXT option specified in the *PutMsgOpts* parameter, but the queue had not been opened with the MQOO_PASS_ALL_CONTEXT option.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_PASS_ALL_CONTEXT (or another option that implies it) when the queue is opened.

2094 (082E) (RC2094): MQRC_NOT_OPEN_FOR_PASS_IDENT:

Explanation

An MQPUT call was issued with the MQPMO_PASS_IDENTITY_CONTEXT option specified in the *PutMsgOpts* parameter, but the queue had not been opened with the MQOO_PASS_IDENTITY_CONTEXT option.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_PASS_IDENTITY_CONTEXT (or another option that implies it) when the queue is opened.

2095 (082F) (RC2095): MQRC_NOT_OPEN_FOR_SET_ALL:

Explanation

An MQPUT call was issued with the MQPMO_SET_ALL_CONTEXT option specified in the *PutMsgOpts* parameter, but the queue had not been opened with the MQOO_SET_ALL_CONTEXT option.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_SET_ALL_CONTEXT when the queue is opened.

2096 (0830) (RC2096): MQRC_NOT_OPEN_FOR_SET_IDENT:

Explanation

An MQPUT call was issued with the MQPMO_SET_IDENTITY_CONTEXT option specified in the *PutMsgOpts* parameter, but the queue had not been opened with the MQOO_SET_IDENTITY_CONTEXT option.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_SET_IDENTITY_CONTEXT (or another option that implies it) when the queue is opened.

2097 (0831) (RC2097): MQRC_CONTEXT_HANDLE_ERROR:

Explanation

On an MQPUT or MQPUT1 call, MQPMO_PASS_IDENTITY_CONTEXT or MQPMO_PASS_ALL_CONTEXT was specified, but the handle specified in the *Context* field of the *PutMsgOpts* parameter is either not a valid queue handle, or it is a valid queue handle but the queue was not opened with MQOO_SAVE_ALL_CONTEXT.

Completion Code

MQCC_FAILED

Programmer response

Specify MQOO_SAVE_ALL_CONTEXT when the queue referred to is opened.

2098 (0832) (RC2098): MQRC_CONTEXT_NOT_AVAILABLE:

Explanation

On an MQPUT or MQPUT1 call, MQPMO_PASS_IDENTITY_CONTEXT or MQPMO_PASS_ALL_CONTEXT was specified, but the queue handle specified in the *Context* field of the *PutMsgOpts* parameter has no context associated with it. This arises if no message has yet been successfully retrieved with the queue handle referred to, or if the last successful MQGET call was a browse.

This condition does not arise if the message that was last retrieved had no context associated with it.

- On z/OS, if a message is received by a message channel agent that is putting messages with the authority of the user identifier in the message, this code is returned in the *Feedback* field of an exception report if the message has no context associated with it.

Completion Code

MQCC_FAILED

Programmer response

Ensure that a successful nonbrowse get call has been issued with the queue handle referred to.

2099 (0833) (RC2099): MQRC_SIGNAL1_ERROR:

Explanation

An MQGET call was issued, specifying MQGMO_SET_SIGNAL in the *GetMsgOpts* parameter, but the *Signal1* field is not valid.

- On z/OS, the address contained in the *Signal1* field is not valid, or points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- On Windows 95, Windows 98, the window handle in the *Signal1* field is not valid.

This reason code occurs only in the following environments: z/OS, Windows 95, Windows 98.

Completion Code

MQCC_FAILED

Programmer response

Correct the setting of the *Signal1* field.

2100 (0834) (RC2100): MQRC_OBJECT_ALREADY_EXISTS:

Explanation

An MQOPEN call was issued to create a dynamic queue, but a queue with the same name as the dynamic queue already exists.

- On z/OS, a rare “race condition” can also give rise to this reason code; see the description of reason code MQRC_NAME_IN_USE for more details.

Completion Code

MQCC_FAILED

Programmer response

If supplying a dynamic queue name in full, ensure that it obeys the naming conventions for dynamic queues; if it does, either supply a different name, or delete the existing queue if it is no longer required. Alternatively, allow the queue manager to generate the name.

If the queue manager is generating the name (either in part or in full), reissue the MQOPEN call.

2101 (0835) (RC2101): MQRC_OBJECT_DAMAGED:

Explanation

The object accessed by the call is damaged and cannot be used. For example, this might be because the definition of the object in main storage is not consistent, or because it differs from the definition of the object on disk, or because the definition on disk cannot be read. The object can be deleted, although it might not be possible to delete the associated user space.

- On z/OS, this reason occurs when the Db2 list header or structure number associated with a shared queue is zero. This situation arises as a result of using the MQSC command DELETE CFSTRUCT to delete the Db2 structure definition. The command resets the list header and structure number to zero for each of the shared queues that references the deleted CF structure.

Completion Code

MQCC_FAILED

Programmer response

It might be necessary to stop and restart the queue manager, or to restore the queue-manager data from backup storage.

- On HP OpenVMS, IBM i, HP Integrity NonStop Server, and UNIX systems, consult the FFST™ record to obtain more detail about the problem.
- On z/OS, delete the shared queue and redefine it using the MQSC command DEFINE QLOCAL. This automatically defines a CF structure and allocates list headers for it.

2102 (0836) (RC2102): MQRC_RESOURCE_PROBLEM:

Explanation

There are insufficient system resources to complete the call successfully. On z/OS this can indicate that Db2 errors occurred when using shared queues, or that the maximum number of shared queues that can be defined in a single coupling facility list structure has been reached.

Completion Code

MQCC_FAILED

Programmer response

Run the application when the machine is less heavily loaded.

- On z/OS, check the operator console for messages that might provide additional information.

- On HP OpenVMS, IBM i, HP Integrity NonStop Server, and UNIX systems, consult the FFST record to obtain more detail about the problem.

2103 (0837) (RC2103): MQRC_ANOTHER_Q_MGR_CONNECTED:

Explanation

An MQCONN or MQCONNX call was issued, but the thread or process is already connected to a different queue manager. The thread or process can connect to only one queue manager at a time.

- On z/OS, this reason code does not occur.
- On Windows, MTS objects do not receive this reason code, as connections to other queue managers are allowed.

Completion Code

MQCC_FAILED

Programmer response

Use the MQDISC call to disconnect from the queue manager that is already connected, and then issue the MQCONN or MQCONNX call to connect to the new queue manager.

Disconnecting from the existing queue manager closes any queues that are currently open; it is suggested that any uncommitted units of work are committed or backed out before the MQDISC call is issued.

2104 (0838) (RC2104): MQRC_UNKNOWN_REPORT_OPTION:

Explanation

An MQPUT or MQPUT1 call was issued, but the *Report* field in the message descriptor MQMD contains one or more options that are not recognized by the local queue manager. The options are accepted.

The options that cause this reason code to be returned depend on the destination of the message; see the description of REPORT in  Report options and message flags (*WebSphere MQ V7.1 Reference*) for more information.

Completion Code

MQCC_WARNING

Programmer response

If this reason code is expected, no corrective action is required. If this reason code is not expected, do the following:

- Ensure that the *Report* field in the message descriptor is initialized with a value when the message descriptor is declared, or is assigned a value prior to the MQPUT or MQPUT1 call.
- Ensure that the report options specified are valid; see the *Report* field described in the description of MQMD in  MQMD – Message descriptor (*WebSphere MQ V7.1 Reference*) for valid report options.
- If multiple report options are being set by adding the individual report options together, ensure that the same report option is not added twice.
- Check that conflicting report options are not specified. For example, do not add both MQRO_EXCEPTION and MQRO_EXCEPTION_WITH_DATA to the *Report* field; only one of these can be specified.

2105 (0839) (RC2105): MQRC_STORAGE_CLASS_ERROR:

Explanation

The MQPUT or MQPUT1 call was issued, but the storage-class object defined for the queue does not exist.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Create the storage-class object required by the queue, or modify the queue definition to use an existing storage class. The name of the storage-class object used by the queue is given by the *StorageClass* queue attribute.

2106 (083A) (RC2106): MQRC_COD_NOT_VALID_FOR_XCF_Q:

Explanation

An MQPUT or MQPUT1 call was issued, but the *Report* field in the message descriptor MQMD specifies one of the MQRO_COD_* options and the target queue is an XCF queue. MQRO_COD_* options cannot be specified for XCF queues.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Remove the relevant MQRO_COD_* option.

2107 (083B) (RC2107): MQRC_XWAIT_CANCELED:

Explanation

An MQXWAIT call was issued, but the call has been canceled because a STOP CHINIT command has been issued (or the queue manager has been stopped, which causes the same effect). See  MQXWAIT – Wait in exit (*WebSphere MQ V7.1 Reference*) for more information about the MQXWAIT call.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Tidy up and terminate.

2108 (083C) (RC2108): MQRC_XWAIT_ERROR:

Explanation

An MQXWAIT call was issued, but the invocation was not valid for one of the following reasons:

- The wait descriptor MQXWD contains data that is not valid.
- The linkage stack level is not valid.
- The addressing mode is not valid.
- There are too many wait events outstanding.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Obey the rules for using the MQXWAIT call. For more information about MQWAIT, see  MQXWAIT – Wait in exit (*WebSphere MQ V7.1 Reference*).

2109 (083D) (RC2109): MQRC_SUPPRESSED_BY_EXIT:

Explanation

On any call other than MQCONN or MQDISC, the API crossing exit suppressed the call.

Completion Code

MQCC_FAILED

Programmer response

Obey the rules for MQI calls that the exit enforces. To find out the rules, see the writer of the exit.

2110 (083E) (RC2110): MQRC_FORMAT_ERROR:

Explanation

An MQGET call was issued with the MQGMO_CONVERT option specified in the *GetMsgOpts* parameter, but the message cannot be converted successfully due to an error associated with the message format. Possible errors include:

- The format name in the message is MQFMT_NONE.
- A user-written exit with the name specified by the *Format* field in the message cannot be found.
- The message contains data that is not consistent with the format definition.

The message is returned unconverted to the application issuing the MQGET call, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

Completion Code

MQCC_WARNING

Programmer response

Check the format name that was specified when the message was put. If this is not one of the built-in formats, check that a suitable exit with the same name as the format is available for the queue manager to load. Verify that the data in the message corresponds to the format expected by the exit.

2111 (083F) (RC2111): MQRC_SOURCE_CCSID_ERROR:

Explanation

The coded character-set identifier from which character data is to be converted is not valid or not supported.

This can occur on the MQGET call when the MQGMO_CONVERT option is included in the *GetMsgOpts* parameter; the coded character-set identifier in error is the *CodedCharSetId* field in the message being retrieved. In this case, the message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

This reason can also occur on the MQGET call when the message contains one or more MQ header structures (MQCIH, MQDLH, MQIIH, MQRMH), and the *CodedCharSetId* field in the message specifies a character set that does not have SBCS characters for the characters that are valid in queue names. MQ header structures containing such characters are not valid, and so the message is returned unconverted. The Unicode character set UCS-2 is an example of such a character set.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

This reason can also occur on the MQXCNVC call; the coded character-set identifier in error is the *SourceCCSID* parameter. Either the *SourceCCSID* parameter specifies a value that is not valid or not supported, or the *SourceCCSID* parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason can also occur on a MQSETMP/MQINQMP/MQDLTMP call when the application issuing the calls does not use Language Environment (LE) and defines CCSID values of MQCCSI_APPL (-3) for message property names and string property values.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Check the character-set identifier that was specified when the message was put, or that was specified for the *SourceCCSID* parameter on the MQXCNVC call. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the specified character set, conversion must be carried out by the application.

If this reason happens as result of a MQSETMP/MQINQMP/MQDLTMP call issued in a non-LE application program that has specified CCSID as MQCCSI_APPL (-3) then applications should be changed to specify the CCSID value used by the application to encode the property names or property string values.

It is recommended that applications override the value of MQCCSI_APPL (-3) with the correct CCSID used as described in  Redefinition of MQCCSI_APPL (*WebSphere MQ V7.1 Reference*) or to set the explicit CCSID value used to encode text strings in MQCHARV or similar structures.

2112 (0840) (RC2112): MQRC_SOURCE_INTEGER_ENC_ERROR:

Explanation

On an MQGET call, with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the message being retrieved specifies an integer encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

This reason code can also occur on the MQXCNVC call, when the *Options* parameter contains an unsupported MQDCC_SOURCE_* value, or when MQDCC_SOURCE_ENC_UNDEFINED is specified for a UCS-2 code page.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Check the integer encoding that was specified when the message was put. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required integer encoding, conversion must be carried out by the application.

2113 (0841) (RC2113): MQRC_SOURCE_DECIMAL_ENC_ERROR:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the message being retrieved specifies a decimal encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

Completion Code

MQCC_WARNING

Programmer response

Check the decimal encoding that was specified when the message was put. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required decimal encoding, conversion must be carried out by the application.

2114 (0842) (RC2114): MQRC_SOURCE_FLOAT_ENC_ERROR:

Explanation

On an MQGET call, with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the message being retrieved specifies a floating-point encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

Completion Code

MQCC_WARNING

Programmer response

Check the floating-point encoding that was specified when the message was put. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required floating-point encoding, conversion must be carried out by the application.

2115 (0843) (RC2115): MQRC_TARGET_CCSID_ERROR:

Explanation

The coded character-set identifier to which character data is to be converted is not valid or not supported.

This can occur on the MQGET call when the MQGMO_CONVERT option is included in the *GetMsgOpts* parameter; the coded character-set identifier in error is the *CodedCharSetId* field in the *MsgDesc* parameter. In this case, the message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

This reason can also occur on the MQGET call when the message contains one or more MQ header structures (MQCIH, MQDLH, MQIIH, MQRMH), and the *CodedCharSetId* field in the *MsgDesc* parameter specifies a character set that does not have SBCS characters for the characters that are valid in queue names. The Unicode character set UCS-2 is an example of such a character set.

This reason can also occur on the MQXCNVC call; the coded character-set identifier in error is the *TargetCCSID* parameter. Either the *TargetCCSID* parameter specifies a value that is not valid or not supported, or the *TargetCCSID* parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Check the character-set identifier that was specified for the *CodedCharSetId* field in the *MsgDesc* parameter on the MQGET call, or that was specified for the *SourceCCSID* parameter on the MQXCNVC call. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the specified character set, conversion must be carried out by the application.

2116 (0844) (RC2116): MQRC_TARGET_INTEGER_ENC_ERROR:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the *MsgDesc* parameter specifies an integer encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message being retrieved, and the call completes with MQCC_WARNING.

This reason code can also occur on the MQXCNVC call, when the *Options* parameter contains an unsupported MQDCC_TARGET_* value, or when MQDCC_TARGET_ENC_UNDEFINED is specified for a UCS-2 code page.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Check the integer encoding that was specified. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required integer encoding, conversion must be carried out by the application.

2117 (0845) (RC2117): MQRC_TARGET_DECIMAL_ENC_ERROR:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the *MsgDesc* parameter specifies a decimal encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

Completion Code

MQCC_WARNING

Programmer response

Check the decimal encoding that was specified. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required decimal encoding, conversion must be carried out by the application.

2118 (0846) (RC2118): MQRC_TARGET_FLOAT_ENC_ERROR:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the *Encoding* value in the *MsgDesc* parameter specifies a floating-point encoding that is not recognized. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

Completion Code

MQCC_WARNING

Programmer response

Check the floating-point encoding that was specified. If this is correct, check that it is one for which queue-manager conversion is supported. If queue-manager conversion is not supported for the required floating-point encoding, conversion must be carried out by the application.

2119 (0847) (RC2119): MQRC_NOT_CONVERTED:

Explanation

An MQGET call was issued with the MQGMO_CONVERT option specified in the *GetMsgOpts* parameter, but an error occurred during conversion of the data in the message. The message data is returned unconverted, the values of the *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter are set to those of the message returned, and the call completes with MQCC_WARNING.

If the message consists of several parts, each of which is described by its own *CodedCharSetId* and *Encoding* fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various *CodedCharSetId* and *Encoding* fields always correctly describe the relevant message data.

This error may also indicate that a parameter to the data-conversion service is not supported.

Completion Code

MQCC_WARNING

Programmer response

Check that the message data is correctly described by the *Format*, *CodedCharSetId* and *Encoding* parameters that were specified when the message was put. Also check that these values, and the *CodedCharSetId* and *Encoding* specified in the *MsgDesc* parameter on the MQGET call, are supported for queue-manager conversion. If the required conversion is not supported, conversion must be carried out by the application.

2120 (0848) (RC2120): MQRC_CONVERTED_MSG_TOO_BIG:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, the message data expanded during data conversion and exceeded the size of the buffer provided by the application. However, the message had already been removed from the queue because prior to conversion the message data could be accommodated in the application buffer without truncation.

The message is returned unconverted, with the *CompCode* parameter of the MQGET call set to MQCC_WARNING. If the message consists of several parts, each of which is described by its own character-set and encoding fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts may be converted and other parts not converted. However, the values returned in the various character-set and encoding fields always correctly describe the relevant message data.

This reason can also occur on the MQXCNVC call, when the *TargetBuffer* parameter is too small to accommodate the converted string, and the string has been truncated to fit in the buffer. The length of valid data returned is given by the *DataLength* parameter; in the case of a DBCS string or mixed SBCS/DBCS string, this length may be *less than* the length of *TargetBuffer*.

Completion Code

MQCC_WARNING

Programmer response

For the MQGET call, check that the exit is converting the message data correctly and setting the output length *DataLength* to the appropriate value. If it is, the application issuing the MQGET call must provide a larger buffer for the *Buffer* parameter.

For the MQXCNVC call, if the string must be converted without truncation, provide a larger output buffer.

2121 (0849) (RC2121): MQRC_NO_EXTERNAL_PARTICIPANTS:

Explanation

An MQBEGIN call was issued to start a unit of work coordinated by the queue manager, but no participating resource managers have been registered with the queue manager. As a result, only changes to MQ resources can be coordinated by the queue manager in the unit of work.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows.

Completion Code

MQCC_WARNING

Programmer response

If the application does not require non-MQ resources to participate in the unit of work, this reason code can be ignored or the MQBEGIN call removed. Otherwise consult your system programmer to determine why the required resource managers have not been registered with the queue manager; the queue manager's configuration file might be in error.

2122 (084A) (RC2122): MQRC_PARTICIPANT_NOT_AVAILABLE:

Explanation

An MQBEGIN call was issued to start a unit of work coordinated by the queue manager, but one or more of the participating resource managers that had been registered with the queue manager is not available. As a result, changes to those resources cannot be coordinated by the queue manager in the unit of work.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows.

Completion Code

MQCC_WARNING

Programmer response

If the application does not require non-MQ resources to participate in the unit of work, this reason code can be ignored. Otherwise consult your system programmer to determine why the required resource managers are not available. The resource manager might have been halted temporarily, or there might be an error in the queue manager's configuration file.

2123 (084B) (RC2123): MQRC_OUTCOME_MIXED:

Explanation

The queue manager is acting as the unit-of-work coordinator for a unit of work that involves other resource managers, but one of the following occurred:

- An MQCMIT or MQDISC call was issued to commit the unit of work, but one or more of the participating resource managers backed-out the unit of work instead of committing it. As a result, the outcome of the unit of work is mixed.
- An MQBACK call was issued to back out a unit of work, but one or more of the participating resource managers had already committed the unit of work.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Examine the queue-manager error logs for messages relating to the mixed outcome; these messages identify the resource managers that are affected. Use procedures local to the affected resource managers to resynchronize the resources.

This reason code does not prevent the application initiating further units of work.

2124 (084C) (RC2124): MQRC_OUTCOME_PENDING:

Explanation

The queue manager is acting as the unit-of-work coordinator for a unit of work that involves other resource managers, and an MQCMIT or MQDISC call was issued to commit the unit of work, but one or more of the participating resource managers has not confirmed that the unit of work was committed successfully.

The completion of the commit operation will happen at some point in the future, but there remains the possibility that the outcome will be mixed.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_WARNING

Programmer response

Use the normal error-reporting mechanisms to determine whether the outcome was mixed. If it was, take appropriate action to resynchronize the resources.

This reason code does not prevent the application initiating further units of work.

2125 (084D) (RC2125): MQRC_BRIDGE_STARTED:

Explanation

The IMS bridge has been started.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2126 (084E) (RC2126): MQRC_BRIDGE_STOPPED:

Explanation

The IMS bridge has been stopped.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2127 (084F) (RC2127): MQRC_ADAPTER_STORAGE_SHORTAGE:**Explanation**

On an MQCONN call, the adapter was unable to acquire storage.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Notify the system programmer. The system programmer should determine why the system is short on storage, and take appropriate action, for example, increase the region size on the step or job card.

2128 (0850) (RC2128): MQRC_UOW_IN_PROGRESS:**Explanation**

An MQBEGIN call was issued to start a unit of work coordinated by the queue manager, but a unit of work is already in existence for the connection handle specified. This may be a global unit of work started by a previous MQBEGIN call, or a unit of work that is local to the queue manager or one of the cooperating resource managers. No more than one unit of work can exist concurrently for a connection handle.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Review the application logic to determine why there is a unit of work already in existence. Move the MQBEGIN call to the appropriate place in the application.

2129 (0851) (RC2129): MQRC_ADAPTER_CONN_LOAD_ERROR:**Explanation**

On an MQCONN call, the connection handling module could not be loaded, so the adapter could not link to it. The connection handling module name is:

- CSQBICON for batch applications
- CSQQCONN or CSQQCON2 for IMS applications

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified in the batch application program execution JCL, and in the queue-manager startup JCL.

2130 (0852) (RC2130): MQRC_ADAPTER_SERV_LOAD_ERROR:

Explanation

On an MQI call, the batch adapter could not load one of the following API service module, and so could not link to it:

- CSQBSRV
- CSQAPEPL
- CSQBCRMH
- CSQBAPPL

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified in the batch application program execution JCL, and in the queue-manager startup JCL.

2131 (0853) (RC2131): MQRC_ADAPTER_DEFS_ERROR:

Explanation

On an MQCONN call, the subsystem definition module (CSQBDEFV for batch and CSQQDEFV for IMS) does not contain the required control block identifier.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check your library concatenation. If this is correct, check that the CSQBDEFV or CSQQDEFV module contains the required subsystem ID.

2132 (0854) (RC2132): MQRC_ADAPTER_DEFS_LOAD_ERROR:

Explanation

On an MQCONN call, the subsystem definition module (CSQBDEFV for batch and CSQQDEFV for IMS) could not be loaded.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified in the application program execution JCL, and in the queue-manager startup JCL.

2133 (0855) (RC2133): MQRC_ADAPTER_CONV_LOAD_ERROR:

Explanation

On an MQGET call, the adapter (batch or IMS) could not load the data conversion services modules.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified in the batch application program execution JCL, and in the queue-manager startup JCL.

2134 (0856) (RC2134): MQRC_BO_ERROR:

Explanation

On an MQBEGIN call, the begin-options structure MQBO is not valid, for one of the following reasons:

- The *StrucId* field is not MQBO_STRUC_ID.
- The *Version* field is not MQBO_VERSION_1.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQBO structure are set correctly.

2135 (0857) (RC2135): MQRC_DH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQDH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQDH_STRUC_ID.
- The *Version* field is not MQDH_VERSION_1.
- The *StrucLength* field specifies a value that is too small to include the structure plus the arrays of MQOR and MQPMR records.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2136 (0858) (RC2136): MQRC_MULTIPLE_REASONS:

Explanation

An MQOPEN, MQPUT or MQPUT1 call was issued to open a distribution list or put a message to a distribution list, but the result of the call was not the same for all of the destinations in the list. One of the following applies:

- The call succeeded for some of the destinations but not others. The completion code is MQCC_WARNING in this case.
- The call failed for all of the destinations, but for differing reasons. The completion code is MQCC_FAILED in this case.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Examine the MQRR response records to identify the destinations for which the call failed, and the reason for the failure. Ensure that sufficient response records are provided by the application on the call to enable the error(s) to be determined. For the MQPUT1 call, the response records must be specified using the MQOD structure, and not the MQPMO structure.

2137 (0859) (RC2137): MQRC_OPEN_FAILED:

Explanation

A queue or other MQ object could not be opened successfully, for one of the following reasons:

- An MQCONN or MQCONNX call was issued, but the queue manager was unable to open an object that is used internally by the queue manager. As a result, processing cannot continue. The error log will contain the name of the object that could not be opened.
- An MQPUT call was issued to put a message to a distribution list, but the message could not be sent to the destination to which this reason code applies because that destination was not opened successfully by the MQOPEN call. This reason occurs only in the *Reason* field of the MQRR response record.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following:

- If the error occurred on the MQCONN or MQCONNX call, ensure that the required objects exist by running the following command and then retrying the application:

```
STRMQM -c qmgr
```

where *qmgr* should be replaced by the name of the queue manager.

- If the error occurred on the MQPUT call, examine the MQRR response records specified on the MQOPEN call to determine the reason that the queue failed to open. Ensure that sufficient response records are provided by the application on the call to enable the error(s) to be determined.

2138 (085A) (RC2138): MQRC_ADAPTER_DISC_LOAD_ERROR:

Explanation

On an MQDISC call, the disconnect handling module (CSQBDSC for batch and CSQQDISC for IMS) could not be loaded, so the adapter could not link to it.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified in the application program execution JCL, and in the queue-manager startup JCL. Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2139 (085B) (RC2139): MQRC_CNO_ERROR:

Explanation

On an MQCONNX call, the connect-options structure MQCNO is not valid, for one of the following reasons:

- The *StrucId* field is not MQCNO_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the parameter pointer points to read-only storage.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQCNO structure are set correctly.

2140 (085C) (RC2140): MQRC_CICS_WAIT_FAILED:

Explanation

On any MQI call, the CICS adapter issued an EXEC CICS WAIT request, but the request was rejected by CICS.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Examine the CICS trace data for actual response codes. The most likely cause is that the task has been canceled by the operator or by the system.

2141 (085D) (RC2141): MQRC_DLH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQDLH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQDLH_STRUC_ID.
- The *Version* field is not MQDLH_VERSION_1.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2142 (085E) (RC2142): MQRC_HEADER_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQ header structure that is not valid. Possible errors include the following:

- The *StrucId* field is not valid.
- The *Version* field is not valid.
- The *StrucLength* field specifies a value that is too small.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2143 (085F) (RC2143): MQRC_SOURCE_LENGTH_ERROR:

Explanation

On the MQXCNCV call, the *SourceLength* parameter specifies a length that is less than zero or not consistent with the string's character set or content (for example, the character set is a double-byte character set, but the length is not a multiple of two). This reason also occurs if the *SourceLength* parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason code can also occur on the MQGET call when the MQGMO_CONVERT option is specified. In this case it indicates that the MQRC_SOURCE_LENGTH_ERROR reason was returned by an MQXCNCV call issued by the data conversion exit.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Specify a length that is zero or greater. If the reason code occurs on the MQGET call, check that the logic in the data-conversion exit is correct.

2144 (0860) (RC2144): MQRC_TARGET_LENGTH_ERROR:

Explanation

On the MQXCNVC call, the *TargetLength* parameter is not valid for one of the following reasons:

- *TargetLength* is less than zero.
- The *TargetLength* parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The MQDCC_FILL_TARGET_BUFFER option is specified, but the value of *TargetLength* is such that the target buffer cannot be filled completely with valid characters. This can occur when *TargetCCSID* is a pure DBCS character set (such as UCS-2), but *TargetLength* specifies a length that is an odd number of bytes.

This reason code can also occur on the MQGET call when the MQGMO_CONVERT option is specified. In this case it indicates that the MQRC_TARGET_LENGTH_ERROR reason was returned by an MQXCNVC call issued by the data conversion exit.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Specify a length that is zero or greater. If the MQDCC_FILL_TARGET_BUFFER option is specified, and *TargetCCSID* is a pure DBCS character set, ensure that *TargetLength* specifies a length that is a multiple of two.

If the reason code occurs on the MQGET call, check that the logic in the data-conversion exit is correct.

2145 (0861) (RC2145): MQRC_SOURCE_BUFFER_ERROR:

Explanation

On the MQXCNVC call, the *SourceBuffer* parameter pointer is not valid, or points to storage that cannot be accessed for the entire length specified by *SourceLength*. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason code can also occur on the MQGET call when the MQGMO_CONVERT option is specified. In this case it indicates that the MQRC_SOURCE_BUFFER_ERROR reason was returned by an MQXCNVC call issued by the data conversion exit.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Specify a valid buffer. If the reason code occurs on the MQGET call, check that the logic in the data-conversion exit is correct.

2146 (0862) (RC2146): MQRC_TARGET_BUFFER_ERROR:

Explanation

On the MQXCNVC call, the *TargetBuffer* parameter pointer is not valid, or points to read-only storage, or to storage that cannot be accessed for the entire length specified by *TargetLength*. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

This reason code can also occur on the MQGET call when the MQGMO_CONVERT option is specified. In this case it indicates that the MQRC_TARGET_BUFFER_ERROR reason was returned by an MQXCNVC call issued by the data conversion exit.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Specify a valid buffer. If the reason code occurs on the MQGET call, check that the logic in the data-conversion exit is correct.

2148 (0864) (RC2148): MQRC_IIH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQIIH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQIIH_STRUC_ID.
- The *Version* field is not MQIIH_VERSION_1.
- The *StrucLength* field is not MQIIH_LENGTH_1.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2149 (0865) (RC2149): MQRC_PCF_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued to put a message containing PCF data, but the length of the message does not equal the sum of the lengths of the PCF structures present in the message. This can occur for messages with the following format names:

- MQFMT_ADMIN
- MQFMT_EVENT
- MQFMT_PCF

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the length of the message specified on the MQPUT or MQPUT1 call equals the sum of the lengths of the PCF structures contained within the message data.

2150 (0866) (RC2150): MQRC_DBCS_ERROR:

Explanation

An error was encountered attempting to convert a double-byte character set (DBCS) string. This can occur in the following cases:

- On the MQXCNVC call, when the *SourceCCSID* parameter specifies the coded character-set identifier of a double-byte character set, but the *SourceBuffer* parameter does not contain a valid DBCS string. This may be because the string contains characters that are not valid DBCS characters, or because the string is a mixed SBCS/DBCS string and the shift-out/shift-in characters are not correctly paired. The completion code is MQCC_FAILED in this case.
- On the MQGET call, when the MQGMO_CONVERT option is specified. In this case it indicates that the MQRC_DBCS_ERROR reason code was returned by an MQXCNVC call issued by the data conversion exit. The completion code is MQCC_WARNING in this case.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Specify a valid string.

If the reason code occurs on the MQGET call, check that the data in the message is valid, and that the logic in the data-conversion exit is correct.

2152 (0868) (RC2152): MQRC_OBJECT_NAME_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued to open a distribution list (that is, the *RecsPresent* field in MQOD is greater than zero), but the *ObjectName* field is neither blank nor the null string.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

If it is intended to open a distribution list, set the *ObjectName* field to blanks or the null string. If it is not intended to open a distribution list, set the *RecsPresent* field to zero.

2153 (0869) (RC2153): MQRC_OBJECT_Q_MGR_NAME_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued to open a distribution list (that is, the *RecsPresent* field in MQOD is greater than zero), but the *ObjectQMgrName* field is neither blank nor the null string.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

If it is intended to open a distribution list, set the *ObjectQMgrName* field to blanks or the null string. If it is not intended to open a distribution list, set the *RecsPresent* field to zero.

2154 (086A) (RC2154): MQRC_RECS_PRESENT_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued, but the call failed for one of the following reasons:

- *RecsPresent* in MQOD is less than zero.
- *ObjectType* in MQOD is not MQOT_Q, and *RecsPresent* is not zero. *RecsPresent* must be zero if the object being opened is not a queue.
- WebSphere MQ Multicast is being used and *RecsPresent* in MQOD is not set to zero. WebSphere MQ Multicast does not use distribution lists.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

If it is intended to open a distribution list, set the *ObjectType* field to MQOT_Q and *RecsPresent* to the number of destinations in the list. If it is not intended to open a distribution list, set the *RecsPresent* field to zero.

2155 (086B) (RC2155): MQRC_OBJECT_RECORDS_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued to open a distribution list (that is, the *RecsPresent* field in MQOD is greater than zero), but the MQOR object records are not specified correctly. One of the following applies:

- *ObjectRecOffset* is zero and *ObjectRecPtr* is zero or the null pointer.
- *ObjectRecOffset* is not zero and *ObjectRecPtr* is not zero and not the null pointer.
- *ObjectRecPtr* is not a valid pointer.
- *ObjectRecPtr* or *ObjectRecOffset* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that one of *ObjectRecOffset* and *ObjectRecPtr* is zero and the other nonzero. Ensure that the field used points to accessible storage.

2156 (086C) (RC2156): MQRC_RESPONSE_RECORDS_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued to open a distribution list (that is, the *RecsPresent* field in MQOD is greater than zero), but the MQRR response records are not specified correctly. One of the following applies:

- *ResponseRecOffset* is not zero and *ResponseRecPtr* is not zero and not the null pointer.
- *ResponseRecPtr* is not a valid pointer.
- *ResponseRecPtr* or *ResponseRecOffset* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that at least one of *ResponseRecOffset* and *ResponseRecPtr* is zero. Ensure that the field used points to accessible storage.

2157 (086D) (RC2157): MQRC_ASID_MISMATCH:

Explanation

On any MQI call, the caller's primary ASID was found to be different from the home ASID.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Correct the application (MQI calls cannot be issued in cross-memory mode). Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2158 (086E) (RC2158): MQRC_PMO_RECORD_FLAGS_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued to put a message, but the *PutMsgRecFields* field in the MQPMO structure is not valid, for one of the following reasons:

- The field contains flags that are not valid.
- The message is being put to a distribution list, and put message records have been provided (that is, *RecsPresent* is greater than zero, and one of *PutMsgRecOffset* or *PutMsgRecPtr* is nonzero), but *PutMsgRecFields* has the value MQPMRF_NONE.
- MQPMRF_ACCOUNTING_TOKEN is specified without either MQPMO_SET_IDENTITY_CONTEXT or MQPMO_SET_ALL_CONTEXT.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that *PutMsgRecFields* is set with the appropriate MQPMRF_* flags to indicate which fields are present in the put message records. If MQPMRF_ACCOUNTING_TOKEN is specified, ensure that either MQPMO_SET_IDENTITY_CONTEXT or MQPMO_SET_ALL_CONTEXT is also specified. Alternatively, set both *PutMsgRecOffset* and *PutMsgRecPtr* to zero.

2159 (086F) (RC2159): MQRC_PUT_MSG_RECORDS_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued to put a message to a distribution list, but the MQPMR put message records are not specified correctly. One of the following applies:

- *PutMsgRecOffset* is not zero and *PutMsgRecPtr* is not zero and not the null pointer.
- *PutMsgRecPtr* is not a valid pointer.
- *PutMsgRecPtr* or *PutMsgRecOffset* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that at least one of *PutMsgRecOffset* and *PutMsgRecPtr* is zero. Ensure that the field used points to accessible storage.

2160 (0870) (RC2160): MQRC_CONN_ID_IN_USE:

Explanation

On an MQCONN call, the connection identifier assigned by the queue manager to the connection between a CICS or IMS allied address space and the queue manager conflicts with the connection identifier of another connected CICS or IMS system. The connection identifier assigned is as follows:

- For CICS, the applid
- For IMS, the IMSID parameter on the IMSCTRL (sysgen) macro, or the IMSID parameter on the execution parameter (EXEC card in IMS control region JCL)
- For batch, the job name
- For TSO, the user ID

A conflict arises only if there are two CICS systems, two IMS systems, or one each of CICS and IMS, having the same connection identifiers. Batch and TSO connections need not have unique identifiers.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the naming conventions used in different systems that might connect to the queue manager do not conflict.

2161 (0871) (RC2161): MQRC_Q_MGR QUIESCING:

Explanation

An MQI call was issued, but the call failed because the queue manager is quiescing (preparing to shut down).

When the queue manager is quiescing, the MQOPEN, MQPUT, MQPUT1, and MQGET calls can still complete successfully, but the application can request that they fail by specifying the appropriate option on the call:

- MQOO_FAIL_IF QUIESCING on MQOPEN
- MQPMO_FAIL_IF QUIESCING on MQPUT or MQPUT1
- MQGMO_FAIL_IF QUIESCING on MQGET

Specifying these options enables the application to become aware that the queue manager is preparing to shut down.

- On z/OS:
 - For batch applications, this reason can be returned to applications running in LPARs that do not have a queue manager installed.
 - For CICS applications, this reason can be returned when no connection was established.
- On IBM i for applications running in compatibility mode, this reason can be returned when no connection was established.

Completion Code

MQCC_FAILED

Programmer response

The application should tidy up and end. If the application specified the MQOO_FAIL_IF QUIESCING, MQPMO_FAIL_IF QUIESCING, or MQGMO_FAIL_IF QUIESCING option on the failing call, the relevant option can be removed and the call reissued. By omitting these options, the application can continue working to complete and commit the current unit of work, but the application does not start a new unit of work.

2162 (0872) (RC2162): MQRC_Q_MGR_STOPPING:

Explanation

An MQI call was issued, but the call failed because the queue manager is shutting down. If the call was an MQGET call with the MQGMO_WAIT option, the wait has been canceled. No more MQI calls can be issued.

For MQ MQI client applications, it is possible that the call did complete successfully, even though this reason code is returned with a *CompCode* of MQCC_FAILED.

- On z/OS, the MQRC_CONNECTION_BROKEN reason may be returned instead if, as a result of system scheduling factors, the queue manager shuts down before the call completes.

Completion Code

MQCC_FAILED

Programmer response

The application should tidy up and end. If the application is in the middle of a unit of work coordinated by an external unit-of-work coordinator, the application should issue the appropriate call to back out the unit of work. Any unit of work that is coordinated by the queue manager is backed out automatically.

2163 (0873) (RC2163): MQRC_DUPLICATE_RECOV_COORD:

Explanation

On an MQCONN or MQCONNX call, a recovery coordinator already exists for the connection name specified on the connection call issued by the adapter.

A conflict arises only if there are two CICS systems, two IMS systems, or one each of CICS and IMS, having the same connection identifiers. Batch and TSO connections need not have unique identifiers.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the naming conventions used in different systems that might connect to the queue manager do not conflict.

2173 (087D) (RC2173): MQRC_PMO_ERROR:

Explanation

On an MQPUT or MQPUT1 call, the MQPMO structure is not valid, for one of the following reasons:

- The *StrucId* field is not MQPMO_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQPMO structure are set correctly.

2182 (0886) (RC2182): MQRC_API_EXIT_NOT_FOUND:

Explanation

The API crossing exit entry point could not be found.

Completion Code

MQCC_FAILED

Programmer response

Check the entry point name is valid for the library module.

2183 (0887) (RC2183): MQRC_API_EXIT_LOAD_ERROR:

Explanation

The API crossing exit module could not be linked. If this reason is returned when the API crossing exit is invoked *after* the call has been executed, the call itself might have executed correctly.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library concatenation has been specified, and that the API crossing exit module is executable and correctly named. Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2184 (0888) (RC2184): MQRC_REMOTE_Q_NAME_ERROR:

Explanation

On an MQOPEN or MQPUT1 call, one of the following occurred:

- A local definition of a remote queue (or an alias to one) was specified, but the *RemoteQName* attribute in the remote queue definition is entirely blank. Note that this error occurs even if the *XmitQName* in the definition is not blank.
- The *ObjectQMgrName* field in the object descriptor is not blank and not the name of the local queue manager, but the *ObjectName* field is blank.

Completion Code

MQCC_FAILED

Programmer response

Alter the local definition of the remote queue and supply a valid remote queue name, or supply a nonblank *ObjectName* in the object descriptor, as appropriate.

2185 (0889) (RC2185): MQRC_INCONSISTENT_PERSISTENCE:

Explanation

An MQPUT call was issued to put a message in a group or a segment of a logical message, but the value specified or defaulted for the *Persistence* field in MQMD is not consistent with the current group and segment information retained by the queue manager for the queue handle. All messages in a group and all segments in a logical message must have the same value for persistence, that is, all must be persistent, or all must be nonpersistent.

If the current call specifies MQPMO_LOGICAL_ORDER, the call fails. If the current call does not specify MQPMO_LOGICAL_ORDER, but the previous MQPUT call for the queue handle did, the call succeeds with completion code MQCC_WARNING.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Modify the application to ensure that the same value of persistence is used for all messages in the group, or all segments of the logical message.

2186 (088A) (RC2186): MQRC_GMO_ERROR:

Explanation

On an MQGET call, the MQGMO structure is not valid, for one of the following reasons:

- The *StrucId* field is not MQGMO_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQGMO structure are set correctly.

2187 (088B) (RC2187): MQRC_CICS_BRIDGE_RESTRICTION:

Explanation

It is not permitted to issue MQI calls from user transactions that are run in an MQ/CICS-bridge environment where the bridge exit also issues MQI calls. The MQI call fails. If it occurs in the bridge exit, it results in a transaction abend. If it occurs in the user transaction, it can result in a transaction abend.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The transaction cannot be run using the MQ/CICS bridge. Refer to the appropriate CICS manual for information about restrictions in the MQ/CICS bridge environment.

2188 (088C) (RC2188): MQRC_STOPPED_BY_CLUSTER_EXIT:

Explanation

An MQOPEN, MQPUT, or MQPUT1 call was issued to open or put a message on a cluster queue, but the cluster workload exit rejected the call.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check the cluster workload exit to ensure that it has been written correctly. Determine why it rejected the call and correct the problem.

2189 (088D) (RC2189): MQRC_CLUSTER_RESOLUTION_ERROR:

Explanation

An MQOPEN, MQPUT, or MQPUT1 call was issued to open or put a message on a cluster queue, but the queue definition could not be resolved correctly because a response was required from the repository manager but none was available.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the repository manager is operating and that the queue and channel definitions are correct.

2190 (088E) (RC2190): MQRC_CONVERTED_STRING_TOO_BIG:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, a string in a fixed-length field in the message expanded during data conversion and exceeded the size of the field. When this happens, the queue manager tries discarding trailing blank characters and characters following the first null character to make the string fit, but in this case there were insufficient characters that could be discarded.

This reason code can also occur for messages with a format name of MQFMT_IMS_VAR_STRING. When this happens, it indicates that the IMS variable string expanded such that its length exceeded the capacity of the 2 byte binary length field contained within the structure of the IMS variable string. (The queue manager never discards trailing blanks in an IMS variable string.)

The message is returned unconverted, with the *CompCode* parameter of the MQGET call set to MQCC_WARNING. If the message consists of several parts, each of which is described by its own character-set and encoding fields (for example, a message with format name MQFMT_DEAD_LETTER_HEADER), some parts might be converted and other parts not converted. However, the values returned in the various character-set and encoding fields always correctly describe the relevant message data.

This reason code does not occur if the string can be made to fit by discarding trailing blank characters.

Completion Code

MQCC_WARNING

Programmer response

Check that the fields in the message contain the correct values, and that the character-set identifiers specified by the sender and receiver of the message are correct. If they are, the layout of the data in the

message must be modified to increase the lengths of the field, or fields so that there is sufficient space to permit the string, or strings to expand when converted.

2191 (088F) (RC2191): MQRC_TMC_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQTMC2 structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQTMC_STRUC_ID.
- The *Version* field is not MQTMC_VERSION_2.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2192 (0890) (RC2192): MQRC_PAGESET_FULL:

Explanation

Former name for MQRC_STORAGE_MEDIUM_FULL.

2192 (0890) (RC2192): MQRC_STORAGE_MEDIUM_FULL:

Explanation

An MQI call or command was issued to operate on an object, but the call failed because the external storage medium is full. One of the following applies:

- A page-set data set is full (nonshared queues only).
- A coupling-facility structure is full (shared queues only).

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check which queues contain messages and look for applications that might be filling the queues unintentionally. Be aware that the queue that has caused the page set or coupling-facility structure to become full is not necessarily the queue referenced by the MQI call that returned MQRC_STORAGE_MEDIUM_FULL.

Check that all of the usual server applications are operating correctly and processing the messages on the queues.

If the applications and servers are operating correctly, increase the number of server applications to cope with the message load, or request the system programmer to increase the size of the page-set data sets.

2193 (0891) (RC2193): MQRC_PAGESET_ERROR:

Explanation

An error was encountered with the page set while attempting to access it for a locally defined queue. This could be because the queue is on a page set that does not exist. A console message is issued that tells you the number of the page set in error. For example if the error occurred in the TEST job, and your user identifier is ABCDEFG, the message is:

```
CSQI041I CSQIALLC JOB TEST USER ABCDEFG HAD ERROR ACCESSING PAGE SET 27
```

If this reason code occurs while attempting to delete a dynamic queue with MQCLOSE, the dynamic queue has not been deleted.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check that the storage class for the queue maps to a valid page set using the DISPLAY Q(xx) STGCLASS, DISPLAY STGCLASS(xx), and DISPLAY USAGE PSID commands. If you are unable to resolve the problem, notify the system programmer who should:

- Collect the following diagnostic information:
 - A description of the actions that led to the error
 - A listing of the application program being run at the time of the error
 - Details of the page sets defined for use by the queue manager
- Attempt to re-create the problem, and take a system dump immediately after the error occurs
- Contact your IBM Support Center

2194 (0892) (RC2194): MQRC_NAME_NOT_VALID_FOR_TYPE:

Explanation

An MQOPEN call was issued to open the queue manager definition, but the *ObjectName* field in the *ObjDesc* parameter is not blank.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the *ObjectName* field is set to blanks.

2195 (0893) (RC2195): MQRC_UNEXPECTED_ERROR:

Explanation

The call was rejected because an unexpected error occurred.

Completion Code

MQCC_FAILED

Programmer response

Check the application's parameter list to ensure, for example, that the correct number of parameters was passed, and that data pointers and storage keys are valid. If the problem cannot be resolved, contact your system programmer.

- On z/OS, check the joblog and logrec, and whether any information has been displayed on the console. If this error occurs on an MQCONN or MQCONNEX call, check that the subsystem named is an active MQ subsystem. In particular, check that it is not a Db2 subsystem. If the problem cannot be resolved, rerun the application with a CSQSNAP DD card (if you have not already got a dump) and send the resulting dump to IBM.
- On IBM i, consult the FFST record to obtain more detail about the problem.
- On HP OpenVMS, HP Integrity NonStop Server, and UNIX systems, consult the FDC file to obtain more detail about the problem.

2196 (0894) (RC2196): MQRC_UNKNOWN_XMIT_Q:

Explanation

On an MQOPEN or MQPUT1 call, a message is to be sent to a remote queue manager. The *ObjectName* or the *ObjectQMGrName* in the object descriptor specifies the name of a local definition of a remote queue (in the latter case queue-manager aliasing is being used), but the *XmitQName* attribute of the definition is not blank and not the name of a locally-defined queue.

Completion Code

MQCC_FAILED

Programmer response

Check the values specified for *ObjectName* and *ObjectQMGrName*. If these are correct, check the queue definitions.

2197 (0895) (RC2197): MQRC_UNKNOWN_DEF_XMIT_Q:

Explanation

An MQOPEN or MQPUT1 call was issued specifying a remote queue as the destination. If a local definition of the remote queue was specified, or if a queue-manager alias is being resolved, the *XmitQName* attribute in the local definition is blank.

Because there is no queue defined with the same name as the destination queue manager, the queue manager has attempted to use the default transmission queue. However, the name defined by the *DefXmitQName* queue-manager attribute is not the name of a locally-defined queue.

Completion Code

MQCC_FAILED

Programmer response

Correct the queue definitions, or the queue-manager attribute.

2198 (0896) (RC2198): MQRC_DEF_XMIT_Q_TYPE_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued specifying a remote queue as the destination. Either a local definition of the remote queue was specified, or a queue-manager alias was being resolved, but in either case the *XmitQName* attribute in the local definition is blank.

Because there is no transmission queue defined with the same name as the destination queue manager, the local queue manager has attempted to use the default transmission queue. However, although there is a queue defined by the *DefXmitQName* queue-manager attribute, it is not a local queue.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following:

- Specify a local transmission queue as the value of the *XmitQName* attribute in the local definition of the remote queue.
- Define a local transmission queue with a name that is the same as that of the remote queue manager.
- Specify a local transmission queue as the value of the *DefXmitQName* queue-manager attribute.

See  *XmitQName (MQCHAR48)* (*WebSphere MQ V7.1 Reference*) for more information about transmission queue names.

2199 (0897) (RC2199): MQRC_DEF_XMIT_Q_USAGE_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued specifying a remote queue as the destination. Either a local definition of the remote queue was specified, or a queue-manager alias was being resolved, but in either case the *XmitQName* attribute in the local definition is blank.

Because there is no transmission queue defined with the same name as the destination queue manager, the local queue manager has attempted to use the default transmission queue. However, the queue defined by the *DefXmitQName* queue-manager attribute does not have a *Usage* attribute of MQUS_TRANSMISSION.

This reason code is returned from MQOPEN or MQPUT1, if the queue manager's Default Transmission Queue is about to be used, but the name of this queue is SYSTEM.CLUSTER.TRANSMIT.QUEUE. This queue is reserved for clustering, so it is not valid to set the queue manager's Default Transmission Queue to this name.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following:

- Specify a local transmission queue as the value of the *XmitQName* attribute in the local definition of the remote queue.
- Define a local transmission queue with a name that is the same as that of the remote queue manager.
- Specify a different local transmission queue as the value of the *DefXmitQName* queue-manager attribute.
- Change the *Usage* attribute of the *DefXmitQName* queue to MQUS_TRANSMISSION.

See  *XmitQName (MQCHAR48)* (*WebSphere MQ V7.1 Reference*) for more information about transmission queue names.

2201 (0899) (RC2201): MQRC_NAME_IN_USE:

Explanation

An MQOPEN call was issued to create a dynamic queue, but a queue with the same name as the dynamic queue already exists. The existing queue is one that is logically deleted, but for which there are still one or more open handles. For more information, see  MQCLOSE – Close object (*WebSphere MQ V7.1 Reference*).

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Either ensure that all handles for the previous dynamic queue are closed, or ensure that the name of the new queue is unique; see the description for reason code MQRC_OBJECT_ALREADY_EXISTS.

2202 (089A) (RC2202): MQRC_CONNECTION QUIESCING:

Explanation

This reason code is issued when the connection to the queue manager is in quiescing state, and an application issues one of the following calls:

- MQCONN or MQCONNX
- MQOPEN, with no connection established, or with MQOO_FAIL_IF QUIESCING included in the *Options* parameter
- MQGET, with MQGMO_FAIL_IF QUIESCING included in the *Options* field of the *GetMsgOpts* parameter
- MQPUT or MQPUT1, with MQPMO_FAIL_IF QUIESCING included in the *Options* field of the *PutMsgOpts* parameter

MQRC_CONNECTION QUIESCING is also issued by the message channel agent (MCA) when the queue manager is in quiescing state.

Completion Code

MQCC_FAILED

Programmer response

The application should tidy up and terminate. Any uncommitted changes in a unit of work should be backed out.

2203 (089B) (RC2203): MQRC_CONNECTION_STOPPING:

Explanation

This reason code is issued when the connection to the queue manager is shutting down, and the application issues an MQI call. No more message-queuing calls can be issued. For the MQGET call, if the MQGMO_WAIT option was specified, the wait is canceled.

Note that the MQRC_CONNECTION_BROKEN reason may be returned instead if, as a result of system scheduling factors, the queue manager shuts down before the call completes.

MQRC_CONNECTION_STOPPING is also issued by the message channel agent (MCA) when the queue manager is shutting down.

For MQ MQI client applications, it is possible that the call did complete successfully, even though this reason code is returned with a *CompCode* of MQCC_FAILED.

Completion Code

MQCC_FAILED

Programmer response

The application should tidy up and terminate. Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2204 (089C) (RC2204): MQRC_ADAPTER_NOT_AVAILABLE:

Explanation

This is issued only for CICS applications, if any call is issued and the CICS adapter (a Task Related User Exit) has been disabled, or has not been enabled.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The application should tidy up and terminate. Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2206 (089E) (RC2206): MQRC_MSG_ID_ERROR:

Explanation

An MQGET call was issued to retrieve a message using the message identifier as a selection criterion, but the call failed because selection by message identifier is not supported on this queue.

- On z/OS, the queue is a shared queue, but the *IndexType* queue attribute does not have an appropriate value:
 - If selection is by message identifier alone, *IndexType* must have the value MQIT_MSG_ID.
 - If selection is by message identifier and correlation identifier combined, *IndexType* must have the value MQIT_MSG_ID or MQIT_CORREL_ID. However, the match-any values of MQCI_NONE and MQMI_NONE respectively are exceptions to this rule, and result in the 2206 MQRC_MSG_ID_ERROR reason code.
- On HP Integrity NonStop Server, a key file is required but has not been defined.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following:

- Modify the application so that it does not use selection by message identifier: set the *MsgId* field to MQMI_NONE and do not specify MQMO_MATCH_MSG_ID in MQGMO.
- On z/OS, change the *IndexType* queue attribute to MQIT_MSG_ID.
- On HP Integrity NonStop Server, define a key file.

2207 (089F) (RC2207): MQRC_CORREL_ID_ERROR:

Explanation

An MQGET call was issued to retrieve a message using the correlation identifier as a selection criterion, but the call failed because selection by correlation identifier is not supported on this queue.

- On z/OS, the queue is a shared queue, but the *IndexType* queue attribute does not have an appropriate value:
 - If selection is by correlation identifier alone, *IndexType* must have the value MQIT_CORREL_ID.
 - If selection is by correlation identifier and message identifier combined, *IndexType* must have the value MQIT_CORREL_ID or MQIT_MSG_ID.
- On HP Integrity NonStop Server, a key file is required but has not been defined.

Completion Code

MQCC_FAILED

Programmer response

Do one of the following:

- On z/OS, change the *IndexType* queue attribute to MQIT_CORREL_ID.
- On HP Integrity NonStop Server, define a key file.
- Modify the application so that it does not use selection by correlation identifier: set the *CorrelId* field to MQCI_NONE and do not specify MQMO_MATCH_CORREL_ID in MQGMO.

2208 (08A0) (RC2208): MQRC_FILE_SYSTEM_ERROR:**Explanation**

An unexpected return code was received from the file system, in attempting to perform an operation on a queue.

This reason code occurs only on VSE/ESA.

Completion Code

MQCC_FAILED

Programmer response

Check the file system definition for the queue that was being accessed. For a VSAM file, check that the control interval is large enough for the maximum message length allowed for the queue.

2209 (08A1) (RC2209): MQRC_NO_MSG_LOCKED:**Explanation**

An MQGET call was issued with the MQGMO_UNLOCK option, but no message was currently locked.

Completion Code

MQCC_WARNING

Programmer response

Check that a message was locked by an earlier MQGET call with the MQGMO_LOCK option for the same handle, and that no intervening call has caused the message to become unlocked.

2210 (08A2) (RC2210): MQRC_SOAP_DOTNET_ERROR:**Explanation**

An exception from the .NET environment (as opposed to WebSphere MQ .NET) has been received and is included as an inner exception.

Completion Code

MQCC_FAILED

Programmer response

Refer to the .NET documentation for details about the inner exception. Follow the corrective action recommended there.

2211 (08A3) (RC2211): MQRC_SOAP_AXIS_ERROR:**Explanation**

An exception from the Axis environment has been received and is included as a chained exception.

Completion Code

MQCC_FAILED

Programmer response

Refer to the Axis documentation for details about the chained exception. Follow the corrective action recommended there.

2212 (08A4) (RC2212): MQRC_SOAP_URL_ERROR:**Explanation**

The SOAP URL has been specified incorrectly.

Completion Code

MQCC_FAILED

Programmer response

Correct the SOAP URL and rerun.

2217 (08A9) (RC2217): MQRC_CONNECTION_NOT_AUTHORIZED:**Explanation**

This reason code arises only for CICS applications. For these, connection to the queue manager is done by the adapter. If that connection fails because the CICS subsystem is not authorized to connect to the queue manager, this reason code is issued whenever an application running under that subsystem subsequently issues an MQI call.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the subsystem is authorized to connect to the queue manager.

2218 (08AA) (RC2218): MQRC_MSG_TOO_BIG_FOR_CHANNEL:

Explanation

A message was put to a remote queue, but the message is larger than the maximum message length allowed by the channel. This reason code is returned in the *Feedback* field in the message descriptor of a report message.

- On z/OS, this return code is issued only if you are not using CICS for distributed queuing. Otherwise, MQRC_MSG_TOO_BIG_FOR_Q_MGR is issued.

Completion Code

MQCC_FAILED

Programmer response

Check the channel definitions. Increase the maximum message length that the channel can accept, or break the message into several smaller messages.

2219 (08AB) (RC2219): MQRC_CALL_IN_PROGRESS:

Explanation

The application issued an MQI call whilst another MQI call was already being processed for that connection. Only one call per application connection can be processed at a time.

Concurrent calls can arise when an application uses multiple threads, or when an exit is invoked as part of the processing of an MQI call. For example, a data-conversion exit invoked as part of the processing of the MQGET call may try to issue an MQI call.

- On z/OS, concurrent calls can arise only with batch or IMS applications; an example is when a subtask ends while an MQI call is in progress (for example, an MQGET that is waiting), and there is an end-of-task exit routine that issues another MQI call.
- On Windows, concurrent calls can also arise if an MQI call is issued in response to a user message while another MQI call is in progress.
- If the application is using multiple threads with shared handles, MQRC_CALL_IN_PROGRESS occurs when the handle specified on the call is already in use by another thread and MQCNO_HANDLE_SHARE_NO_BLOCK was specified on the MQCONN call.

Completion Code

MQCC_FAILED

Programmer response

Ensure that an MQI call cannot be issued while another one is active. Do not issue MQI calls from within a data-conversion exit.

- On z/OS, if you want to provide a subtask to allow an application that is waiting for a message to arrive to be canceled, wait for the message by using MQGET with MQGMO_SET_SIGNAL, rather than MQGMO_WAIT.

2220 (08AC) (RC2220): MQRC_RMH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQRMH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQRMH_STRUC_ID.
- The *Version* field is not MQRMH_VERSION_1.
- The *StrucLength* field specifies a value that is too small to include the structure plus the variable-length data at the end of the structure.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2222 (08AE) (RC2222): MQRC_Q_MGR_ACTIVE:

Explanation

This condition is detected when a queue manager becomes active.

- On z/OS, this event is not generated for the first start of a queue manager, only on subsequent restarts.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2223 (08AF) (RC2223): MQRC_Q_MGR_NOT_ACTIVE:

Explanation

This condition is detected when a queue manager is requested to stop or quiesce.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2224 (08B0) (RC2224): MQRC_Q_DEPTH_HIGH:**Explanation**

An MQPUT or MQPUT1 call has caused the queue depth to be incremented to or above the limit specified in the *QDepthHighLimit* attribute.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2225 (08B1) (RC2225): MQRC_Q_DEPTH_LOW:**Explanation**

An MQGET call has caused the queue depth to be decremented to or below the limit specified in the *QDepthLowLimit* attribute.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2226 (08B2) (RC2226): MQRC_Q_SERVICE_INTERVAL_HIGH:**Explanation**

No successful gets or puts have been detected within an interval that is greater than the limit specified in the *QServiceInterval* attribute.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2227 (08B3) (RC2227): MQRC_Q_SERVICE_INTERVAL_OK:**Explanation**

A successful get has been detected within an interval that is less than or equal to the limit specified in the *QServiceInterval* attribute.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2228 (08B4) (RC2228): MQRC_RFH_HEADER_FIELD_ERROR:

Explanation

An expected RFH header field was not found or had an invalid value. If this error occurs in a WebSphere MQ SOAP listener, the missing or erroneous field is either the *contentType* field or the *transportVersion* field or both.

Completion Code

MQCC_FAILED

Programmer response

If this error occurs in a WebSphere MQ SOAP listener, and you are using the IBM-supplied sender, contact your IBM Support Center. If you are using a bespoke sender, check the associated error message, and that the RFH2 section of the SOAP/MQ request message contains all the mandatory fields, and that these fields have valid values.

2229 (08B5) (RC2229): MQRC_RAS_PROPERTY_ERROR:

Explanation

There is an error related to the RAS property file. The file might be missing, it might be not accessible, or the commands in the file might be incorrect.

Completion Code

MQCC_FAILED

Programmer response

Look at the associated error message, which explains the error in detail. Correct the error and try again.

2232 (08B8) (RC2232): MQRC_UNIT_OF_WORK_NOT_STARTED:

Explanation

An MQGET, MQPUT or MQPUT1 call was issued to get or put a message within a unit of work, but no TM/MP transaction had been started. If MQGMO_NO_SYNCPOINT is not specified on MQGET, or MQPMO_NO_SYNCPOINT is not specified on MQPUT or MQPUT1 (the default), the call requires a unit of work.

Completion Code

MQCC_FAILED

Programmer response

Ensure a TM/MP transaction is available, or issue the MQGET call with the MQGMO_NO_SYNCPOINT option, or the MQPUT or MQPUT1 call with the MQPMO_NO_SYNCPOINT option, which will cause a transaction to be started automatically.

2233 (08B9) (RC2233): MQRC_CHANNEL_AUTO_DEF_OK:**Explanation**

This condition is detected when the automatic definition of a channel is successful. The channel is defined by the MCA.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2234 (08BA) (RC2234): MQRC_CHANNEL_AUTO_DEF_ERROR:**Explanation**

This condition is detected when the automatic definition of a channel fails; this might be because an error occurred during the definition process, or because the channel automatic-definition exit inhibited the definition. Additional information is returned in the event message indicating the reason for the failure.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING

Programmer response

Examine the additional information returned in the event message to determine the reason for the failure.

2235 (08BB) (RC2235): MQRC_CFH_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFH structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2236 (08BC) (RC2236): MQRC_CFIL_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFIL or MQRCFIL64 structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2237 (08BD) (RC2237): MQRC_CFIN_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFIN or MQCFIN64 structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2238 (08BE) (RC2238): MQRC_CFSL_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFSL structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2239 (08BF) (RC2239): MQRC_CFST_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFST structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2241 (08C1) (RC2241): MQRC_INCOMPLETE_GROUP:

Explanation

An operation was attempted on a queue using a queue handle that had an incomplete message group. This reason code can arise in the following situations:

- On the MQPUT call, when the application specifies MQPMO_LOGICAL_ORDER and attempts to put a message that is not in a group. The completion code is MQCC_FAILED in this case.
- On the MQPUT call, when the application does *not* specify MQPMO_LOGICAL_ORDER, but the previous MQPUT call for the queue handle did specify MQPMO_LOGICAL_ORDER. The completion code is MQCC_WARNING in this case.
- On the MQGET call, when the application does *not* specify MQGMO_LOGICAL_ORDER, but the previous MQGET call for the queue handle did specify MQGMO_LOGICAL_ORDER. The completion code is MQCC_WARNING in this case.
- On the MQCLOSE call, when the application attempts to close the queue that has the incomplete message group. The completion code is MQCC_WARNING in this case.

If there is an incomplete logical message as well as an incomplete message group, reason code MQRC_INCOMPLETE_MSG is returned in preference to MQRC_INCOMPLETE_GROUP.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

If this reason code is expected, no corrective action is required. Otherwise, ensure that the MQPUT call for the last message in the group specifies MQMF_LAST_MSG_IN_GROUP.

2242 (08C2) (RC2242): MQRC_INCOMPLETE_MSG:

Explanation

An operation was attempted on a queue using a queue handle that had an incomplete logical message. This reason code can arise in the following situations:

- On the MQPUT call, when the application specifies MQPMO_LOGICAL_ORDER and attempts to put a message that is not a segment, or that has a setting for the MQMF_LAST_MSG_IN_GROUP flag that is different from the previous message. The completion code is MQCC_FAILED in this case.
- On the MQPUT call, when the application does *not* specify MQPMO_LOGICAL_ORDER, but the previous MQPUT call for the queue handle did specify MQPMO_LOGICAL_ORDER. The completion code is MQCC_WARNING in this case.
- On the MQGET call, when the application does *not* specify MQGMO_LOGICAL_ORDER, but the previous MQGET call for the queue handle did specify MQGMO_LOGICAL_ORDER. The completion code is MQCC_WARNING in this case.
- On the MQCLOSE call, when the application attempts to close the queue that has the incomplete logical message. The completion code is MQCC_WARNING in this case.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

If this reason code is expected, no corrective action is required. Otherwise, ensure that the MQPUT call for the last segment specifies MQMF_LAST_SEGMENT.

2243 (08C3) (RC2243): MQRC_INCONSISTENT_CCIDS:

Explanation

An MQGET call was issued specifying the MQGMO_COMPLETE_MSG option, but the message to be retrieved consists of two or more segments that have differing values for the *CodedCharSetId* field in MQMD. This can arise when the segments take different paths through the network, and some of those paths have MCA sender conversion enabled. The call succeeds with a completion code of MQCC_WARNING, but only the first few segments that have identical character-set identifiers are returned.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING

Programmer response

Remove the MQGMO_COMPLETE_MSG option from the MQGET call and retrieve the remaining message segments one by one.

2244 (08C4) (RC2244): MQRC_INCONSISTENT_ENCODINGS:

Explanation

An MQGET call was issued specifying the MQGMO_COMPLETE_MSG option, but the message to be retrieved consists of two or more segments that have differing values for the *Encoding* field in MQMD. This can arise when the segments take different paths through the network, and some of those paths have MCA sender conversion enabled. The call succeeds with a completion code of MQCC_WARNING, but only the first few segments that have identical encodings are returned.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING

Programmer response

Remove the MQGMO_COMPLETE_MSG option from the MQGET call and retrieve the remaining message segments one by one.

2245 (08C5) (RC2245): MQRC_INCONSISTENT_UOW:

Explanation

One of the following applies:

- An MQPUT call was issued to put a message in a group or a segment of a logical message, but the value specified or defaulted for the MQPMO_SYNCPOINT option is not consistent with the current group and segment information retained by the queue manager for the queue handle.
If the current call specifies MQPMO_LOGICAL_ORDER, the call fails. If the current call does not specify MQPMO_LOGICAL_ORDER, but the previous MQPUT call for the queue handle did, the call succeeds with completion code MQCC_WARNING.
- An MQGET call was issued to remove from the queue a message in a group or a segment of a logical message, but the value specified or defaulted for the MQGMO_SYNCPOINT option is not consistent with the current group and segment information retained by the queue manager for the queue handle.
If the current call specifies MQGMO_LOGICAL_ORDER, the call fails. If the current call does not specify MQGMO_LOGICAL_ORDER, but the previous MQGET call for the queue handle did, the call succeeds with completion code MQCC_WARNING.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

Modify the application to ensure that the same unit-of-work specification is used for all messages in the group, or all segments of the logical message.

2246 (08C6) (RC2246): MQRC_INVALID_MSG_UNDER_CURSOR:

Explanation

An MQGET call was issued specifying the MQGMO_COMPLETE_MSG option with either MQGMO_MSG_UNDER_CURSOR or MQGMO_BROWSE_MSG_UNDER_CURSOR, but the message that is under the cursor has an MQMD with an *Offset* field that is greater than zero. Because MQGMO_COMPLETE_MSG was specified, the message is not valid for retrieval.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Reposition the browse cursor so that it is located on a message with an *Offset* field in MQMD that is zero. Alternatively, remove the MQGMO_COMPLETE_MSG option.

2247 (08C7) (RC2247): MQRC_MATCH_OPTIONS_ERROR:

Explanation

An MQGET call was issued, but the value of the *MatchOptions* field in the *GetMsgOpts* parameter is not valid, for one of the following reasons:

- An undefined option is specified.
- All of the following are true:
 - MQGMO_LOGICAL_ORDER is specified.
 - There is a current message group or logical message for the queue handle.
 - Neither MQGMO_BROWSE_MSG_UNDER_CURSOR nor MQGMO_MSG_UNDER_CURSOR is specified.
 - One or more of the MQMO_* options is specified.
 - The values of the fields in the *MsgDesc* parameter corresponding to the MQMO_* options specified, differ from the values of those fields in the MQMD for the message to be returned next.
- On z/OS, one or more of the options specified is not valid for the index type of the queue.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that only valid options are specified for the field.

2248 (08C8) (RC2248): MQRC_MDE_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQMDE structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQMDE_STRUC_ID.
- The *Version* field is not MQMDE_VERSION_2.
- The *StrucLength* field is not MQMDE_LENGTH_2.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2249 (08C9) (RC2249): MQRC_MSG_FLAGS_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the *MsgFlags* field in the message descriptor MQMD contains one or more message flags that are not recognized by the local queue manager. The message flags that cause this reason code to be returned depend on the destination of the message; see the description of REPORT in  Report options and message flags (*WebSphere MQ V7.1 Reference*) for more information.

This reason code can also occur in the *Feedback* field in the MQMD of a report message, or in the *Reason* field in the MQDLH structure of a message on the dead-letter queue; in both cases it indicates that the destination queue manager does not support one or more of the message flags specified by the sender of the message.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Do the following:

- Ensure that the *MsgFlags* field in the message descriptor is initialized with a value when the message descriptor is declared, or is assigned a value prior to the MQPUT or MQPUT1 call. Specify MQMF_NONE if no message flags are needed.

- Ensure that the message flags specified are valid; see the *MsgFlags* field described in the description of MQMD in  *MsgFlags (MQLONG)* (*WebSphere MQ V7.1 Reference*) for valid message flags.
- If multiple message flags are being set by adding the individual message flags together, ensure that the same message flag is not added twice.
- On z/OS, ensure that the message flags specified are valid for the index type of the queue; see the description of the *MsgFlags* field in MQMD for further details.

2250 (08CA) (RC2250): MQRC_MSG_SEQ_NUMBER_ERROR:

Explanation

An MQGET, MQPUT, or MQPUT1 call was issued, but the value of the *MsgSeqNumber* field in the MQMD or MQMDE structure is less than one or greater than 999 999 999.

This error can also occur on the MQPUT call if the *MsgSeqNumber* field would have become greater than 999 999 999 as a result of the call.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify a value in the range 1 through 999 999 999. Do not attempt to create a message group containing more than 999 999 999 messages.

2251 (08CB) (RC2251): MQRC_OFFSET_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the value of the *Offset* field in the MQMD or MQMDE structure is less than zero or greater than 999 999 999.

This error can also occur on the MQPUT call if the *Offset* field would have become greater than 999 999 999 as a result of the call.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify a value in the range 0 through 999 999 999. Do not attempt to create a message segment that would extend beyond an offset of 999 999 999.

2252 (08CC) (RC2252): MQRC_ORIGINAL_LENGTH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued to put a report message that is a segment, but the *OriginalLength* field in the MQMD or MQMDE structure is either:

- Less than the length of data in the message, or
- Less than one (for a segment that is not the last segment), or
- Less than zero (for a segment that is the last segment)

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify a value that is greater than zero. Zero is valid only for the last segment.

2253 (08CD) (RC2253): MQRC_SEGMENT_LENGTH_ZERO:

Explanation

An MQPUT or MQPUT1 call was issued to put the first or an intermediate segment of a logical message, but the length of the application message data in the segment (excluding any MQ headers that may be present) is zero. The length must be at least one for the first or intermediate segment.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check the application logic to ensure that segments are put with a length of one or greater. Only the last segment of a logical message is permitted to have a length of zero.

2255 (08CF) (RC2255): MQRC_UOW_NOT_AVAILABLE:

Explanation

An MQGET, MQPUT, or MQPUT1 call was issued to get or put a message outside a unit of work, but the options specified on the call required the queue manager to process the call within a unit of work. Because there is already a user-defined unit of work in existence, the queue manager was unable to create a temporary unit of work for the duration of the call.

This reason occurs in the following circumstances:

- On an MQGET call, when the MQGMO_COMPLETE_MSG option is specified in MQGMO and the logical message to be retrieved is persistent and consists of two or more segments.
- On an MQPUT or MQPUT1 call, when the MQMF_SEGMENTATION_ALLOWED flag is specified in MQMD and the message requires segmentation.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Issue the MQGET, MQPUT, or MQPUT1 call inside the user-defined unit of work. Alternatively, for the MQPUT or MQPUT1 call, reduce the size of the message so that it does not require segmentation by the queue manager.

2256 (08D0) (RC2256): MQRC_WRONG_GMO_VERSION:

Explanation

An MQGET call was issued specifying options that required an MQGMO with a version number not less than MQGMO_VERSION_2, but the MQGMO supplied did not satisfy this condition.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Modify the application to pass a version-2 MQGMO. Check the application logic to ensure that the *Version* field in MQGMO has been set to MQGMO_VERSION_2. Alternatively, remove the option that requires the version-2 MQGMO.

2257 (08D1) (RC2257): MQRC_WRONG_MD_VERSION:

Explanation

An MQGET, MQPUT, or MQPUT1 call was issued specifying options that required an MQMD with a version number not less than MQMD_VERSION_2, but the MQMD supplied did not satisfy this condition.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Modify the application to pass a version-2 MQMD. Check the application logic to ensure that the *Version* field in MQMD has been set to MQMD_VERSION_2. Alternatively, remove the option that requires the version-2 MQMD.

2258 (08D2) (RC2258): MQRC_GROUP_ID_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued to put a distribution-list message that is also a message in a group, a message segment, or has segmentation allowed, but an invalid combination of options and values was specified. All of the following are true:

- MQPMO_LOGICAL_ORDER is not specified in the *Options* field in MQPMO.
- Either there are too few MQPMR records provided by MQPMO, or the *GroupId* field is not present in the MQPMR records.
- One or more of the following flags is specified in the *MsgFlags* field in MQMD or MQMDE:
 - MQMF_SEGMENTATION_ALLOWED
 - MQMF_*_MSG_IN_GROUP
 - MQMF_*_SEGMENT
- The *GroupId* field in MQMD or MQMDE is not MQGI_NONE.

This combination of options and values would result in the same group identifier being used for all of the destinations in the distribution list; this is not permitted by the queue manager.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify MQGI_NONE for the *GroupId* field in MQMD or MQMDE. Alternatively, if the call is MQPUT specify MQPMO_LOGICAL_ORDER in the *Options* field in MQPMO.

2259 (08D3) (RC2259): MQRC_INCONSISTENT_BROWSE:

Explanation

An MQGET call was issued with the MQGMO_BROWSE_NEXT option specified, but the specification of the MQGMO_LOGICAL_ORDER option for the call is different from the specification of that option for the previous call for the queue handle. Either both calls must specify MQGMO_LOGICAL_ORDER, or neither call must specify MQGMO_LOGICAL_ORDER.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Add or remove the MQGMO_LOGICAL_ORDER option as appropriate. Alternatively, to switch between logical order and physical order, specify the MQGMO_BROWSE_FIRST option to restart the scan from the beginning of the queue, omitting or specifying MQGMO_LOGICAL_ORDER as required.

2260 (08D4) (RC2260): MQRC_XQH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQXQH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQXQH_STRUC_ID.
- The *Version* field is not MQXQH_VERSION_1.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2261 (08D5) (RC2261): MQRC_SRC_ENV_ERROR:

Explanation

This reason occurs when a channel exit that processes reference messages detects an error in the source environment data of a reference message header (MQRMH). One of the following is true:

- *SrcEnvLength* is less than zero.
- *SrcEnvLength* is greater than zero, but there is no source environment data.
- *SrcEnvLength* is greater than zero, but *SrcEnvOffset* is negative, zero, or less than the length of the fixed part of MQRMH.
- *SrcEnvLength* is greater than zero, but *SrcEnvOffset* plus *SrcEnvLength* is greater than *StrucLength*.

The exit returns this reason in the *Feedback* field of the MQCXP structure. If an exception report is requested, it is copied to the *Feedback* field of the MQMD associated with the report.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify the source environment data correctly.

2262 (08D6) (RC2262): MQRC_SRC_NAME_ERROR:

Explanation

This reason occurs when a channel exit that processes reference messages detects an error in the source name data of a reference message header (MQRMH). One of the following is true:

- *SrcNameLength* is less than zero.
- *SrcNameLength* is greater than zero, but there is no source name data.
- *SrcNameLength* is greater than zero, but *SrcNameOffset* is negative, zero, or less than the length of the fixed part of MQRMH.
- *SrcNameLength* is greater than zero, but *SrcNameOffset* plus *SrcNameLength* is greater than *StrucLength*.

The exit returns this reason in the *Feedback* field of the MQCXP structure. If an exception report is requested, it is copied to the *Feedback* field of the MQMD associated with the report.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify the source name data correctly.

2263 (08D7) (RC2263): MQRC_DEST_ENV_ERROR:

Explanation

This reason occurs when a channel exit that processes reference messages detects an error in the destination environment data of a reference message header (MQRMH). One of the following is true:

- *DestEnvLength* is less than zero.
- *DestEnvLength* is greater than zero, but there is no destination environment data.
- *DestEnvLength* is greater than zero, but *DestEnvOffset* is negative, zero, or less than the length of the fixed part of MQRMH.
- *DestEnvLength* is greater than zero, but *DestEnvOffset* plus *DestEnvLength* is greater than *StrucLength*.

The exit returns this reason in the *Feedback* field of the MQCXP structure. If an exception report is requested, it is copied to the *Feedback* field of the MQMD associated with the report.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify the destination environment data correctly.

2264 (08D8) (RC2264): MQRC_DEST_NAME_ERROR:

Explanation

This reason occurs when a channel exit that processes reference messages detects an error in the destination name data of a reference message header (MQRMH). One of the following is true:

- *DestNameLength* is less than zero.
- *DestNameLength* is greater than zero, but there is no destination name data.
- *DestNameLength* is greater than zero, but *DestNameOffset* is negative, zero, or less than the length of the fixed part of MQRMH.
- *DestNameLength* is greater than zero, but *DestNameOffset* plus *DestNameLength* is greater than *StrucLength*.

The exit returns this reason in the *Feedback* field of the MQCXP structure. If an exception report is requested, it is copied to the *Feedback* field of the MQMD associated with the report.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Specify the destination name data correctly.

2265 (08D9) (RC2265): MQRC_TM_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQTM structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQTM_STRUC_ID.
- The *Version* field is not MQTM_VERSION_1.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2266 (08DA) (RC2266): MQRC_CLUSTER_EXIT_ERROR:

Explanation

An MQOPEN, MQPUT, or MQPUT1 call was issued to open or put a message on a cluster queue, but the cluster workload exit defined by the queue-manager's *ClusterWorkloadExit* attribute failed unexpectedly or did not respond in time. Subsequent MQOPEN, MQPUT, and MQPUT1 calls for this queue handle are processed as though the *ClusterWorkloadExit* attribute were blank.

- On z/OS, a message giving more information about the error is written to the system log, for example message CSQV455E or CSQV456E.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check the cluster workload exit to ensure that it has been written correctly.

2267 (08DB) (RC2267): MQRC_CLUSTER_EXIT_LOAD_ERROR:

Explanation

An MQCONN or MQCONNX call was issued to connect to a queue manager, but the queue manager was unable to load the cluster workload exit. Execution continues without the cluster workload exit.

- On z/OS, if the cluster workload exit cannot be loaded, a message is written to the system log, for example message CSQV453I. Processing continues as though the *ClusterWorkloadExit* attribute had been blank.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_WARNING

Programmer response

Ensure that the queue-manager's *ClusterWorkloadExit* attribute has the correct value, and that the exit has been installed into the correct location.

2268 (08DC) (RC2268): MQRC_CLUSTER_PUT_INHIBITED:

Explanation

An MQOPEN call with the MQOO_OUTPUT and MQOO_BIND_ON_OPEN options in effect was issued for a cluster queue, but the call failed because all of the following are true:

- All instances of the cluster queue are currently put-inhibited (that is, all of the queue instances have the *InhibitPut* attribute set to MQQA_PUT_INHIBITED).
- There is no local instance of the queue. (If there is a local instance, the MQOPEN call succeeds, even if the local instance is put-inhibited.)

- There is no cluster workload exit for the queue, or there is a cluster workload exit but it did not choose a queue instance. (If the cluster workload exit does choose a queue instance, the MQOPEN call succeeds, even if that instance is put-inhibited.)

If the MQOO_BIND_NOT_FIXED option is specified on the MQOPEN call, the call can succeed even if all of the queues in the cluster are put-inhibited. However, a subsequent MQPUT call may fail if all of the queues are still put-inhibited at the time of the MQPUT call.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

If the system design allows put requests to be inhibited for short periods, retry the operation later. If the problem persists, determine why all of the queues in the cluster are put-inhibited.

2269 (08DD) (RC2269): MQRC_CLUSTER_RESOURCE_ERROR:

Explanation

An MQOPEN, MQPUT, or MQPUT1 call was issued for a cluster queue, but an error occurred whilst trying to use a resource required for clustering.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Do the following:

- Check that the SYSTEM.CLUSTER.* queues are not put inhibited or full.
- Check the event queues for any events relating to the SYSTEM.CLUSTER.* queues, as these may give guidance as to the nature of the failure.
- Check that the repository queue manager is available.
- On z/OS, check the console for signs of the failure, such as full page sets.

2270 (08DE) (RC2270): MQRC_NO_DESTINATIONS_AVAILABLE:

Explanation

An MQPUT or MQPUT1 call was issued to put a message on a cluster queue, but at the time of the call there were no longer any instances of the queue in the cluster. The message therefore could not be sent.

This situation can occur when MQOO_BIND_NOT_FIXED is specified on the MQOPEN call that opens the queue, or MQPUT1 is used to put the message.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check the queue definition and queue status to determine why all instances of the queue were removed from the cluster. Correct the problem and rerun the application.

2271 (08DF) (RC2271): MQRC_CONN_TAG_IN_USE:

Explanation

An MQCONNX call was issued specifying one of the MQCNO_*_CONN_TAG_* options, but the call failed because the connection tag specified by *ConnTag* in MQCNO is in use by an active process or thread, or there is an unresolved unit of work that references this connection tag.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The problem is likely to be transitory. The application should wait a short while and then retry the operation.

2272 (08E0) (RC2272): MQRC_PARTIALLY_CONVERTED:

Explanation

On an MQGET call with the MQGMO_CONVERT option included in the *GetMsgOpts* parameter, one or more MQ header structures in the message data could not be converted to the specified target character set or encoding. In this situation, the MQ header structures are converted to the queue-manager's character set and encoding, and the application data in the message is converted to the target character set and encoding. On return from the call, the values returned in the various *CodedCharSetId* and *Encoding* fields in the *MsgDesc* parameter and MQ header structures indicate the character set and encoding that apply to each part of the message. The call completes with MQCC_WARNING.

This reason code usually occurs when the specified target character set is one that causes the character strings in the MQ header structures to expand beyond the lengths of their fields. Unicode character set UCS-2 is an example of a character set that causes this to happen.

Completion Code

MQCC_FAILED

Programmer response

If this is an expected situation, no corrective action is required.

If this is an unexpected situation, check that the MQ header structures contain valid data. If they do, specify as the target character set a character set that does not cause the strings to expand.

2273 (08E1) (RC2273): MQRC_CONNECTION_ERROR:

Explanation

An MQCONN or MQCONNX call failed for one of the following reasons:

- The installation and customization options chosen for WebSphere MQ do not allow connection by the type of application being used.
- The system parameter module is not at the same release level as the queue manager.
- The channel initiator is not at the same release level as the queue manager.
- An internal error was detected by the queue manager.

Completion Code

MQCC_FAILED

Programmer response

None, if the installation and customization options chosen for WebSphere MQ do not allow all functions to be used.

Otherwise, if this occurs while starting the channel initiator, ensure that the queue manager and the channel initiator are both at the same release level and that their started task JCL procedures both specify the same level of WebSphere MQ program libraries; if this occurs while starting the queue manager, relinkedit the system parameter module (CSQZPARM) to ensure that it is at the correct level. If the problem persists, contact your IBM support center.

2274 (08E2) (RC2274): MQRC_OPTION_ENVIRONMENT_ERROR:

Explanation

An MQGET call with the MQGMO_MARK_SKIP_BACKOUT option specified was issued from a Db2 Stored Procedure. The call failed because the MQGMO_MARK_SKIP_BACKOUT option cannot be used from a Db2 Stored Procedure.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Remove the MQGMO_MARK_SKIP_BACKOUT option from the MQGET call.

2277 (08E5) (RC2277): MQRC_CD_ERROR:

Explanation

An MQCONNX call was issued to connect to a queue manager, but the MQCD channel definition structure addressed by the *ClientConnOffset* or *ClientConnPtr* field in MQCNO contains data that is not valid. Consult the error log for more information about the nature of the error.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Ensure that input fields in the MQCD structure are set correctly.

2278 (08E6) (RC2278): MQRC_CLIENT_CONN_ERROR:

Explanation

An MQCONNX call was issued to connect to a queue manager, but the MQCD channel definition structure is not specified correctly. One of the following applies:

- *ClientConnOffset* is not zero and *ClientConnPtr* is not zero and not the null pointer.
- *ClientConnPtr* is not a valid pointer.
- *ClientConnPtr* or *ClientConnOffset* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems. It also occurs in Java applications when a client channel definition table (CCDT) is specified to determine the name of the channel, but the table itself cannot be found.

Completion Code

MQCC_FAILED

Programmer response

Ensure that at least one of *ClientConnOffset* and *ClientConnPtr* is zero. Ensure that the field used points to accessible storage. Ensure that the URL of the client channel definition table is correct.

2279 (08E7) (RC2279): MQRC_CHANNEL_STOPPED_BY_USER:

Explanation

This condition is detected when the channel has been stopped by an operator. The reason qualifier identifies the reasons for stopping.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2280 (08E8) (RC2280): MQRC_HCONFIG_ERROR:

Explanation

The configuration handle *Hconfig* specified on the MQXEP call or MQZEP call is not valid. The MQXEP call is issued by an API exit function; the MQZEP call is issued by an installable service.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Specify the configuration handle that was provided by the queue manager:

- On the MQXEP call, use the handle passed in the *Hconfig* field of the MQAXP structure.
- On the MQZEP call, use the handle passed to the installable service's configuration function on the component initialization call. For more information about installable services, see  [Installable services and components for UNIX, Linux and Windows \(WebSphere MQ V7.1 Programming Guide\)](#).

2281 (08E9) (RC2281): MQRC_FUNCTION_ERROR:

Explanation

An MQXEP or MQZEP call was issued, but the function identifier *Function* specified on the call is not valid, or not supported by the installable service being configured.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Do the following:

- For the MQXEP call, specify one of the MQXF_* values.
- For the MQZEP call, specify an MQZID_* value that is valid for the installable service being configured. See  [MQZEP – Add component entry point \(WebSphere MQ V7.1 Reference\)](#) to determine which values are valid.

2282 (08EA) (RC2282): MQRC_CHANNEL_STARTED:

Explanation

One of the following has occurred:

- An operator has issued a Start Channel command.
- An instance of a channel has been successfully established. This condition is detected when Initial Data negotiation is complete and resynchronization has been performed where necessary such that message transfer can proceed.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2283 (08EB) (RC2283): MQRC_CHANNEL_STOPPED:**Explanation**

This condition is detected when the channel has been stopped. The reason qualifier identifies the reasons for stopping.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2284 (08EC) (RC2284): MQRC_CHANNEL_CONV_ERROR:**Explanation**

This condition is detected when a channel is unable to do data conversion and the MQGET call to get a message from the transmission queue resulted in a data conversion error. The conversion reason code identifies the reason for the failure.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2285 (08ED) (RC2285): MQRC_SERVICE_NOT_AVAILABLE:**Explanation**

This reason should be returned by an installable service component when the requested action cannot be performed because the required underlying service is not available.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Make the underlying service available.

2286 (08EE) (RC2286): MQRC_INITIALIZATION_FAILED:

Explanation

This reason should be returned by an installable service component when the component is unable to complete initialization successfully.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Correct the error and retry the operation.

2287 (08EF) (RC2287): MQRC_TERMINATION_FAILED:

Explanation

This reason should be returned by an installable service component when the component is unable to complete termination successfully.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Correct the error and retry the operation.

2288 (08F0) (RC2288): MQRC_UNKNOWN_Q_NAME:

Explanation

This reason should be returned by the MQZ_LOOKUP_NAME installable service component when the name specified for the *QName* parameter is not recognized.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

None. See  Installable services and components for UNIX, Linux and Windows (*WebSphere MQ V7.1 Programming Guide*) for more information about installable services.

2289 (08F1) (RC2289): MQRC_SERVICE_ERROR:

Explanation

This reason should be returned by an installable service component when the component encounters an unexpected error.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Correct the error and retry the operation.

2290 (08F2) (RC2290): MQRC_Q_ALREADY_EXISTS:

Explanation

This reason should be returned by the MQZ_INSERT_NAME installable service component when the queue specified by the *QName* parameter is already defined to the name service.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

None. See  Installable services and components for UNIX, Linux and Windows (*WebSphere MQ V7.1 Programming Guide*) for more information about installable services.

2291 (08F3) (RC2291): MQRC_USER_ID_NOT_AVAILABLE:

Explanation

This reason should be returned by the MQZ_FIND_USERID installable service component when the user ID cannot be determined.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

None. See  Installable services and components for UNIX, Linux and Windows (*WebSphere MQ V7.1 Programming Guide*) for more information about installable services.

2292 (08F4) (RC2292): MQRC_UNKNOWN_ENTITY:

Explanation

This reason should be returned by the authority installable service component when the name specified by the *EntityName* parameter is not recognized.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the entity is defined.

2294 (08F6) (RC2294): MQRC_UNKNOWN_REF_OBJECT:

Explanation

This reason should be returned by the MQZ_COPY_ALL_AUTHORITY installable service component when the name specified by the *RefObjectName* parameter is not recognized.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the reference object is defined. See  Installable services and components for UNIX, Linux and Windows (*WebSphere MQ V7.1 Programming Guide*) for more information about installable services.

2295 (08F7) (RC2295): MQRC_CHANNEL_ACTIVATED:

Explanation

This condition is detected when a channel that has been waiting to become active, and for which a Channel Not Activated event has been generated, is now able to become active because an active slot has been released by another channel.

This event is not generated for a channel that is able to become active without waiting for an active slot to be released.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2296 (08F8) (RC2296): MQRC_CHANNEL_NOT_ACTIVATED:

Explanation

This condition is detected when a channel is required to become active, either because it is starting or because it is about to make another attempt to establish connection with its partner. However, it is unable to do so because the limit on the number of active channels has been reached.

- On z/OS, the maximum number of active channels is given by the ACTCHL queue manager attribute.
- In other environments, the maximum number of active channels is given by the MaxActiveChannels parameter in the qm.ini file.

The channel waits until it is able to take over an active slot released when another channel ceases to be active. At that time a Channel Activated event is generated.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2297 (08F9) (RC2297): MQRC_UOW_CANCELED:

Explanation

An MQI call was issued, but the unit of work (TM/MP transaction) being used for the MQ operation had been canceled. This may have been done by TM/MP itself (for example, due to the transaction running for too long, or exceeding audit trail sizes), or by the application program issuing an ABORT_TRANSACTION. All updates performed to resources owned by the queue manager are backed out.

Completion Code

MQCC_FAILED

Programmer response

Refer to the operating system's *Transaction Management Operations Guide* to determine how the Transaction Manager can be tuned to avoid the problem of system limits being exceeded.

2298 (08FA) (RC2298): MQRC_FUNCTION_NOT_SUPPORTED:

Explanation

The function requested is not available in the current environment.

Completion Code

MQCC_FAILED

Programmer response

Remove the call from the application.

This reason code can be used when the call requires resources or functionality that is restricted by the queue manager OPMODE setting.

If you get this reason code with CICS group connect, check that the queue manager attribute GROUPPUR is enabled.

2299 (08FB) (RC2299): MQRC_SELECTOR_TYPE_ERROR:

Explanation

The *Selector* parameter has the wrong data type; it must be of type Long.

Completion Code

MQCC_FAILED

Programmer response

Declare the *Selector* parameter as Long.

2300 (08FC) (RC2300): MQRC_COMMAND_TYPE_ERROR:

Explanation

The mqExecute call was issued, but the value of the MQIASY_TYPE data item in the administration bag is not MQCFT_COMMAND.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the MQIASY_TYPE data item in the administration bag has the value MQCFT_COMMAND.

2301 (08FD) (RC2301): MQRC_MULTIPLE_INSTANCE_ERROR:

Explanation

The *Selector* parameter specifies a system selector (one of the MQIASY_* values), but the value of the *ItemIndex* parameter is not MQIND_NONE. Only one instance of each system selector can exist in the bag.

Completion Code

MQCC_FAILED

Programmer response

Specify MQIND_NONE for the *ItemIndex* parameter.

2302 (08FE) (RC2302): MQRC_SYSTEM_ITEM_NOT_ALTERABLE:

Explanation

A call was issued to modify the value of a system data item in a bag (a data item with one of the MQIASY_* selectors), but the call failed because the data item is one that cannot be altered by the application.

Completion Code

MQCC_FAILED

Programmer response

Specify the selector of a user-defined data item, or remove the call.

2303 (08FF) (RC2303): MQRC_BAG_CONVERSION_ERROR:

Explanation

The mqBufferToBag or mqGetBag call was issued, but the data in the buffer or message could not be converted into a bag. This occurs when the data to be converted is not valid PCF.

Completion Code

MQCC_FAILED

Programmer response

Check the logic of the application that created the buffer or message to ensure that the buffer or message contains valid PCF.

If the message contains PCF that is not valid, the message cannot be retrieved using the mqGetBag call:

- If one of the MQGMO_BROWSE_* options was specified, the message remains on the queue and can be retrieved using the MQGET call.
- In other cases, the message has already been removed from the queue and discarded. If the message was retrieved within a unit of work, the unit of work can be backed out and the message retrieved using the MQGET call.

2304 (0900) (RC2304): MQRC_SELECTOR_OUT_OF_RANGE:

Explanation

The *Selector* parameter has a value that is outside the valid range for the call. If the bag was created with the MQCBO_CHECK_SELECTORS option:

- For the mqAddInteger call, the value must be within the range MQIA_FIRST through MQIA_LAST.
- For the mqAddString call, the value must be within the range MQCA_FIRST through MQCA_LAST.

If the bag was not created with the MQCBO_CHECK_SELECTORS option:

- The value must be zero or greater.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value.

2305 (0901) (RC2305): MQRC_SELECTOR_NOT_UNIQUE:

Explanation

The *ItemIndex* parameter has the value MQIND_NONE, but the bag contains more than one data item with the selector value specified by the *Selector* parameter. MQIND_NONE requires that the bag contain only one occurrence of the specified selector.

This reason code also occurs on the mqExecute call when the administration bag contains two or more occurrences of a selector for a required parameter that permits only one occurrence.

Completion Code

MQCC_FAILED

Programmer response

Check the logic of the application that created the bag. If correct, specify for *ItemIndex* a value that is zero or greater, and add application logic to process all of the occurrences of the selector in the bag.

Review the description of the administration command being issued, and ensure that all required parameters are defined correctly in the bag.

2306 (0902) (RC2306): MQRC_INDEX_NOT_PRESENT:

Explanation

The specified index is not present:

- For a bag, this means that the bag contains one or more data items that have the selector value specified by the *Selector* parameter, but none of them has the index value specified by the *ItemIndex* parameter. The data item identified by the *Selector* and *ItemIndex* parameters must exist in the bag.
- For a namelist, this means that the index parameter value is too large, and outside the range of valid values.

Completion Code

MQCC_FAILED

Programmer response

Specify the index of a data item that does exist in the bag or namelist. Use the mqCountItems call to determine the number of data items with the specified selector that exist in the bag, or the nameCount method to determine the number of names in the namelist.

2307 (0903) (RC2307): MQRC_STRING_ERROR:

Explanation

The *String* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2308 (0904) (RC2308): MQRC_ENCODING_NOT_SUPPORTED:

Explanation

The *Encoding* field in the message descriptor MQMD contains a value that is not supported:

- For the mqPutBag call, the field in error resides in the *MsgDesc* parameter of the call.
- For the mqGetBag call, the field in error resides in:
 - The *MsgDesc* parameter of the call if the MQGMO_CONVERT option was specified.
 - The message descriptor of the message about to be retrieved if MQGMO_CONVERT was *not* specified.

Completion Code

MQCC_FAILED

Programmer response

The value must be MQENC_NATIVE.

If the value of the *Encoding* field in the message is not valid, the message cannot be retrieved using the mqGetBag call:

- If one of the MQGMO_BROWSE_* options was specified, the message remains on the queue and can be retrieved using the MQGET call.
- In other cases, the message has already been removed from the queue and discarded. If the message was retrieved within a unit of work, the unit of work can be backed out and the message retrieved using the MQGET call.

2309 (0905) (RC2309): MQRC_SELECTOR_NOT_PRESENT:

Explanation

The *Selector* parameter specifies a selector that does not exist in the bag.

Completion Code

MQCC_FAILED

Programmer response

Specify a selector that does exist in the bag.

2310 (0906) (RC2310): MQRC_OUT_SELECTOR_ERROR:

Explanation

The *OutSelector* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2311 (0907) (RC2311): MQRC_STRING_TRUNCATED:

Explanation

The string returned by the call is too long to fit in the buffer provided. The string has been truncated to fit in the buffer.

Completion Code

MQCC_FAILED

Programmer response

If the entire string is required, provide a larger buffer. On the *mqInquireString* call, the *StringLength* parameter is set by the call to indicate the size of the buffer required to accommodate the string without truncation.

2312 (0908) (RC2312): MQRC_SELECTOR_WRONG_TYPE:

Explanation

A data item with the specified selector exists in the bag, but has a data type that conflicts with the data type implied by the call being used. For example, the data item might have an integer data type, but the call being used might be *mqSetString*, which implies a character data type.

This reason code also occurs on the *mqBagToBuffer*, *mqExecute*, and *mqPutBag* calls when *mqAddString* or *mqSetString* was used to add the MQIACF_INQUIRY data item to the bag.

Completion Code

MQCC_FAILED

Programmer response

For the `mqSetInteger` and `mqSetString` calls, specify `MQIND_ALL` for the *ItemIndex* parameter to delete from the bag all existing occurrences of the specified selector before creating the new occurrence with the required data type.

For the `mqInquireBag`, `mqInquireInteger`, and `mqInquireString` calls, use the `mqInquireItemInfo` call to determine the data type of the item with the specified selector, and then use the appropriate call to determine the value of the data item.

For the `mqBagToBuffer`, `mqExecute`, and `mqPutBag` calls, ensure that the `MQIACF_INQUIRY` data item is added to the bag using the `mqAddInteger` or `mqSetInteger` calls.

2313 (0909) (RC2313): MQRC_INCONSISTENT_ITEM_TYPE:

Explanation

The `mqAddInteger` or `mqAddString` call was issued to add another occurrence of the specified selector to the bag, but the data type of this occurrence differed from the data type of the first occurrence.

This reason can also occur on the `mqBufferToBag` and `mqGetBag` calls, where it indicates that the PCF in the buffer or message contains a selector that occurs more than once but with inconsistent data types.

Completion Code

MQCC_FAILED

Programmer response

For the `mqAddInteger` and `mqAddString` calls, use the call appropriate to the data type of the first occurrence of that selector in the bag.

For the `mqBufferToBag` and `mqGetBag` calls, check the logic of the application that created the buffer or sent the message to ensure that multiple-occurrence selectors occur with only one data type. A message that contains a mixture of data types for a selector cannot be retrieved using the `mqGetBag` call:

- If one of the `MQGMO_BROWSE_*` options was specified, the message remains on the queue and can be retrieved using the `MQGET` call.
- In other cases, the message has already been removed from the queue and discarded. If the message was retrieved within a unit of work, the unit of work can be backed out and the message retrieved using the `MQGET` call.

2314 (090A) (RC2314): MQRC_INDEX_ERROR:

Explanation

An index parameter to a call or method has a value that is not valid. The value must be zero or greater. For bag calls, certain `MQIND_*` values can also be specified:

- For the `mqDeleteItem`, `mqSetInteger` and `mqSetString` calls, `MQIND_ALL` and `MQIND_NONE` are valid.
- For the `mqInquireBag`, `mqInquireInteger`, `mqInquireString`, and `mqInquireItemInfo` calls, `MQIND_NONE` is valid.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value.

2315 (090B) (RC2315): MQRC_SYSTEM_BAG_NOT_ALTERABLE:

Explanation

A call was issued to add a data item to a bag, modify the value of an existing data item in a bag, or retrieve a message into a bag, but the call failed because the bag is one that had been created by the system as a result of a previous mqExecute call. System bags cannot be modified by the application.

Completion Code

MQCC_FAILED

Programmer response

Specify the handle of a bag created by the application, or remove the call.

2316 (090C) (RC2316): MQRC_ITEM_COUNT_ERROR:

Explanation

The mqTruncateBag call was issued, but the *ItemCount* parameter specifies a value that is not valid. The value is either less than zero, or greater than the number of user-defined data items in the bag.

This reason also occurs on the mqCountItems call if the parameter pointer is not valid, or points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value. Use the mqCountItems call to determine the number of user-defined data items in the bag.

2317 (090D) (RC2317): MQRC_FORMAT_NOT_SUPPORTED:

Explanation

The *Format* field in the message descriptor MQMD contains a value that is not supported:

- In an administration message, the format value must be one of the following: MQFMT_ADMIN, MQFMT_EVENT, MQFMT_PCF. For the mqPutBag call, the field in error resides in the *MsgDesc* parameter of the call. For the mqGetBag call, the field in error resides in the message descriptor of the message about to be retrieved.
- On z/OS, the message was put to the command input queue with a format value of MQFMT_ADMIN, but the version of MQ being used does not support that format for commands.

Completion Code

MQCC_FAILED

Programmer response

If the error occurred when putting a message, correct the format value.

If the error occurred when getting a message, the message cannot be retrieved using the mqGetBag call:

- If one of the MQGMO_BROWSE_* options was specified, the message remains on the queue and can be retrieved using the MQGET call.
- In other cases, the message has already been removed from the queue and discarded. If the message was retrieved within a unit of work, the unit of work can be backed out and the message retrieved using the MQGET call.

2318 (090E) (RC2318): MQRC_SELECTOR_NOT_SUPPORTED:

Explanation

The *Selector* parameter specifies a value that is a system selector (a value that is negative), but the system selector is not one that is supported by the call.

Completion Code

MQCC_FAILED

Programmer response

Specify a selector value that is supported.

2319 (090F) (RC2319): MQRC_ITEM_VALUE_ERROR:

Explanation

The mqInquireBag or mqInquireInteger call was issued, but the *ItemValue* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2320 (0910) (RC2320): MQRC_HBAG_ERROR:

Explanation

A call was issued that has a parameter that is a bag handle, but the handle is not valid. For output parameters, this reason also occurs if the parameter pointer is not valid, or points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2321 (0911) (RC2321): MQRC_PARAMETER_MISSING:

Explanation

An administration message requires a parameter that is not present in the administration bag. This reason code occurs only for bags created with the MQCBO_ADMIN_BAG or MQCBO_REORDER_AS_REQUIRED options.

Completion Code

MQCC_FAILED

Programmer response

Review the description of the administration command being issued, and ensure that all required parameters are present in the bag.

2322 (0912) (RC2322): MQRC_CMD_SERVER_NOT_AVAILABLE:

Explanation

The command server that processes administration commands is not available.

Completion Code

MQCC_FAILED

Programmer response

Start the command server.

2323 (0913) (RC2323): MQRC_STRING_LENGTH_ERROR:

Explanation

The *StringLength* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2324 (0914) (RC2324): MQRC_INQUIRY_COMMAND_ERROR:

Explanation

The mqAddInquiry call was used previously to add attribute selectors to the bag, but the command code to be used for the mqBagToBuffer, mqExecute, or mqPutBag call is not recognized. As a result, the correct PCF message cannot be generated.

Completion Code

MQCC_FAILED

Programmer response

Remove the mqAddInquiry calls and use instead the mqAddInteger call with the appropriate MQIACF_*_ATTRS or MQIACH_*_ATTRS selectors.

2325 (0915) (RC2325): MQRC_NESTED_BAG_NOT_SUPPORTED:

Explanation

A bag that is input to the call contains nested bags. Nested bags are supported only for bags that are output from the call.

Completion Code

MQCC_FAILED

Programmer response

Use a different bag as input to the call.

2326 (0916) (RC2326): MQRC_BAG_WRONG_TYPE:

Explanation

The *Bag* parameter specifies the handle of a bag that has the wrong type for the call. The bag must be an administration bag, that is, it must be created with the MQCBO_ADMIN_BAG option specified on the mqCreateBag call.

Completion Code

MQCC_FAILED

Programmer response

Specify the MQCBO_ADMIN_BAG option when the bag is created.

2327 (0917) (RC2327): MQRC_ITEM_TYPE_ERROR:**Explanation**

The `mqInquireItemInfo` call was issued, but the *ItemType* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2328 (0918) (RC2328): MQRC_SYSTEM_BAG_NOT_DELETABLE:**Explanation**

An `mqDeleteBag` call was issued to delete a bag, but the call failed because the bag is one that had been created by the system as a result of a previous `mqExecute` call. System bags cannot be deleted by the application.

Completion Code

MQCC_FAILED

Programmer response

Specify the handle of a bag created by the application, or remove the call.

2329 (0919) (RC2329): MQRC_SYSTEM_ITEM_NOT_DELETABLE:**Explanation**

A call was issued to delete a system data item from a bag (a data item with one of the `MQIASY_*` selectors), but the call failed because the data item is one that cannot be deleted by the application.

Completion Code

MQCC_FAILED

Programmer response

Specify the selector of a user-defined data item, or remove the call.

2330 (091A) (RC2330): MQRC_CODED_CHAR_SET_ID_ERROR:**Explanation**

The *CodedCharSetId* parameter is not valid. Either the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2331 (091B) (RC2331): MQRC_MSG_TOKEN_ERROR:

Explanation

An MQGET call was issued to retrieve a message using the message token as a selection criterion, but the options specified are not valid, because MQMO_MATCH_MSG_TOKEN was specified with either MQGMO_WAIT or MQGMO_SET_SIGNAL.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Remove the MQMO_MATCH_MSG_TOKEN option from the MQGET call.

2332 (091C) (RC2332): MQRC_MISSING_WIH:

Explanation

An MQPUT or MQPUT1 call was issued to put a message on a queue with an *IndexType* attribute that had the value MQIT_MSG_TOKEN, but the *Format* field in the MQMD was not MQFMT_WORK_INFO_HEADER. This error occurs only when the message arrives at the destination queue manager.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Modify the application to ensure that it places an MQWIH structure at the start of the message data, and sets the *Format* field in the MQMD to MQFMT_WORK_INFO_HEADER. Alternatively, change the *ApplType* attribute of the process definition used by the destination queue to be MQAT_WLM, and specify the required service name and service step name in its *EnvData* attribute.

2333 (091D) (RC2333): MQRC_WIH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQWIH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQWIH_STRUC_ID.
- The *Version* field is not MQWIH_VERSION_1.
- The *StrucLength* field is not MQWIH_LENGTH_1.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).
- On z/OS, this error also occurs when the *IndexType* attribute of the queue is MQIT_MSG_TOKEN, but the message data does not begin with an MQWIH structure.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

- On z/OS, if the queue has an *IndexType* of MQIT_MSG_TOKEN, ensure that the message data begins with an MQWIH structure.

2334 (091E) (RC2334): MQRC_RFH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQRFH or MQRFH2 structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQRFH_STRUC_ID.
- The *Version* field is not MQRFH_VERSION_1 (MQRFH), or MQRFH_VERSION_2 (MQRFH2).
- The *StrucLength* field specifies a value that is too small to include the structure plus the variable-length data at the end of the structure.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure (the structure extends beyond the end of the message).

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value (note: MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are *not* valid in this field).

2335 (091F) (RC2335): MQRC_RFH_STRING_ERROR:

Explanation

The contents of the *NameValueString* field in the MQRFH structure are not valid. *NameValueString* must adhere to the following rules:

- The string must consist of zero or more name/value pairs separated from each other by one or more blanks; the blanks are not significant.
- If a name or value contains blanks that are significant, the name or value must be enclosed in double quotation marks.
- If a name or value itself contains one or more double quotation marks, the name or value must be enclosed in double quotation marks, and each embedded double quotation mark must be doubled.
- A name or value can contain any characters other than the null, which acts as a delimiter. The null and characters following it, up to the defined length of *NameValueString*, are ignored.

The following is a valid *NameValueString*:

```
Famous_Words "The program displayed ""Hello World"""
```

Completion Code

MQCC_FAILED

Programmer response

Modify the application that generated the message to ensure that it places in the *NameValueString* field data that adheres to the rules. Check that the *StrucLength* field is set to the correct value.

2336 (0920) (RC2336): MQRC_RFH_COMMAND_ERROR:

Explanation

The message contains an MQRFH structure, but the command name contained in the *NameValueString* field is not valid.

Completion Code

MQCC_FAILED

Programmer response

Modify the application that generated the message to ensure that it places in the *NameValueString* field a command name that is valid.

2337 (0921) (RC2337): MQRC_RFH_PARM_ERROR:

Explanation

The message contains an MQRFH structure, but a parameter name contained in the *NameValueString* field is not valid for the command specified.

Completion Code

MQCC_FAILED

Programmer response

Modify the application that generated the message to ensure that it places in the *NameValueString* field only parameters that are valid for the specified command.

2338 (0922) (RC2338): MQRC_RFH_DUPLICATE_PARM:

Explanation

The message contains an MQRFH structure, but a parameter occurs more than once in the *NameValueString* field when only one occurrence is valid for the specified command.

Completion Code

MQCC_FAILED

Programmer response

Modify the application that generated the message to ensure that it places in the *NameValueString* field only one occurrence of the parameter.

2339 (0923) (RC2339): MQRC_RFH_PARM_MISSING:

Explanation

The message contains an MQRFH structure, but the command specified in the *NameValueString* field requires a parameter that is not present.

Completion Code

MQCC_FAILED

Programmer response

Modify the application that generated the message to ensure that it places in the *NameValueString* field all parameters that are required for the specified command.

2340 (0924) (RC2340): MQRC_CHAR_CONVERSION_ERROR:

Explanation

This reason code is returned by the Java MQQueueManager constructor when a required character-set conversion is not available. The conversion required is between two nonUnicode character sets.

This reason code occurs in the following environment: MQ Classes for Java on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the National Language Resources component of the z/OS Language Environment is installed, and that conversion between the IBM-1047 and ISO8859-1 character sets is available.

2341 (0925) (RC2341): MQRC_UCS2_CONVERSION_ERROR:

Explanation

This reason code is returned by the Java MQQueueManager constructor when a required character set conversion is not available. The conversion required is between the UCS-2 Unicode character set and the character set of the queue manager which defaults to IBM-500 if no specific value is available.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the relevant Unicode conversion tables are available for the JVM. For z/OS ensure that the Unicode conversion tables are available to the z/OS Language Environment. The conversion tables should be installed as part of the z/OS C/C++ optional feature. Refer to the *z/OS C/C++ Programming Guide* for more information about enabling UCS-2 conversions.

2342 (0926) (RC2342): MQRC_DB2_NOT_AVAILABLE:

Explanation

An MQOPEN, MQPUT1, or MQSET call, or a command, was issued to access a shared queue, but it failed because the queue manager is not connected to a Db2 subsystem. As a result, the queue manager is unable to access the object definition relating to the shared queue.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Configure the Db2 subsystem so that the queue manager can connect to it.

2343 (0927) (RC2343): MQRC_OBJECT_NOT_UNIQUE:

Explanation

An MQOPEN or MQPUT1 call, or a command, was issued to access a queue, but the call failed because the queue specified cannot be resolved unambiguously. There exists a shared queue with the specified name, and a nonshared queue with the same name.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

One of the queues must be deleted. If the queue to be deleted contains messages, use the MQSC command MOVE QLOCAL to move the messages to a different queue, and then use the command DELETE QLOCAL to delete the queue.

2344 (0928) (RC2344): MQRC_CONN_TAG_NOT_RELEASED:**Explanation**

An MQDISC call was issued when there was a unit of work outstanding for the connection handle. For CICS, IMS, and RRS connections, the MQDISC call does not commit or back out the unit of work. As a result, the connection tag associated with the unit of work is not yet available for reuse. The tag becomes available for reuse only when processing of the unit of work has been completed.

This reason code occurs only on z/OS.

Completion Code

MQCC_WARNING

Programmer response

Do not try to reuse the connection tag immediately. If the MQCONN call is issued with the same connection tag, and that tag is still in use, the call fails with reason code MQRC_CONN_TAG_IN_USE.

2345 (0929) (RC2345): MQRC_CF_NOT_AVAILABLE:**Explanation**

An MQOPEN or MQPUT1 call was issued to access a shared queue, but the allocation of the coupling-facility structure specified in the queue definition failed because there is no suitable coupling facility to hold the structure, based on the preference list in the active CFRM policy.

This reason code can also occur when the API call requires a capability that is not supported by the CF level defined in the coupling-facility structure object. For example, this reason code is returned by an attempt to open a shared queue that has a index type of MQIT_GROUP_ID, but the coupling-facility structure for the queue has a CF level lower than three.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Make available a coupling facility with one of the names specified in the CFRM policy, or modify the CFRM policy to specify the names of coupling facilities that are available.

2346 (092A) (RC2346): MQRC_CF_STRUC_IN_USE:**Explanation**

An MQI call or command was issued to operate on a shared queue, but the call failed because the coupling-facility structure specified in the queue definition is unavailable. The coupling-facility structure can be unavailable because a structure dump is in progress, or new connectors to the structure are currently inhibited, or an existing connector to the structure failed or disconnected abnormally and clean-up is not yet complete.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Typically, this is a temporary problem: wait for a while then retry the operation.

If the problem does not resolve itself, then connectivity problems experienced during the recovery of structures in the coupling facility could have occurred. In this case, restart the queue manager which reported the error. Resolve all the connectivity problems concerning the coupling facility before restarting the queue manager.

2347 (092B) (RC2347): MQRC_CF_STRUC_LIST_HDR_IN_USE:

Explanation

An MQGET, MQOPEN, MQPUT1, or MQSET call was issued to access a shared queue, but the call failed because the list header associated with the coupling-facility structure specified in the queue definition is temporarily unavailable. The list header is unavailable because it is undergoing recovery processing.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The problem is temporary; wait a short while and then retry the operation.

2348 (092C) (RC2348): MQRC_CF_STRUC_AUTH_FAILED:

Explanation

An MQOPEN or MQPUT1 call was issued to access a shared queue, but the call failed because the user is not authorized to access the coupling-facility structure specified in the queue definition.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Modify the security profile for the user identifier used by the application so that the application can access the coupling-facility structure specified in the queue definition.

2349 (092D) (RC2349): MQRC_CF_STRUC_ERROR:

Explanation

An MQOPEN or MQPUT1 call was issued to access a shared queue, but the call failed because the coupling-facility structure name specified in the queue definition is not defined in the CFRM data set, or is not the name of a list structure.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Modify the queue definition to specify the name of a coupling-facility list structure that is defined in the CFRM data set.

2350 (092E) (RC2350): MQRC_CONN_TAG_NOT_USABLE:

Explanation

An MQCONN call was issued specifying one of the MQCNO_*_CONN_TAG_* options, but the call failed because the connection tag specified by *ConnTag* in MQCNO is being used by the queue manager for recovery processing, and this processing is delayed pending recovery of the coupling facility.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The problem is likely to persist. Consult the system programmer to ascertain the cause of the problem.

2351 (092F) (RC2351): MQRC_GLOBAL_UOW_CONFLICT:

Explanation

An attempt was made to use inside a global unit of work a connection handle that is participating in another global unit of work. This can occur when an application passes connection handles between objects where the objects are involved in different DTC transactions. Because transaction completion is asynchronous, it is possible for this error to occur *after* the application has finalized the first object and committed its transaction.

This error does not occur for nontransactional MQI calls.

This reason code occurs only on Windows and z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check that the “MTS Transaction Support” attribute defined for the object’s class is set correctly. If necessary, modify the application so that the connection handle is not used by objects participating in different units of work.

2352 (0930) (RC2352): MQRC_LOCAL_UOW_CONFLICT:

Explanation

An attempt was made to use inside a global unit of work a connection handle that is participating in a queue-manager coordinated local unit of work. This can occur when an application passes connection handles between objects where one object is involved in a DTC transaction and the other is not.

This error does not occur for nontransactional MQI calls.

This reason code occurs only on Windows and z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check that the “MTS Transaction Support” attribute defined for the object’s class is set correctly. If necessary, modify the application so that the connection handle is not used by objects participating in different units of work.

2353 (0931) (RC2353): MQRC_HANDLE_IN_USE_FOR_UOW:

Explanation

An attempt was made to use outside a unit of work a connection handle that is participating in a global unit of work.

This error can occur when an application passes connection handles between objects where one object is involved in a DTC transaction and the other is not. Because transaction completion is asynchronous, it is possible for this error to occur *after* the application has finalized the first object and committed its transaction.

This error can also occur when a single object that was created and associated with the transaction loses that association whilst the object is running. The association is lost when DTC terminates the transaction independently of MTS. This might be because the transaction timed out, or because DTC shut down.

This error does not occur for nontransactional MQI calls.

This reason code occurs only on Windows.

Completion Code

MQCC_FAILED

Programmer response

Check that the “MTS Transaction Support” attribute defined for the object’s class is set correctly. If necessary, modify the application so that objects executing within different units of work do not try to

use the same connection handle.

2354 (0932) (RC2354): MQRC_UOW_ENLISTMENT_ERROR:

Explanation

This reason code can occur for various reasons and occurs only on Windows, and HP Integrity NonStop Server.

On Windows, the most likely reason is that an object created by a DTC transaction does not issue a transactional MQI call until after the DTC transaction timed out. (If the DTC transaction times out after a transactional MQI call has been issued, reason code MQRC_HANDLE_IN_USE_FOR_UOW is returned by the failing MQI call.)

On HP Integrity NonStop Server, this reason occurs:

- On a transactional MQI call when the client encounters a configuration error preventing it from enlisting with the TMF/Gateway, therefore preventing participation within a global unit of work that is coordinated by the Transaction Management Facility (TMF).
- If a client application makes an enlistment request before the TMF/Gateway completes recovery of in-doubt transactions, the request is held for up to 1 second. If recovery does not complete within that time, the enlistment is rejected.

Another cause of MQRC_UOW_ENLISTMENT_ERROR is incorrect installation; On Windows, for example, the Windows NT Service pack must be installed after the Windows NT Option pack.

Completion Code

MQCC_FAILED

Programmer response

On Windows, check the DTC "Transaction timeout" value. If necessary, verify the Windows NT installation order.

On HP Integrity NonStop Server this might be a configuration error. The client issues a message to the client error log providing extra information about the configuration error. Contact your system administrator to resolve the indicated error.

2355 (0933) (RC2355): MQRC_UOW_MIX_NOT_SUPPORTED:

Explanation

This reason code occurs only on Windows when you are running a version of the queue manager before version 5.2., and on HP Integrity NonStop Server.

On Windows, the following explanations might apply:

- The mixture of calls that is used by the application to perform operations within a unit of work is not supported. In particular, it is not possible to mix within the same process a local unit of work that is coordinated by the queue manager with a global unit of work that is coordinated by DTC (Distributed Transaction Coordinator).
- An application might cause this mixture to arise if some objects in a package are coordinated by DTC and others are not. It can also occur if transactional MQI calls from an MTS client are mixed with transactional MQI calls from a library package transactional MTS object.
- No problem arises if all transactional MQI calls originate from transactional MTS objects, or all transactional MQI calls originate from non-transactional MTS objects. But when a mixture of styles is

used, the first style that is used fixes the style for the unit of work, and subsequent attempts to use the other style within the process fail with reason code MQRC_UOW_MIX_NOT_SUPPORTED.

- When an application is run twice, scheduling factors in the operating system mean that it is possible for the queue-manager-coordinated transactional calls to fail in one run, and for the DTC-coordinated transactional calls to fail in the other run.

On HP Integrity NonStop Server it is not possible, within a single IBM WebSphere MQ connection, to issue transactional MQI calls under the coordination of the Transaction Management Facility (TMF) if transactional MQI calls have already been made within a local unit of work that is coordinated by the queue manager until the local unit of work is completed by issuing either MQCMIT or MQBACK.

Completion Code

MQCC_FAILED

Programmer response

On Windows, check that the “MTS Transaction Support” attribute defined for the object’s class is set correctly. If necessary, modify the application so that objects that run within different units of work do not try to use the same connection handle.

On HP Integrity NonStop Server, if a local unit of work that is coordinated by the queue manager is in progress, it must either be completed by issuing MQCMIT, or rolled back by issuing MQBACK before issuing any transactional MQI calls under the coordination of TMF.

2356 (0934) (RC2356): MQRC_WXP_ERROR:

Explanation

An MQXCLWLN call was issued from a cluster workload exit to obtain the address of the next record in the chain, but the workload exit parameter structure *ExitParms* is not valid, for one of the following reasons:

- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The *StrucId* field is not MQWXP_STRUC_ID.
- The *Version* field is not MQWXP_VERSION_2.
- The *CacheContext* field does not contain the value passed to the exit by the queue manager.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the parameter specified for *ExitParms* is the MQWXP structure that was passed to the exit when the exit was invoked.

2357 (0935) (RC2357): MQRC_CURRENT_RECORD_ERROR:

Explanation

An MQXCLWLN call was issued from a cluster workload exit to obtain the address of the next record in the chain, but the address specified by the *CurrentRecord* parameter is not the address of a valid record. *CurrentRecord* must be the address of a destination record (MQWDR), queue record (MQWQR), or cluster record (MQWCR) residing within the cluster cache.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the cluster workload exit passes the address of a valid record residing in the cluster cache.

2358 (0936) (RC2358): MQRC_NEXT_OFFSET_ERROR:

Explanation

An MQXCLWLN call was issued from a cluster workload exit to obtain the address of the next record in the chain, but the offset specified by the *NextOffset* parameter is not valid. *NextOffset* must be the value of one of the following fields:

- *ChannelDefOffset* field in MQWDR
- *ClusterRecOffset* field in MQWDR
- *ClusterRecOffset* field in MQWQR
- *ClusterRecOffset* field in MQWCR

Completion Code

MQCC_FAILED

Programmer response

Ensure that the value specified for the *NextOffset* parameter is the value of one of the fields listed.

2359 (0937) (RC2359): MQRC_NO_RECORD_AVAILABLE:

Explanation

An MQXCLWLN call was issued from a cluster workload exit to obtain the address of the next record in the chain, but the current record is the last record in the chain.

Completion Code

MQCC_FAILED

Programmer response

None.

2360 (0938) (RC2360): MQRC_OBJECT_LEVEL_INCOMPATIBLE:

Explanation

An MQOPEN or MQPUT1 call, or a command, was issued, but the definition of the object to be accessed is not compatible with the queue manager to which the application has connected. The object definition was created or modified by a different version of the queue manager.

If the object to be accessed is a queue, the incompatible object definition could be the object specified, or one of the object definitions used to resolve the specified object (for example, the base queue to which an alias queue resolves, or the transmission queue to which a remote queue or queue-manager alias resolves).

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The application must be run on a queue manager that is compatible with the object definition. See

 Migration paths: IBM WebSphere MQ for z/OS for more information about compatibility and migration between different versions of the queue manager.

2361 (0939) (RC2361): MQRC_NEXT_RECORD_ERROR:

Explanation

An MQXCLWLN call was issued from a cluster workload exit to obtain the address of the next record in the chain, but the address specified for the *NextRecord* parameter is either null, not valid, or the address of read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Specify a valid address for the *NextRecord* parameter.

2362 (093A) (RC2362): MQRC_BACKOUT_THRESHOLD_REACHED:

Explanation

This reason code occurs only in the *Reason* field in an MQDLH structure, or in the *Feedback* field in the MQMD of a report message.

A JMS ConnectionConsumer found a message that exceeds the queue's backout threshold. The queue does not have a backout requeue queue defined, so the message was processed as specified by the disposition options in the *Report* field in the MQMD of the message.

On queue managers that do not support the *BackoutThreshold* and *BackoutRequeueQName* queue attributes, JMS ConnectionConsumer uses a value of 20 for the backout threshold. When the *BackoutCount* of a message reaches this threshold, the message is processed as specified by the disposition options.

If the *Report* field specifies one of the MQRO_EXCEPTION_* options, this reason code appears in the *Feedback* field of the report message. If the *Report* field specifies MQRO_DEAD_LETTER_Q, or the disposition report options are left as default, this reason code appears in the *Reason* field of the MQDLH.

Completion Code

None

Programmer response

Investigate the cause of the backout count being greater than the threshold. To correct this, define the backout queue for the queue concerned.

2363 (093B) (RC2363): MQRC_MSG_NOT_MATCHED:

Explanation

This reason code occurs only in the *Reason* field in an MQDLH structure, or in the *Feedback* field in the MQMD of a report message.

While performing Point-to-Point messaging, JMS encountered a message matching none of the selectors of ConnectionConsumers monitoring the queue. To maintain performance, the message was processed as specified by the disposition options in the *Report* field in the MQMD of the message.

If the *Report* field specifies one of the MQRO_EXCEPTION_* options, this reason code appears in the *Feedback* field of the report message. If the *Report* field specifies MQRO_DEAD_LETTER_Q, or the disposition report options are left as default, this reason code appears in the *Reason* field of the MQDLH.

Completion Code

None

Programmer response

To correct this, ensure that the ConnectionConsumers monitoring the queue provide a complete set of selectors. Alternatively, set the QueueConnectionFactory to retain messages.

2364 (093C) (RC2364): MQRC_JMS_FORMAT_ERROR:

Explanation

This reason code is generated by JMS applications that use either:

- ConnectionConsumers
- Activation Specifications
- WebSphere Application Server Listener Ports

and connect to a WebSphere MQ queue manager using WebSphere MQ messaging provider migration mode. When the WebSphere MQ classes for JMS encounter a message that cannot be parsed (for example, the message contains an invalid RFH2 header) the message is processed as specified by the disposition options in the *Report* field in the MQMD of the message.

If the *Report* field specifies one of the MQRO_EXCEPTION_* options, this reason code appears in the *Feedback* field of the report message. If the *Report* field specifies MQRO_DEAD_LETTER_Q, or the disposition report options are left as default, this reason code appears in the *Reason* field of the MQDLH.

Completion Code

None

Programmer response

Investigate the origin of the message.

2365 (093D) (RC2365): MQRC_SEGMENTS_NOT_SUPPORTED:

Explanation

An MQPUT call was issued to put a segment of a logical message, but the queue on which the message is to be placed has an *IndexType* of MQIT_GROUP_ID. Message segments cannot be placed on queues with this index type.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Modify the application to put messages that are not segments; ensure that the MQMF_SEGMENT and MQMF_LAST_SEGMENT flags in the *MsgFlags* field in MQMD are not set, and that the *Offset* is zero. Alternatively, change the index type of the queue.

2366 (093E) (RC2366): MQRC_WRONG_CF_LEVEL:

Explanation

An MQOPEN or MQPUT1 call was issued specifying a shared queue, but the queue requires a coupling-facility structure with a different level of capability.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the coupling-facility structure used for the queue is at the level required to support the capabilities that the queue provides.

You can use the DISPLAY CFSTRUCT command to display the level, and ALTER CFSTRUCT() CFLEVEL() command to modify the level; see  The MQSC commands (*WebSphere MQ V7.1 Reference*).

2367 (093F) (RC2367): MQRC_CONFIG_CREATE_OBJECT:

Explanation

This condition is detected when an object is created.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2368 (0940) (RC2368): MQRC_CONFIG_CHANGE_OBJECT:

Explanation

This condition is detected when an object is changed.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2369 (0941) (RC2369): MQRC_CONFIG_DELETE_OBJECT:

Explanation

This condition is detected when an object is deleted.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2370 (0942) (RC2370): MQRC_CONFIG_REFRESH_OBJECT:

Explanation

This condition is detected when an object is refreshed.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2371 (0943) (RC2371): MQRC_CHANNEL_SSL_ERROR:

Explanation

This condition is detected when a connection cannot be established due to an SSL key-exchange or authentication failure.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2373 (0945) (RC2373): MQRC_CF_STRUC_FAILED:

Explanation

An MQI call or command was issued to access a shared queue, but the call failed because the coupling-facility structure used for the shared queue had failed.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Report the problem to the operator or administrator, who should use the MQSC command RECOVER CFSTRUCT to initiate recovery of the coupling-facility structure

2374 (0946) (RC2374): MQRC_API_EXIT_ERROR:

Explanation

An API exit function returned an invalid response code, or failed in some other way.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Check the exit logic to ensure that the exit is returning valid values in the *ExitResponse* and *ExitResponse2* fields of the MQAXP structure. Consult the FFST record to see if it contains more detail about the problem.

2375 (0947) (RC2375): MQRC_API_EXIT_INIT_ERROR:

Explanation

The queue manager encountered an error while attempting to initialize the execution environment for an API exit function.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Consult the FFST record to obtain more detail about the problem.

2376 (0948) (RC2376): MQRC_API_EXIT_TERM_ERROR:

Explanation

The queue manager encountered an error while attempting to terminate the execution environment for an API exit function.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Consult the FFST record to obtain more detail about the problem.

2377 (0949) (RC2377): MQRC_EXIT_REASON_ERROR:

Explanation

An MQXEP call was issued by an API exit function, but the value specified for the *ExitReason* parameter is either not valid, or not supported for the specified function identifier *Function*.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Modify the exit function to specify a value for *ExitReason* that is valid for the specified value of *Function*.

2378 (094A) (RC2378): MQRC_RESERVED_VALUE_ERROR:

Explanation

An MQXEP call was issued by an API exit function, but the value specified for the *Reserved* parameter is not valid. The value must be the null pointer.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Modify the exit to specify the null pointer as the value of the *Reserved* parameter.

2379 (094B) (RC2379): MQRC_NO_DATA_AVAILABLE:

Explanation

This reason should be returned by the MQZ_ENUMERATE_AUTHORITY_DATA installable service component when there is no more authority data to return to the invoker of the service component.

- On z/OS, this reason code does not occur.

Completion Code

MQCC_FAILED

Programmer response

None.

2380 (094C) (RC2380): MQRC_SCO_ERROR:

Explanation

On an MQCONNX call, the MQSCO structure is not valid for one of the following reasons:

- The *StrucId* field is not MQSCO_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Correct the definition of the MQSCO structure.

2381 (094D) (RC2381): MQRC_KEY_REPOSITORY_ERROR:

Explanation

On an MQCONN or MQCONNX call, the location of the key repository is either not specified, not valid, or results in an error when used to access the key repository. The location of the key repository is specified by one of the following:

- The value of the MQSSLKEYR environment variable (MQCONN or MQCONNX call), or
- The value of the *KeyRepository* field in the MQSCO structure (MQCONNX call only).

For the MQCONNX call, if both MQSSLKEYR and *KeyRepository* are specified, the latter is used.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid location for the key repository.

2382 (094E) (RC2382): MQRC_CRYPTO_HARDWARE_ERROR:

Explanation

On an MQCONN or MQCONNX call, the configuration string for the cryptographic hardware is not valid, or results in an error when used to configure the cryptographic hardware. The configuration string is specified by one of the following:

- The value of the MQSSLCRYP environment variable (MQCONN or MQCONNX call), or
- The value of the *CryptoHardware* field in the MQSCO structure (MQCONNX call only).

For the MQCONNX call, if both MQSSLCRYP and *CryptoHardware* are specified, the latter is used.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid configuration string for the cryptographic hardware.

2383 (094F) (RC2383): MQRC_AUTH_INFO_REC_COUNT_ERROR:

Explanation

On an MQCONNX call, the *AuthInfoRecCount* field in the MQSCO structure specifies a value that is less than zero.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a value for *AuthInfoRecCount* that is zero or greater.

2384 (0950) (RC2384): MQRC_AUTH_INFO_REC_ERROR:

Explanation

On an MQCONNX call, the MQSCO structure does not specify the address of the MQAIR records correctly. One of the following applies:

- *AuthInfoRecCount* is greater than zero, but *AuthInfoRecOffset* is zero and *AuthInfoRecPtr* is the null pointer.
- *AuthInfoRecOffset* is not zero and *AuthInfoRecPtr* is not the null pointer.
- *AuthInfoRecPtr* is not a valid pointer.
- *AuthInfoRecOffset* or *AuthInfoRecPtr* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Ensure that one of *AuthInfoRecOffset* or *AuthInfoRecPtr* is zero and the other nonzero. Ensure that the field used points to accessible storage.

2385 (0951) (RC2385): MQRC_AIR_ERROR:

Explanation

On an MQCONNX call, an MQAIR record is not valid for one of the following reasons:

- The *StrucId* field is not MQAIR_STRUC_ID.
- The *Version* field specifies a value that is not valid or not supported.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Correct the definition of the MQAIR record.

2386 (0952) (RC2386): MQRC_AUTH_INFO_TYPE_ERROR:**Explanation**

On an MQCONNX call, the *AuthInfoType* field in an MQAIR record specifies a value that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify MQAIT_CRL_LDAP for *AuthInfoType*.

2387 (0953) (RC2387): MQRC_AUTH_INFO_CONN_NAME_ERROR:**Explanation**

On an MQCONNX call, the *AuthInfoConnName* field in an MQAIR record specifies a value that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid connection name.

2388 (0954) (RC2388): MQRC_LDAP_USER_NAME_ERROR:**Explanation**

On an MQCONNX call, an LDAP user name in an MQAIR record is not specified correctly. One of the following applies:

- *LDAPUserNameLength* is greater than zero, but *LDAPUserNameOffset* is zero and *LDAPUserNamePtr* is the null pointer.
- *LDAPUserNameOffset* is nonzero and *LDAPUserNamePtr* is not the null pointer.
- *LDAPUserNamePtr* is not a valid pointer.
- *LDAPUserNameOffset* or *LDAPUserNamePtr* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Ensure that one of *LDAPUserNameOffset* or *LDAPUserNamePtr* is zero and the other nonzero. Ensure that the field used points to accessible storage.

2389 (0955) (RC2389): MQRC_LDAP_USER_NAME_LENGTH_ERR:**Explanation**

On an MQCONN call, the *LDAPUserNameLength* field in an MQAIR record specifies a value that is less than zero.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a value for *LDAPUserNameLength* that is zero or greater.

2390 (0956) (RC2390): MQRC_LDAP_PASSWORD_ERROR:**Explanation**

On an MQCONN call, the *LDAPPassword* field in an MQAIR record specifies a value when no value is allowed.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Specify a value that is blank or null.

2391 (0957) (RC2391): MQRC_SSL_ALREADY_INITIALIZED:**Explanation**

An MQCONN or MQCONN call was issued when a connection is already open to the same queue manager. There is a conflict between the SSL options of the connections for one of three reasons:

- The SSL configuration options are different between the first and second connections.
- The existing connection was specified without SSL configuration options, but the second connection has SSL configuration options specified.
- The existing connection was specified with SSL configuration options, but the second connection does not have any SSL configuration options specified.

The connection to the queue manager completed successfully, but the SSL configuration options specified on the call were ignored; the existing SSL environment was used instead.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_WARNING

Programmer response

If the application must be run with the SSL configuration options defined on the MQCONN or MQCONNX call, use the MQDISC call to sever the connection to the queue manager and then stop the application. Alternatively run the application later when the SSL environment has not been initialized.

2392 (0958) (RC2392): MQRC_SSL_CONFIG_ERROR:

Explanation

On an MQCONNX call, the MQCNO structure does not specify the MQSCO structure correctly. One of the following applies:

- *SSLConfigOffset* is nonzero and *SSLConfigPtr* is not the null pointer.
- *SSLConfigPtr* is not a valid pointer.
- *SSLConfigOffset* or *SSLConfigPtr* points to storage that is not accessible.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Ensure that one of *SSLConfigOffset* or *SSLConfigPtr* is zero and the other nonzero. Ensure that the field used points to accessible storage.

2393 (0959) (RC2393): MQRC_SSL_INITIALIZATION_ERROR:

Explanation

An MQCONN or MQCONNX call was issued with SSL configuration options specified, but an error occurred during the initialization of the SSL environment.

This reason code occurs in the following environments: AIX, HP-UX, Solaris, Windows.

Completion Code

MQCC_FAILED

Programmer response

Check that the SSL installation is correct.

2394 (095A) (RC2394): MQRC_Q_INDEX_TYPE_ERROR:

Explanation

An MQGET call was issued specifying one or more of the following options:

- MQGMO_ALL_MSGS_AVAILABLE
- MQGMO_ALL_SEGMENTS_AVAILABLE
- MQGMO_COMPLETE_MSG
- MQGMO_LOGICAL_ORDER

but the call failed because the queue is not indexed by group identifier. These options require the queue to have an *IndexType* of MQIT_GROUP_ID.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

Redefine the queue to have an *IndexType* of MQIT_GROUP_ID. Alternatively, modify the application to avoid using the options listed.

2395 (095B) (RC2395): MQRC_CFBS_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFBS structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2396 (095C) (RC2396): MQRC_SSL_NOT_ALLOWED:

Explanation

A connection to a queue manager was requested, specifying SSL encryption. However, the connection mode requested is one that does not support SSL (for example, bindings connect).

Completion Code

MQCC_FAILED

Programmer response

Modify the application to request client connection mode, or to disable SSL encryption.

Note: Using a non null setting, including blanks, for the connection's cipher suite property can also cause this error.

2397 (095D) (RC2397): MQRC_JSSE_ERROR:**Explanation**

JSSE reported an error (for example, while connecting to a queue manager using SSL encryption). The MQException object containing this reason code references the Exception thrown by JSSE; this can be obtained by using the MQException.getCause() method. From JMS, the MQException is linked to the thrown JMSEException.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Inspect the causal exception to determine the JSSE error.

2398 (095E) (RC2398): MQRC_SSL_PEER_NAME_MISMATCH:**Explanation**

The application attempted to connect to the queue manager using SSL encryption, but the distinguished name presented by the queue manager does not match the specified pattern.

Completion Code

MQCC_FAILED

Programmer response

Check the certificates used to identify the queue manager. Also check the value of the sslPeerName property specified by the application.

2399 (095F) (RC2399): MQRC_SSL_PEER_NAME_ERROR:**Explanation**

The application specified a peer name of incorrect format.

Completion Code

MQCC_FAILED

Programmer response

Check the value of the sslPeerName property specified by the application.

2400 (0960) (RC2400): MQRC_UNSUPPORTED_CIPHER_SUITE:

Explanation

A connection to a queue manager was requested, specifying SSL encryption. However, JSSE reported that it does not support the CipherSuite specified by the application.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Check the CipherSuite specified by the application. Note that the names of JSSE CipherSuites differ from their equivalent CipherSpecs used by the queue manager.

Also, check that JSSE is correctly installed.

2401 (0961) (RC2401): MQRC_SSL_CERTIFICATE_REVOKED:

Explanation

A connection to a queue manager was requested, specifying SSL encryption. However, the certificate presented by the queue manager was found to be revoked by one of the specified CertStores.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Check the certificates used to identify the queue manager.

2402 (0962) (RC2402): MQRC_SSL_CERT_STORE_ERROR:

Explanation

A connection to a queue manager was requested, specifying SSL encryption. However, none of the CertStore objects provided by the application could be searched for the certificate presented by the queue manager. The MQException object containing this reason code references the Exception encountered when searching the first CertStore; this can be obtained using the MQException.getCause() method. From JMS, the MQException is linked to the thrown JMSEException.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Inspect the causal exception to determine the underlying error. Check the CertStore objects provided by your application. If the causal exception is a java.lang.NoSuchElementException, ensure that your application is not specifying an empty collection of CertStore objects.

2406 (0966) (RC2406): MQRC_CLIENT_EXIT_LOAD_ERROR:

Explanation

The external user exit required for a client connection could not be loaded because the shared library specified for it cannot be found, or the entry point specified for it cannot be found.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct library has been specified, and that the path variable for the machine environment includes the relevant directory. Ensure also that the entry point has been named properly and that the named library does export it.

2407 (0967) (RC2407): MQRC_CLIENT_EXIT_ERROR:

Explanation

A failure occurred while executing a non-Java user exit for a client connection.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Check that the non-Java user exit can accept the parameters and message being passed to it and that it can handle error conditions, and that any information that the exit requires, such as user data, is correct and available.

2409 (0969) (RC2409): MQRC_SSL_KEY_RESET_ERROR:

Explanation

On an MQCONN or MQCONNX call, the value of the SSL key reset count is not in the valid range of 0 through 999 999 999.

The value of the SSL key reset count is specified by either the value of the MQSSLRESET environment variable (MQCONN or MQCONNX call), or the value of the *KeyResetCount* field in the MQSCO structure (MQCONNX call only). For the MQCONNX call, if both MQSSLRESET and *KeyResetCount* are specified, the latter is used. MQCONN or MQCONNX

If you specify an SSL/TLS secret key reset count in the range 1 byte through 32Kb, SSL/TLS channels will use a secret key reset count of 32Kb. This is to avoid the overhead of excessive key resets which would occur for small SSL/TLS secret key reset values.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure and the MQSSLRESET environment variable are set correctly.

2411 (096B) (RC2411): MQRC_LOGGER_STATUS:

Explanation

This condition is detected when a logger event occurs.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2412 (096C) (RC2412): MQRC_COMMAND_MQSC:

Explanation

This condition is detected when an MQSC command is executed.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2413 (096D) (RC2413): MQRC_COMMAND_PCF:

Explanation

This condition is detected when a PCF command is executed.

Completion Code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2414 (096E) (RC2414): MQRC_CFIF_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFIF structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2415 (096F) (RC2415): MQRC_CFSF_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFSF structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2416 (0970) (RC2416): MQRC_CFGR_ERROR:**Explanation**

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFGR structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, z/OS, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2417 (0971) (RC2417): MQRC_MSG_NOT_ALLOWED_IN_GROUP:

An explanation of the error, completion code, and programmer response.

Explanation

An MQPUT or MQPUT1 call was issued to put a message in a group but it is not valid to put such a message in a group. An example of an invalid message is a PCF message where the Type is MQCFT_TRACE_ROUTE.

You cannot use grouped or segmented messages with Publish/Subscribe.

Completion Code

MQCC_FAILED

Programmer response

Remove the invalid message from the group.

2418 (0972) (RC2418): MQRC_FILTER_OPERATOR_ERROR:**Explanation**

The **Operator** parameter supplied is not valid.

If it is an input variable then the value is not one of the MQCFOP_* constant values. If it is an output variable then the parameter pointer is not valid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer response

Correct the parameter.

2419 (0973) (RC2419): MQRC_NESTED_SELECTOR_ERROR:**Explanation**

An mqAddBag call was issued, but the bag to be nested contained a data item with an inconsistent selector. This reason only occurs if the bag into which the nested bag was to be added was created with the MQCBO_CHECK_SELECTORS option.

Completion Code

MQCC_FAILED

Programmer response

Ensure that all data items within the bag to be nested have selectors that are consistent with the data type implied by the item.

2420 (0974) (RC2420): MQRC_EPH_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQEPH structure that is not valid. Possible errors include the following:

- The *StrucId* field is not MQEPH_STRUC_ID.
- The *Version* field is not MQEPH_VERSION_1.
- The *StrucLength* field specifies a value that is too small to include the structure plus the variable-length data at the end of the structure.
- The *CodedCharSetId* field is zero, or a negative value that is not valid.
- The *Flags* field contains an invalid combination of MQEPH_* values.
- The *BufferLength* parameter of the call has a value that is too small to accommodate the structure, so the structure extends beyond the end of the message.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly. Ensure that the application sets the *CodedCharSetId* field to a valid value; note that MQCCSI_DEFAULT, MQCCSI_EMBEDDED, MQCCSI_Q_MGR, and MQCCSI_UNDEFINED are not valid in this field.

2421 (0975) (RC2421): MQRC_RFH_FORMAT_ERROR:

Explanation

The message contains an MQRFH structure, but its format is incorrect. If you are using WebSphere MQ SOAP, the error is in an incoming SOAP/MQ request message.

Completion Code

MQCC_FAILED

Programmer response

If you are using WebSphere MQ SOAP with the IBM-supplied sender, contact your IBM support center. If you are using WebSphere MQ SOAP with a bespoke sender, check that the RFH2 section of the SOAP/MQ request message is in valid RFH2 format.

2422 (0976) (RC2422): MQRC_CFBF_ERROR:

Explanation

An MQPUT or MQPUT1 call was issued, but the message data contains an MQCFBF structure that is not valid.

This reason code occurs in the following environments: AIX, HP-UX, IBM i, Solaris, Windows, plus WebSphere MQ clients connected to these systems.

Completion Code

MQCC_FAILED

Programmer response

Check that the fields in the structure are set correctly.

2423 (0977) (RC2423): MQRC_CLIENT_CHANNEL_CONFLICT:

Explanation

A client channel definition table (CCDT) was specified for determining the name of the channel, but the name has already been defined.

This reason code occurs only with Java applications.

Completion Code

MQCC_FAILED

Programmer response

Change the channel name to blank and try again.

2424 (0978) (RC2424): MQRC_SD_ERROR:

Explanation

On the MQSUB call, the Subscription Descriptor MQSD is not valid, for one of the following reasons:

- The StrucId field is not MQSD_SCTRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid (it is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results can occur).
- The queue manager cannot copy the changes structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQSD structure are set correctly.

2425 (0979) (RC2425): MQRC_TOPIC_STRING_ERROR:

Explanation

On the MQOPEN or MQPUT1 call in the Object Descriptor MQOD, or on the MQSUB call in the Subscription Descriptor MQSD the resultant full topic string is not valid.

One of the following applies:

- ObjectName contains the name of a TOPIC object with a TOPICSTR attribute that contains an empty topic string.
- The fully resolved topic string contains the escape character '%' and it is not followed by one of the characters, '*', '?', or '%', and the MQSO_WILDCARD_CHAR option has been used on an MQSUB call.

- On an MQOPEN, conversion cannot be performed using the CCSID specified in the MQOD structure.
- The topic string is greater than 255 characters when using WebSphere MQ Multicast messaging.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that there are no invalid topic string characters in either ObjectString or ObjectName.

If using WebSphere MQ Multicast messaging, ensure that the topic string is less than 255 characters.

2426 (097A) (RC2426): MQRC_STS_ERROR:

Explanation

On an MQSTAT call, the MQSTS structure is not valid, for one of the following reasons:

- The StrucId field is not MQSTS_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQSTS structure are set correctly.

2428 (097C) (RC2428): MQRC_NO_SUBSCRIPTION:

Explanation

An MQSUB call using option MQSO_RESUME was made specifying a full subscription name that does not match any existing subscription.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the subscription exists and that the full subscription name is correctly specified in your application. The full subscription name is built from the ConnTag field specified at connection time in the MQCNO structure and the SubName field specified at MQSUB time in the MQSD structure.

2429 (097D) (RC2429): MQRC_SUBSCRIPTION_IN_USE:

Explanation

An MQSUB call using option MQSO_RESUME was made specifying a full subscription name that is in use.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the subscription name is correctly specified in your application. The subscription name is specified in the SubName field in the MQSD structure.

2430 (097E) (RC2430): MQRC_STAT_TYPE_ERROR:

Explanation

The STS parameter contains options that are not valid for the MQSTAT call. This reason also occurs if the parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Programmer response

Specify a valid MQSTS structure as a parameter on the call to MQSTAT.

2431 (097F) (RC2431): MQRC_SUB_USER_DATA_ERROR:

Explanation

On the MQSUB call in the Subscription Descriptor MQSD the SubUserData field is not valid. One of the following applies:

- SubUserData.VSLength is greater than zero, but SubUserData.VSOffset is zero and SubUserData.VSPtr is the null pointer.
- SubUserData.VSOffset is nonzero and SubUserData.VSPtr is not the null pointer (that is, it appears both fields are being used where only one is allowed).
- SubUserData.VSPtr is not a valid pointer.
- SubUserData.VSOffset or SubUserData.VSPtr points to storage that is not accessible.
- SubUserData.VSLength exceeds the maximum length allowed for this field.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that one of SubUserData.VSOffset or SubUserData.VSPtr is zero and the other nonzero. Ensure that the field used points to accessible storage. Specify a length that does not exceed the maximum length allowed for this field.

2432 (0980) (RC2432): MQRC_SUB_ALREADY_EXISTS:**Explanation**

An MQSUB call was issued to create a subscription, using the MQSO_CREATE option, but a subscription using the same SubName and ObjectString already exists.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the SubName and ObjectString input fields in the MQSD structure are set correctly, or use the MQSO_RESUME option to get a handle for the subscription that already exists.

2434 (0982) (RC2434): MQRC_IDENTITY_MISMATCH:**Explanation**

An MQSUB call using either MQSO_RESUME or MQSO_ALTER was made against a subscription that has the MQSO_FIXED_USERID option set, by a userid other than the one recorded as owning the subscription.

Completion Code

MQCC_FAILED

Programmer Response

Correct the full subscription name to one that is unique, or update the existing subscription to allow different userids to use it by using the MQSO_ANY_USERID option from an application running under the owning userid.

2435 (0983) (RC2435): MQRC_ALTER_SUB_ERROR:**Explanation**

An MQSUB call using option MQSO_ALTER was made changing a subscription that was created with the MQSO_IMMUTABLE option.

Completion Code

MQCC_FAILED

Programmer Response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly.

2436 (0984) (RC2436): MQRC_DURABILITY_NOT_ALLOWED:

Explanation

An MQSUB call using the MQSO_DURABLE option failed. This can be for one of the following reasons:

- The topic subscribed to is defined as DURSUB(NO).
- The queue named SYSTEM.DURABLE.SUBSCRIBER.QUEUE is not available.
- The topic subscribed to is defined as both MCAST(ONLY) and DURSUB(YES) (or DURSUB(ASPARENT) and the parent is DURSUB(YES)).

Completion Code

MQCC_FAILED

Programmer Response

Durable subscriptions are stored on the SYSTEM.DURABLE.SUBSCRIBER.QUEUE. Ensure that this queue is available for use. Possible reasons for failure include the queue being full, the queue being put inhibited, the queue not existing, or (on z/OS) the pageset the queue is defined to use doesn't exist.

If the topic subscribed to is defined as DURSUB(NO) either alter the administrative topic node to use DURSUB(YES) or use the MQSO_NON_DURABLE option instead.

If the topic subscribed to is defined as MCAST(ONLY) when using WebSphere MQ Multicast messaging, alter the topic to use DURSUB(NO).

2437 (0985) (RC2437): MQRC_NO_RETAINED_MSG:

Explanation

An MQSUBRQ call was made to a topic to request that any retained publications for this topic are sent to the subscriber. However, there are no retained publications currently stored for this topic.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that publishers to the topic are marking their publication to be retained and that publications are being made to this topic.

2438 (0986) (RC2438): MQRC_SRO_ERROR:

Explanation

On the MQSUBRQ call, the Subscription Request Options MQSRO is not valid, for one of the following reasons:

- The StrucId field is not MQSRO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQSRO structure are set correctly.

2440 (0988) (RC2440): MQRC_SUB_NAME_ERROR:

Explanation

On the MQSUB call in the Subscription Descriptor MQSD the SubName field is not valid or has been omitted. This is required if the MQSD option MQSO_DURABLE is specified, but may also be used if MQSO_DURABLE is not specified.

One of the following applies:

- SubName.VSLength is greater than zero, but SubName.VSOffset is zero and SubName.VSPtr is the null pointer.
- SubName.VSOffset is nonzero and SubName.VSPtr is not the null pointer (that is, it appears both fields are being used where only one is allowed).
- SubName.VSPtr is not a valid pointer.
- SubName.VSOffset or SubName.VSPtr points to storage that is not accessible.
- SubName.VSLength is zero but this field is required.
- SubName.VSLength exceeds the maximum length allowed for this field.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that SubName is specified and SubName.VSLength is nonzero. Ensure that one of SubName.VSOffset or SubName.VSPtr is zero and the other nonzero. Ensure that the field used points to accessible storage. Specify a length that does not exceed the maximum length allowed for this field.

This code can be returned if the sd.Options flags MQSO_CREATE and MQSO_RESUME are set together and sd.SubName is not initialized. You must also initialize the MQCHARV structure for sd.SubName, even if there is no subscription to resume; see  Example 2: Managed MQ subscriber (*WebSphere MQ V7.1 Programming Guide*) for more details.

2441 (0989) (RC2441): MQRC_OBJECT_STRING_ERROR:

Explanation

On the MQOPEN or MQPUT1 call in the Object Descriptor MQOD, or on the MQSUB call in the Subscription Descriptor MQSD the ObjectString field is not valid.

One of the following applies:

- ObjectString.VSLength is greater than zero, but ObjectString.VSOffset is zero and ObjectString.VSPtr is the null pointer.
- ObjectString.VSOffset is nonzero and ObjectString.VSPtr is not the null pointer (that is, it appears both fields are being used where only one is allowed).

- ObjectString.VSPtr is not a valid pointer.
- ObjectString.VSOffset or ObjectString.VSPtr points to storage that is not accessible.
- ObjectString.VSLength exceeds the maximum length allowed for this field.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that one of ObjectString.VSOffset or ObjectString.VSPtr is zero and the other nonzero. Ensure that the field used points to accessible storage. Specify a length that does not exceed the maximum length allowed for this field.

2442 (098A) (RC2442): MQRC_PROPERTY_NAME_ERROR:

Explanation

An attempt was made to set a property with an invalid or restricted name. With validation set to off, only restricted names will result in this error. With validation set to on, any of the following property names may result in this error:

- The name contains a logical grouping or "." on a MQSETMP call specifying MQSMPO_SET_PROP_BEFORE_CURSOR or MQSMPO_SET_PROP_AFTER_CURSOR, with a grouping that is different to the cursor location.
- The name contains an invalid character.
- The name begins "JMS" or "usr.JMS" and the JMS property is not recognized.
- The name begins "mq" in any mixture of lowercase or uppercase and is not "mq_usr" and contains more than one "." character (U+002E). Multiple "." characters are not allowed in properties with those prefixes.
- The name is "NULL", "TRUE", "FALSE", "NOT", "AND", "OR", "BETWEEN", "LIKE", "IN", "IS" and "ESCAPE" or is one of these keywords prefixed by "usr."
- The name begins with "Body" or "Root" (except for names beginning "Root.MQMD.").
- The name is a restricted name, see  Property name restrictions (*WebSphere MQ V7.1 Programming Guide*).
- A "." character must not be followed immediately by another "." character.
- The "." character cannot be the last character in a property name.

Completion Code

MQCC_FAILED

Programmer Response

Valid property names are described in the WebSphere MQ documentation. Ensure that all properties in the message have valid names before reissuing the call.

2443 (098B) (RC2443): MQRC_SEGMENTATION_NOT_ALLOWED:**Explanation**

An MQPUT or MQPUT1 call was issued to put a segmented message or a message that may be broken up into smaller segments (MQMF_SEGMENTATION_ALLOWED). The message was found to contain one or more MQ-defined properties in the message data; MQ-defined properties are not valid in the message data of a segmented message.

WebSphere MQ Multicast cannot use segmented messages.

Completion Code

MQCC_FAILED

Programmer Response

Remove the invalid properties from the message data or prevent the message from being segmented.

2444 (098C) (RC2444): MQRC_CBD_ERROR:**Explanation**

a MQCB call the MQCBD structure is not valid for one of the following reasons:

- The StrucId field is not MQCBD_STRUC_ID
- The Version field is specifies a value that is not valid or is not supported
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQCBD structure are set correctly.

2445 (098D) (RC2445): MQRC_CTLO_ERROR:**Explanation**

On a MQCTL call the MQCTLO structure is not valid for one of the following reasons:

- The StrucId field is not MQCTLO_STRUC_ID
- The Version field is specifies a value that is not valid or is not supported
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQCTLO structure are set correctly.

2446 (098E) (RC2446): MQRC_NO_CALLBACKS_ACTIVE:**Explanation**

An MQCTL call was made with an Operation of MQOP_START_WAIT and has returned because there are no currently defined callbacks which are not suspended.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that there is at least one registered, resumed consumer function.

2448 (0990) (RC2448): MQRC_CALLBACK_NOT_REGISTERED:**Explanation**

An attempt to issue an MQCB call has been made against an object handle which does not currently have a registered callback.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that a callback has been registered against the object handle.

2449 (0991) (RC2449): MQRC_OPERATION_NOT_ALLOWED:**Explanation**

An MQCTL call was made with an Operation that is not allowed because of the state of asynchronous consumption on the hConn is currently in.

If Operation was MQOP_RESUME, the operation is not allowed because the state of asynchronous consumption on the hConn is STOPPED. Re-issue MQCTL with the MQOP_START Operation.

If Operation was MQOP_SUSPEND, the operation is not allowed because the state of asynchronous consumption on the hConn is STOPPED. If you need to get your hConn into a SUSPENDED state, issue MQCTL with the MQOP_START Operation followed by MQCTL with MQOP_SUSPEND.

If Operation was MQOP_START, the operation is not allowed because the state of asynchronous consumption on the hConn is SUSPENDED. Re-issue MQCTL with the MQOP_RESUME Operation.

If Operation was MQOP_START_WAIT, the operation is not allowed because either

- The state of asynchronous consumption on the hConn is SUSPENDED. Re-issue MQCTL with the MQOP_RESUME Operation.
- The state of asynchronous consumption on the hConn is already STARTED. Do not mix the use of MQOP_START and MQOP_START_WAIT within one application.

Completion Code

MQCC_FAILED

Programmer Response

Re-issue the MQCTL call with the correct Operation.

2457 (0999) (RC2457): MQRC_OPTIONS_CHANGED:

Explanation

An MQGET call on a queue handle opened using MQOO_READ_AHEAD (or resolved to that value through the queue's default value) has altered an option that is required to be consistent between MQGET calls.

Completion Code

MQCC_FAILED

Programmer Response

Keep all required MQGET options the same between invocations of MQGET, or use MQOO_NO_READ_AHEAD when opening the queue.

2458 (099A) (RC2458): MQRC_READ_AHEAD_MSGS:

Explanation

On an MQCLOSE call, the option MQCO_QUIESCE was used and there are still messages stored in client read ahead buffer that were sent to the client ahead of an application requesting them and have not yet been consumed by the application.

Completion Code

MQCC_WARNING

Programmer Response

Continue to consume messages using the queue handle until there are no more available and then issue the MQCLOSE again, or choose to discard these messages by issuing the MQCLOSE call with the MQCO_IMMEDIATE option instead.

2459 (099B) (RC2459): MQRC_SELECTOR_SYNTAX_ERROR:

Explanation

An MQOPEN, MQPUT1 or MQSUB call was issued but a selection string was specified which contained a syntax error.

Completion Code

MQCC_FAILED

Programmer Response

See  Message selector syntax (*WebSphere MQ V7.1 Programming Guide*) and ensure that you have correctly followed the rules for specifying selection strings. Correct any syntax errors and resubmit the MQ API call for which the error occurred.

2460 (099C) (RC2460): MQRC_HMSG_ERROR:

Explanation

On an MQCRTMH, MQDLTMH, MQSETMP, MQINQMP or MQDLT call, a message handle supplied is not valid, for one of the following reasons:

- The parameter pointer is not valid, or (for the MQCRTMH call) points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The value specified was not returned by a preceding MQCRTMH call.
- The value specified has been made invalid by a preceding MQDLTMH call.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that a successful MQCRTMH call is performed for the connection, and that an MQDLTMH call has not already been performed for it. Ensure that the handle is being used within its valid scope (see the description of MQCRTMH in the WebSphere MQ documentation).

2461 (099D) (RC2461): MQRC_CMHO_ERROR:

Explanation

On an MQCRTMH call, the create message handle options structure MQCMHO is not valid, for one of the following reasons:

- The StrucId field is not MQCMHO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQCMHO structure are set correctly.

2462 (099E) (RC2462): MQRC_DMHO_ERROR:

Explanation

On an MQDLTMH call, the delete message handle options structure MQDMHO is not valid, for one of the following reasons:

- The StrucId field is not MQDMHO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQDMHO structure are set correctly.

2463 (099F) (RC2463): MQRC_SMPO_ERROR:

Explanation

On an MQSETMP call, the set message property options structure MQSMPO is not valid, for one of the following reasons:

- The StrucId field is not MQSMPO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQSMPO structure are set correctly.

2464 (09A0) (RC2464): MQRC_IMPO_ERROR:

Explanation

On an MQINQMP call, the inquire message property options structure MQIMPO is not valid, for one of the following reasons:

- The StrucId field is not MQIMPO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The queue manager cannot copy the changed structure to application storage, even though the call is successful. This can occur, for example, if the pointer points to read-only storage.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQIMPO structure are set correctly.

2465 (09A1) (RC2465): MQRC_PROPERTY_NAME_TOO_BIG:

Explanation

On an MQINQMP call, WebSphere MQ attempted to copy the name of the inquired property into the location indicated by the ReturnedName field of the InqPropOpts parameter but the buffer was too small to contain the full property name. The call failed but the VSLength field of the ReturnedName of the InqPropOpts parameter indicates how large the ReturnedName buffer needs to be.

Completion Code

MQCC_FAILED

Programmer response

The full property name can be retrieved by calling MQINQMP again with a larger buffer for the returned name, also specifying the MQIMPO_INQ_PROP_UNDER_CURSOR option. This will inquire on the same property.

2466 (09A2) (RC2466): MQRC_PROP_VALUE_NOT_CONVERTED:

Explanation

An MQINQMP call was issued with the MQIMPO_CONVERT_VALUE option specified in the InqPropOpts parameter, but an error occurred during conversion of the value of the property. The property value is returned unconverted, the values of the ReturnedCCSID and ReturnedEncoding fields in the InqPropOpts parameter are set to those of the value returned.

Completion Code

MQCC_FAILED

Programmer Response

Check that the property value is correctly described by the ValueCCSID and ValueEncoding parameters that were specified when the property was set. Also check that these values, and the RequestedCCSID and RequestedEncoding specified in the InqPropOpts parameter of the MQINQMP call, are supported for MQ conversion. If the required conversion is not supported, conversion must be carried out by the application.

2467 (09A3) (RC2467): MQRC_PROP_TYPE_NOT_SUPPORTED:

Explanation

An MQINQMP call was issued and the property inquired has an unsupported data type. A string representation of the value is returned and the TypeString field of the InqPropOpts parameter can be used to determine the data type of the property.

Completion Code

MQCC_WARNING

Programmer Response

Check whether the property value was intended to have a data type indicated by the TypeString field. If so the application must decide how to interpret the value. If not modify the application that set the

property to give it a supported data type.

2469 (09A5) (RC2469): MQRC_PROPERTY_VALUE_TOO_BIG:

Explanation

On an MQINQMP call, the property value was too large to fit into the supplied buffer. The DataLength field is set to the length of the property value before truncation and the Value parameter contains as much of the value as fits.

On an MQMHBUF call, the BufferLength was less than the size of the properties to be put in the buffer. In this case the call fails. The DataLength field is set to the length of the properties before truncation.

Completion Code

MQCC_WARNING

MQCC_FAILED

Programmer Response

Supply a buffer that is at least as large as DataLength if all of the property value data is required and call MQINQMP again with the MQIMPO_INQ_PROP_UNDER_CURSOR option specified.

2470 (09A6) (RC2470): MQRC_PROP_CONV_NOT_SUPPORTED:

Explanation

On an MQINQMP call, the MQIMPO_CONVERT_TYPE option was specified to request that the property value be converted to the supplied data type before the call returned. Conversion between the actual and requested property data types is not supported. The Type parameter indicates the data type of the property value.

Completion Code

MQCC_FAILED

Programmer Response

Either call MQINQMP again without MQIMPO_CONVERT_TYPE specified, or request a data type for which conversion is supported.

2471 (09A7) (RC2471): MQRC_PROPERTY_NOT_AVAILABLE:

Explanation

On an MQINQMP call, no property could be found that matched the specified name. When iterating through multiple properties, possibly using a name containing a wildcard character, this indicates that all properties matching the name have now been returned.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the correct property name was specified. If the name contains a wildcard character specify option MQIMPO_INQ_FIRST to begin iterating over the properties again.

2472 (09A8) (RC2472): MQRC_PROP_NUMBER_FORMAT_ERROR:

Explanation

On an MQINQMP call, conversion of the property value was requested. The format of the property is invalid for conversion to the requested data type.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the correct property name and data type were specified. Ensure that the application setting the property gave it the correct format. See the documentation for the MQINQMP call for details on the formats required for data conversion of property values.

2473 (09A9) (RC2473): MQRC_PROPERTY_TYPE_ERROR:

Explanation

On an MQSETMP call, the Type parameter does not specify a valid MQTYPE_* value. For properties beginning "Root.MQMD." or "JMS" the specified Type must correspond to the data type of the matching MQMD or JMS header field:

- For MQCHARn or Java String fields use MQTYPE_STRING.
- For MQLONG or Java int fields use MQTYPE_INT32.
- For MQBYTEn fields use MQTYPE_BYTE_STRING.
- For Java long fields use MQTYPE_INT64.

On an MQINQMP call, the Type parameter is not valid. Either the parameter pointer is not valid, the value is invalid, or it points to read-only storage. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Correct the parameter.

2478 (09AE) (RC2478): MQRC_PROPERTIES_TOO_BIG:

Explanation

An MQPUT or MQPUT1 call was issued to put a message on a queue, but the properties of the message were too large. The length of the properties cannot exceed the value of the MaxPropertiesLength queue manager attribute.

This return code will also be issued if a message with headers greater than 511K is put to a shared queue using OFFLOAD(DB2)

Completion Code

MQCC_FAILED

Programmer Response

Consider one of the following actions:

- Reduce the number or the size of the properties associated with the message. This could include moving some of the properties into the application data.
- Increase the value of the MaxPropertiesLength queue manager attribute.

If the message has a header greater than 511K and was written to a shared queue, consider using shared message data sets (SMDS) instead of Db2 as a means of offloading the data.

2479 (09AF) (RC2479): MQRC_PUT_NOT_RETAINED:

Explanation

An MQPUT or MQPUT1 call was issued to publish a message on a topic, using the MQPMO_RETAIN option, but the publication was unable to be retained. The publication is not published to any matching subscribers.

Completion Code

MQCC_FAILED

Programmer Response

Retained publications are stored on the SYSTEM.RETAINED.PUB.QUEUE. Ensure that this queue is available for use by the application. Possible reasons for failure include the queue being full, the queue being put inhibited, or the queue not existing.

2480 (09B0) (RC2480): MQRC_ALIAS_TARGTYPE_CHANGED:

Explanation

An MQPUT or MQPUT1 call was issued to publish a message on a topic. One of the subscriptions matching this topic was made with a destination queue that was an alias queue which originally referenced a queue, but now references a topic object, which is not allowed. In this situation the reason code MQRC_ALIAS_TARGTYPE_CHANGED is returned in the Feedback field in the MQMD of a report message, or in the Reason field in the MQDLH structure of a message on the dead-letter queue.

Completion Code

MQCC_FAILED

Programmer Response

Find the subscriber that is using an alias queue which references a topic object and change it to reference a queue again, or change the subscription to reference a different queue.

2481 (09B1) (RC2481): MQRC_DMPO_ERROR:

Explanation

On an MQDLTMP call, the delete message property options structure MQDMPO is not valid, for one of the following reasons:

- The StrucId field is not MQDMPO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQDMPO structure are set correctly.

2482 (09B2) (RC2482): MQRC_PD_ERROR:

Explanation

On an MQSETMP or MQINQMP call, the property descriptor structure MQPD is not valid, for one of the following reasons:

- The StrucId field is not MQPD_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)
- The Context field contains an unrecognized value.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQPD structure are set correctly.

2483 (09B3) (RC2483): MQRC_CALLBACK_TYPE_ERROR:

Explanation

An MQCB call was made with an Operation of MQOP_REGISTER with an incorrect value for CallbackType

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the CallbackType field of the MQCBDO is specified correctly.

2484 (09B4) (RC2484): MQRC_CBD_OPTIONS_ERROR:**Explanation**

An MQCB call was made with an Operation of MQOP_REGISTER with an incorrect value for the Options field of the MQCBD.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the Options are specified correctly.

2485 (09B5) (RC2485): MQRC_MAX_MSG_LENGTH_ERROR:**Explanation**

An MQCB call was made with an Operation of MQOP_REGISTER with an incorrect value for the MaxMsgLength field of the MQCBD.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the MaxMsgLength are specified correctly.

2486 (09B6) (RC2486): MQRC_CALLBACK_ROUTINE_ERROR:**Explanation**

An MQCB call was made with an Operation of MQOP_REGISTER failed for one of the following reasons:

- Both CallbackName and CallbackFunction are specified. Only one must be specified on the call.
- The call was made from an environment not supporting function pointers.
- A programming language that does not support Function pointer references.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the CallbackName value is specified correctly.

2487 (09B7) (RC2487): MQRC_CALLBACK_LINK_ERROR:

Explanation

On an MQCTL call, the callback handling module (CSQBMCSM or CSQBMCSX for batch and DFHMQMCM for CICS) could not be loaded, so the adapter could not link to it.

This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the correct library concatenation has been specified in the application program execution JCL, and in the queue-manager startup JCL. Any uncommitted changes in a unit of work should be backed out. A unit of work that is coordinated by the queue manager is backed out automatically.

2488 (09B8) (RC2488): MQRC_OPERATION_ERROR:

Explanation

An MQCTL or MQCB call was made with an invalid **Operation** parameter, no registered consumers when using MQOP_START or MQOP_START_WAIT parameter, and trying to use non-threaded libraries with asynchronous API calls. parameter.

There is a conflict with the value specified for **Operation** parameter.

This error can be caused by an invalid value in the **Operation** parameter, no registered consumers when using MQOP_START or MQOP_START_WAIT parameter, and trying to use non-threaded libraries with asynchronous API calls.

Completion Code

MQCC_FAILED

Programmer Response

Investigate the application program and verify the **Operation** parameter options are correct. Ensure you have link edited the application with the correct version of the threading libraries for asynchronous functions.

2489 (09B9) (RC2489): MQRC_BMHO_ERROR:

Explanation

On an MQBFMH call, the buffer to message handle options structure MQBMHO is not valid, for one of the following reasons:

- The StruId field is not MQBMHO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQBMHO structure are set correctly.

2490 (09BA) (RC2490): MQRC_UNSUPPORTED_PROPERTY:**Explanation**

A message was found to contain a property that the queue manager does not support. The operation that failed required all the properties to be supported by the queue manager. This can occur on the MQPUT/MQPUT1 call or when a message is about to be sent down a channel to a queue manager that does not support message properties.

Completion Code

MQCC_FAILED

Programmer Response

Determine which property of the message is not supported by the queue manager and decide whether to remove the property from the message or connect to a queue manager which does support the property.

2492 (09BC) (RC2492): MQRC_PROP_NAME_NOT_CONVERTED:**Explanation**

An MQINQMP call was issued with the MQIMPO_CONVERT_VALUE option specified in the InqPropOpts parameter, but an error occurred during conversion of the returned name of the property. The returned name is unconverted.

Completion Code

MQCC_WARNING

Programmer Response

Check that the character set of the returned name was correctly described when the property was set. Also check that these values, and the RequestedCCSID and RequestedEncoding specified in the InqPropOpts parameter of the MQINQMP call, are supported for MQ conversion. If the required conversion is not supported, conversion must be carried out by the application.

2494 (09BE) (RC2494): MQRC_GET_ENABLED:**Explanation**

This reason code is returned to an asynchronous consumer at the time a queue that was previously inhibited for get has been re-enabled for get.

Completion Code

MQCC_WARNING

Programmer Response

None. This reason code is used to inform the application of the change in state of the queue.

2495 (09BF) (RC2495): MQRC_MODULE_NOT_FOUND:

Explanation

A native shared library could not be loaded.

Completion Code

MQCC_FAILED

Programmer Response

This problem could be caused by either of the two following reasons:

- A MQCB call was made with an Operation of MQOP_REGISTER specifying a *CallbackName* which could not be found. Ensure that the *CallbackName* value is specified correctly.
- The Java MQ code could not load a Java native shared library. Check the associated Exception stack and FFST. Ensure that the JNI shared library is specified correctly.

2496 (09C0) (RC2496): MQRC_MODULE_INVALID:

Explanation

An MQCB call was made with an Operation of MQOP_REGISTER, specifying a *CallbackName* which is not a valid load module.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the *CallbackName* value is specified correctly.

2497 (09C1) (RC2497): MQRC_MODULE_ENTRY_NOT_FOUND:

Explanation

An MQCB call was made with an Operation of MQOP_REGISTER and the *CallbackName* identifies a function name which can't be found in the specified library.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the *CallbackName* value is specified correctly.

2498 (09C2) (RC2498): MQRC_MIXED_CONTENT_NOT_ALLOWED:**Explanation**

An attempt was made to set a property with mixed content. For example, if an application set the property "x.y" and then attempted to set the property "x.y.z" it is unclear whether in the property name hierarchy "y" contains a value or another logical grouping. Such a hierarchy would be "mixed content" and this is not supported. Setting a property which would cause mixed content is not allowed. A hierarchy within a property name is created using the "." character (U+002E).

Completion Code

MQCC_FAILED

Programmer Response

Valid property names are described in the WebSphere MQ documentation. Change the property name hierarchy so that it no longer contains mixed content before re-issuing the call.

2499 (09C3) (RC2499): MQRC_MSG_HANDLE_IN_USE:**Explanation**

A message property call was called (MQCRTMH, MQDLTMH, MQSETMP, MQINQMP, MQDLTMP or MQMHBUF) specifying a message handle that is already in use on another API call. A message handle may only be used on one call at a time.

Concurrent use of a message handle can arise, for example, when an application uses multiple threads.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the message handle cannot be used while another call is in progress.

2500 (09C4) (RC2500): MQRC_HCONN_ASYNC_ACTIVE:**Explanation**

An attempt to issue an MQI call has been made while the connection is started.

Completion Code

MQCC_FAILED

Programmer Response

Stop or suspend the connection using the MQCTL call and retry the operation.

2501 (09C5) (RC2501): MQRC_MHBO_ERROR:

Explanation

On an MQMHBUF call, the message handle to buffer options structure MQMHBO is not valid, for one of the following reasons:

- The StrucId field is not MQMHBO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQMHBO structure are set correctly.

2502 (09C6) (RC2502): MQRC_PUBLICATION_FAILURE:

Explanation

An MQPUT or MQPUT1 call was issued to publish a message on a topic. Delivery of the publication to one of the subscribers failed and due to the combination of the syncpoint option used and either:

- • The PMSGDLV attribute on the administrative TOPIC object if it was a persistent message.
- • The NPMSGDLV attribute on the administrative TOPIC object if it was a non-persistent message.

The publication has not been delivered to any of the subscribers.

Completion Code

MQCC_FAILED

Programmer response

Find the subscriber or subscribers who are having problems with their subscription queue and resolve the problem, or change the setting of the PMSGDLV or NPMSGDLV attributes on the TOPIC so that problems with one subscriber do not have an effect on other subscribers. Retry the MQPUT.

2503 (09C7) (RC2503): MQRC_SUB_INHIBITED:

Explanation

MQSUB calls are currently inhibited for the topic subscribed to.

Completion Code

MQCC_FAILED

Programmer Response

If the system design allows subscription requests to be inhibited for short periods, retry the operation later.

2504 (09C8) (RC2504): MQRC_SELECTOR_ALWAYS_FALSE:**Explanation**

An MQOPEN, MQPUT1 or MQSUB call was issued but a selection string was specified which will never select a message

Completion Code

MQCC_FAILED

Programmer Response

Verify that the logic of the selection string which was passed in on the API is as expected. Make any necessary corrections to the logic of the string and resubmit the MQ API call for which the message occurred.

2507 (09CB) (RC2507): MQRC_XEPO_ERROR:**Explanation**

On an MQXEP call, the exit options structure MQXEPO is not valid, for one of the following reasons:

- The StrucId field is not MQXEPO_STRUC_ID.
- The Version field specifies a value that is not valid or not supported.
- The parameter pointer is not valid. (It is not always possible to detect parameter pointers that are not valid; if not detected, unpredictable results occur.)

Completion Code

MQCC_FAILED

Programmer Response

Ensure that input fields in the MQXEPO structure are set correctly.

2509 (09CD) (RC2509): MQRC_DURABILITY_NOT_ALTERABLE:**Explanation**

An MQSUB call using option MQSO_ALTER was made changing the durability of the subscription. The durability of a subscription cannot be changed.

Completion Code

MQCC_FAILED

Programmer Response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly, or change the durability option used on the MQSUB call so that it matches the existing subscription.

2510 (09CE) (RC2510): MQRC_TOPIC_NOT_ALTERABLE:**Explanation**

An MQSUB call using option MQSO_ALTER was made changing the one or more of the fields in the MQSD that provide the topic being subscribed to. These fields are the ObjectName, ObjectString, or wildcard options. The topic subscribed to cannot be changed.

Completion Code

MQCC_FAILED

Programmer Response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly, or change the attributes and options used on the MQSUB call so that it matches the existing subscription.

2512 (09D0) (RC2512): MQRC_SUBLEVEL_NOT_ALTERABLE:**Explanation**

An MQSUB call using option MQSO_ALTER was made changing the SubLevel of the subscription. The SubLevel of a subscription cannot be changed.

Completion Code

MQCC_FAILED

Programmer Response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly, or change the SubLevel field used on the MQSUB call so that it matches the existing subscription.

2513 (09D1) (RC2513): MQRC_PROPERTY_NAME_LENGTH_ERR:**Explanation**

An attempt was made to set, inquire or delete a property with an invalid name. This is for one of the following reasons:

- The VSLength field of the property name was set to less than or equal to zero.
- The VSLength field of the property name was set to greater than the maximum allowed value (see constant MQ_MAX_PROPERTY_NAME_LENGTH).
- The VSLength field of the property name was set to MQVS_NULL_TERMINATED and the property name was greater than the maximum allowed value.

Completion Code

MQCC_FAILED

Programmer Response

Valid property names are described in the WebSphere MQ documentation. Ensure that the property has a valid name length before issuing the call again.

2514 (09D2) (RC2514): MQRC_DUPLICATE_GROUP_SUB:**Explanation**

An MQSUB call using option MQSO_GROUP_SUB was made creating a new grouped subscription but, although it has a unique SubName, it matches the Full topic name of an existing subscription in the group.

Completion Code

MQCC_FAILED

Programmer Response

Correct the Full topic name used so that it does not match any existing subscription in the group, or correct the grouping attributes if, either a different group was intended or the subscription was not intended to be grouped at all.

2515 (09D3) (RC2515): MQRC_GROUPING_NOT_ALTERABLE:**Explanation**

An MQSUB call was made using option MQSO_ALTER on a grouped subscription, that is one made with the option MQSO_GROUP_SUB. Grouping of subscriptions is not alterable.

Completion Code

MQCC_FAILED

Programmer response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly, or change the various grouping fields used on the MQSUB call so that it matches the existing subscription.

2516 (09D4) (RC2516): MQRC_SELECTOR_INVALID_FOR_TYPE:**Explanation**

A SelectionString may only be specified in the MQOD for an MQOPEN/MQPUT1 if the following is true:

- ObjectType is MQOT_Q
- The queue is being opened using one of the MQOO_INPUT_* open options.

Completion Code

MQCC_FAILED

Programmer Response

Modify the value of ObjectType to be MQOT_Q and ensure that the queue is being opened using one of the MQOO_INPUT_* options.

2517 (09D5) (RC2517): MQRC_HOBJ QUIESCED:

Explanation

The HOBJ has been quiesced but there are no messages in the read ahead buffer which match the current selection criteria. This reason code indicates that the read ahead buffer is not empty.

Completion Code

MQCC_FAILED

Programmer Response

This reason code indicates that all messages with the current selection criteria have been processed. Do one of the following:

- If no further messages need to be processed issue an MQCLOSE without the MQCO_QUIESCE option. Any messages in the read ahead buffer will be discarded.
- Relax the current selection criteria by modifying the values in the MQGMO and reissue the call. Once all messages have been consumed the call will return MQRC_HOBJ_QUIESCED_NO_MSGS.

2518 (09D6) (RC2518): MQRC_HOBJ_QUIESCED_NO_MSGS:

Explanation

The HOBJ has been quiesced and the read ahead buffer is now empty. No further messages will be delivered to this HOBJ

Completion Code

MQCC_FAILED

Programmer Response

Issue MQCLOSE against the HOBJ.

2519 (09D7) (RC2519): MQRC_SELECTION_STRING_ERROR:

Explanation

The SelectionString must be specified according to the description of how to use an MQCHARV structure. Examples of why this error was returned:

- SelectionString.VSLength is greater than zero, but SelectionString.VSOffset is zero and SelectionString.VSPtr is a null pointer.
- SelectionString.VSOffset is nonzero and SelectionString.VSPtr is not the null pointer (that is, it appears both fields are being used where only one is allowed).
- SelectionString.VSPtr is not a valid pointer.
- SelectionString.VSOffset or SelectionString.VSPtr points to storage that is not accessible.
- SelectionString.VSLength exceeds the maximum length allowed for this field. The maximum length is determined by  MQ_SELECTOR_LENGTH (*WebSphere MQ V7.1 Reference*).

Completion Code

MQCC_FAILED

Programmer Response

Modify the fields of the MQCHARV so that it follows the rules for a valid MQCHARV structure.

2520 (09D8) (RC2520): MQRC_RES_OBJECT_STRING_ERROR:

Explanation

On the MQOPEN or MQPUT1 call in the Object Descriptor MQOD, or on the MQSUB call in the Subscription Descriptor MQSD the ResObjectString field is not valid.

One of the following applies:

- ResObjectString.VSLength is greater than zero, but ResObjectString.VSOffset is zero and ResObjectString.VSPtr is the null pointer.
- ResObjectString.VSOffset is nonzero and ResObjectString.VSPtr is not the null pointer (that is, it appears both fields are being used where only one is allowed).
- ResObjectString.VSPtr is not a valid pointer.
- ResObjectString.VSOffset or ResObjectString.VSPtr points to storage that is not accessible.
- ResObjectString.VSBufSize is MQVS_USE_VSLENGTH and one of ResObjectString.VSOffset or ResObjectString.VSPtr have been provided.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that one of ResObjectString.VSOffset or ResObjectString.VSPtr is zero and the other nonzero and that the buffer length is provided in ResObjectString.VSBufSize. Ensure that the field used points to accessible storage.

2521 (09D9) (RC2521): MQRC_CONNECTION_SUSPENDED:

Explanation

An MQCTL call with Operation MQOP_START_WAIT has returned because the asynchronous consumption of messages has been suspended. This can be for the following reasons:

- The connection was explicitly suspended using MQCTL with Operation MQOP_SUSPEND
- All consumers have been either unregistered or suspended.

Completion Code

MQCC_WARNING

Programmer Response

If this is an expected condition, no corrective action required. If this is an unexpected condition check that:

- At least one consumer is registered and not suspended
- The connection has not been suspended

2522 (09DA) (RC2522): MQRC_INVALID_DESTINATION:

Explanation

An MQSUB call failed because of a problem with the destination where publications messages are to be sent, so an object handle cannot be returned to the application and the subscription is not made. This can be for one of the following reasons:

- The MQSUB call used MQSO_CREATE, MQSO_MANAGED and MQSO_NON_DURABLE and the model queue referred to by MNDURMDL on the administrative topic node does not exist
- The MQSUB call used MQSO_CREATE, MQSO_MANAGED and MQSO_DURABLE and the model queue referred to by MDURMDL on the administrative topic node does not exist, or has been defined with a DEFTYPE of TEMPDYN.
- The MQSUB call used MQSO_CREATE or MQSO_ALTER on a durable subscription and the object handle provided referred to a temporary dynamic queue. This is not an appropriate destination for a durable subscription.
- The MQSUB call used MQSO_RESUME and a Hobj of MQHO_NONE, to resume an administratively created subscription, but the queue name provided in the DEST parameter of the subscription does not exist.
- The MQSUB call used MQSO_RESUME and a Hobj of MQHO_NONE, to resume a previously created API subscription, but the queue previously used no longer exists.

Completion Code

MQCC_FAILED

Programmer Response

Ensure that the model queues referred to by MNDURMDL and MDURMDL exist and have an appropriate DEFTYPE. Create the queue referred to by the DEST parameter in an administrative subscription if one is being used. Alter the subscription to use an existing queue if the previously used one does not exist.

2523 (09DB) (RC2523): MQRC_INVALID_SUBSCRIPTION:

Explanation

An MQSUB call using MQSO_RESUME or MQSO_ALTER failed because the subscription named is not valid for use by applications. This can be for one of the following reasons:

- The subscription is the SYSTEM.DEFAULT.SUB subscription which is not a valid subscription and should only be used to fill in the default values on DEFINE SUB commands.
- The subscription is a proxy type subscription which is not a valid subscription for an application to resume and is only used to enable publications to be forwarded between queue managers.
- The subscription has expired and is no longer valid for use.

Completion Code

MQCC_FAILED

Programmer Response

Ensure the subscription named in SubName field is not one of the invalid ones listed. If you have a handle open to the subscription already it must have expired. Use MQCLOSE to close the handle and then if necessary create a new subscription.

2524 (09DC) (RC2524): MQRC_SELECTOR_NOT_ALTERABLE:

Explanation

An MQSUB call was issued with the MQSO_ALTER option and the MQSD contained a SelectionString. It is not valid to alter the SelectionString of a subscription.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the SelectionString field of the MQSD does not contain a valid VSPtr and that the VSLength is set to zero when making a call to MQSUB.

2525 (09DD) (RC2525): MQRC_RETAINED_MSG_Q_ERROR:

Explanation

An MQSUB call which did not use the MQSO_NEW_PUBLICATIONS_ONLY option, or an MQSUBRQ call, failed because the retained publications which exist for the topic string subscribed to cannot be retrieved from the SYSTEM.RETAINED.PUB.QUEUE. This can be for one of the following reasons:

- The queue has become damaged or has been deleted.
- The queue has been set to GET(DISABLED).
- Messages have been removed from this queue directly.

An error message will be written to the log giving more details about the problem with the SYSTEM.RETAINED.PUB.QUEUE.

When this return code occurs on an MQSUB call, it can only occur using the MQSO_CREATE option, and in this case the subscription is not created.

Completion Code

MQCC_FAILED

Programmer Response

If this occurs on an MQSUB call, re-issue the MQSUB call using the option MQSO_NEW_PUBLICATIONS_ONLY, which will mean no previously retained publications are sent to this subscription, or fix the SYSTEM.RETAINED.PUB.QUEUE so that messages can be retrieved from it and re-issue the MQSUB call.

If this occurs on an MQSUBRQ call, fix the SYSTEM.RETAINED.PUB.QUEUE so that messages can be retrieved from it and re-issue the MQSUBRQ call.

2526 (09DE) (RC2526): MQRC_RETAINED_NOT_DELIVERED:

Explanation

An MQSUB call which did not use the MQSO_NEW_PUBLICATIONS_ONLY option or an MQSUBRQ call, failed because the retained publications which exist for the topic string subscribed to cannot be delivered to the subscription destination queue and have subsequently failed to be delivered to the dead-letter queue.

When this return code occurs on an MQSUB call, it can only occur using the MQSO_CREATE option, and in this case the subscription is not created.

Completion Code

MQCC_FAILED

Programmer response

Fix the problems with the destination queue and the dead-letter queue and re-issue the MQSUB or MQSUBRQ call.

2527 (09DF) (RC2527): MQRC_RFH_RESTRICTED_FORMAT_ERR:

Explanation

A message was put to a queue containing an MQRFH2 header which included a folder with a restricted format. However, the folder was not in the required format. These restrictions are:

- If NameValueCCSID of the folder is 1208 then only single byte UTF-8 characters are allowed in the folder, group or element names.
- Groups are not allowed in the folder.
- The values of properties may not contain any characters that require escaping.
- Only Unicode character U+0020 will be treated as white space within the folder.
- The folder tag does not contain the content attribute.
- The folder must not contain a property with a null value.

The <mq> folder requires formatting of this restricted form.

Completion Code

MQCC_FAILED

Programmer Response

Change the message to include valid MQRFH2 folders.

2528 (09E0) (RC2528): MQRC_CONNECTION_STOPPED:

Explanation

An MQCTL call was issued to start the asynchronous consumption of messages, but before the connection was ready to consume messages it was stopped by one of the message consumers.

Completion Code

MQCC_FAILED

Programmer Response

If this is an expected condition, no corrective action required. If this is an unexpected condition check whether an MQCTL with Operation MQOP_STOP was issued during the MQCBCT_START callback function.

2529 (09E1) (RC2529): MQRC_ASYNC_UOW_CONFLICT:

Explanation

An MQCTL call with Operation MQOP_START was issued to start the asynchronous consumption of messages, but the connection handle used already has a global unit of work outstanding. MQCTL cannot be used to start asynchronous consumption of messages while a unit of work is in existence unless the MQOP_START_WAIT Operation is used

Completion Code

MQCC_FAILED

Programmer Response

Issue an MQCMIT on the connection handle to commit the unit of work and then reissue the MQCTL call, or issue an MQCTL call using Operation MQOP_START_WAIT to use the unit of work from within the asynchronous consumption callback functions.

2530 (09E2) (RC2530): MQRC_ASYNC_XA_CONFLICT:

Explanation

An MQCTL call with Operation MQOP_START was issued to start the asynchronous consumption of messages, but an external XA syncpoint coordinator has already issued an xa_open call for this connection handle. XA transactions must be done using the MQOP_START_WAIT Operation.

Completion Code

MQCC_FAILED

Programmer Response

Reissue the MQCTL call using Operation MQOP_START_WAIT.

2531 (09E3) (RC2531): MQRC_PUBSUB_INHIBITED:

Explanation

MQSUB, MQOPEN, MQPUT, and MQPUT1 calls are currently inhibited for all publish/subscribe topics, either with the queue manager attribute PSMODE or because processing of publish/subscribe state at queue manager start-up has failed, or has not yet completed.

Completion Code

MQCC_FAILED

Programmer Response

If this queue manager does not intentionally inhibit publish/subscribe, investigate any error messages that describe the failure at queue manager start-up, or wait until start-up processing completes. If the queue manager is a member of cluster, then start-up is not complete until the channel initiator has also started. On z/OS, if you get this return code from the Chinit for the SYSTEM.BROKER.DEFAULT.STREAM queue or topic, then the Chinit is busy processing work, and the pubsub task starts later. Use the DISPLAY PUBSUB command to check the status of the publish/subscribe engine to ensure that it is ready for use. Additionally, on z/OS you might receive an information message CSQM076I.

2532 (09E4) (RC2532): MQRC_MSG_HANDLE_COPY_FAILURE:

Explanation

An MQGET call was issued specifying a valid MsgHandle in which to retrieve any properties of the message. After the message had been removed from the queue the application could not allocate enough storage for the properties of the message. The message data is available to the application but the properties are not. Check the queue manager error logs for more information about how much storage was required.

Completion Code

MQCC_WARNING

Programmer response

Raise the memory limit of the application to allow it store the properties.

2533 (09E5) (RC2533): MQRC_DEST_CLASS_NOT_ALTERABLE:

Explanation

An MQSUB call using option MQSO_ALTER was made changing the use of the MQSO_MANAGED option on the subscription. The destination class of a subscription cannot be changed. When the MQSO_MANAGED option is not used, the queue provided can be changed, but the class of destination (managed or not) cannot be changed.

Completion Code

MQCC_FAILED

Programmer Response

Remove the subscription using MQCLOSE and re-create it with MQSUB with the attributes set correctly, or change the use of the MQSO_MANAGED option used on the MQSUB call so that it matches the existing subscription.

2534 (09E6) (RC2534): MQRC_OPERATION_NOT_ALLOWED:

Explanation

An MQCTL call was made with an Operation that is not allowed because of the state of asynchronous consumption on the hConn is currently in.

If Operation was MQOP_RESUME, the operation is not allowed because the state of asynchronous consumption on the hConn is STOPPED. Re-issue MQCTL with the MQOP_START Operation.

If Operation was MQOP_SUSPEND, the operation is not allowed because the state of asynchronous consumption on the hConn is STOPPED. If you need to get your hConn into a SUSPENDED state, issue MQCTL with the MQOP_START Operation followed by MQCTL with MQOP_SUSPEND.

If Operation was MQOP_START, the operation is not allowed because the state of asynchronous consumption on the hConn is SUSPENDED. Re-issue MQCTL with the MQOP_RESUME Operation.

If Operation was MQOP_START_WAIT, the operation is not allowed because either:

- The state of asynchronous consumption on the hConn is SUSPENDED. Re-issue MQCTL with the MQOP_RESUME Operation.
- The state of asynchronous consumption on the hConn is already STARTED. Do not mix the use of MQOP_START and MQOP_START_WAIT within one application.

Completion Code

MQCC_FAILED

Programmer Response

Re-issue the MQCTL call with the correct Operation.

2535 (09E7): MQRC_ACTION_ERROR:

Explanation

An MQPUT call was issued, but the value of the Action field in the PutMsgOpts parameter is not a valid MQACTP_* value.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid value for the field.

2537 (09E9) (RC2537): MQRC_CHANNEL_NOT_AVAILABLE:

Explanation

An MQCONN call was issued from a client to connect to a queue manager but the channel is not currently available. Common causes of this reason code are:

- The channel is currently in stopped state.
- The channel has been stopped by a channel exit.
- The queue manager has reached its maximum allowable limit for this channel from this client.
- The queue manager has reached its maximum allowable limit for this channel.

- The queue manager has reached its maximum allowable limit for all channels.

Completion Code

MQCC_FAILED

Programmer Response

Examine the queue manager and client error logs for messages explaining the cause of the problem.

2538 (09EA) (RC2538): MQRC_HOST_NOT_AVAILABLE:

Explanation

An MQCONN call was issued from a client to connect to a queue manager but the attempt to allocate a conversation to the remote system failed. Common causes of this reason code are:

- The listener has not been started on the remote system.
- The connection name in the client channel definition is incorrect.
- The network is currently unavailable.
- A firewall blocking the port, or protocol-specific traffic.
- The security call initializing the WebSphere MQ client is blocked by a security exit on the SVRCONN channel at the server.

Completion Code

MQCC_FAILED

Programmer Response

Examine the client error log for messages explaining the cause of the problem.

If you are using a Linux server, and receiving a 2538 return code when trying to connect to a queue manager, ensure that you check your internal firewall configuration.

To diagnose the problem, issue the following commands to temporarily turn off the internal Linuxfirewall :

```
/etc/init.d/iptables save  
/etc/init.d/iptables stop
```

To turn the internal Linux firewall back on, issue the command:

```
/etc/init.d/iptables start
```

To permanently turn off the internal Linux firewall, issue the command:

```
chkconfig iptables off
```

2539 (09EB) (RC2539): MQRC_CHANNEL_CONFIG_ERROR:**Explanation**

An MQCONN call was issued from a client to connect to a queue manager but the attempt to establish communication failed. Common causes of this reason code are:

- The server and client cannot agree on the channel attributes to use.
- There are errors in one or both of the QM.INI or MQCLIENT.INI configuration files.
- The server machine does not support the code page used by the client.

Completion Code

MQCC_FAILED

Programmer Response

Examine the queue manager and client error logs for messages explaining the cause of the problem.

2540 (09EC) (RC2540): MQRC_UNKNOWN_CHANNEL_NAME:**Explanation**

An MQCONN call was issued from a client to connect to a queue manager but the attempt to establish communication failed because the queue manager did not recognize the channel name.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the client is configured to use the correct channel name.

2541 (09ED) (RC2541): MQRC_LOOPING_PUBLICATION:**Explanation**

A Distributed Pub/Sub topology has been configured with a combination of Pub/Sub clusters and Pub/Sub Hierarchies such that some, or all, of the queue managers have been connected in a loop. A looping publication has been detected and put onto the dead-letter queue.

Completion Code

MQCC_FAILED

Programmer response

Examine the hierarchy and correct the loop.

2543 (09EF) (RC2543): MQRC_STANDBY_Q_MGR:**Explanation**

The application attempted to connect to a standby queue manager instance.

Standby queue manager instances do not accept connections. To connect to the queue manager, you must connect to its active instance.

Completion Code

MQCC_FAILED

Programmer response

Connect the application to an active queue manager instance.

2544 (09F0) (RC2544): MQRC_RECONNECTING:**Explanation**

The connection has started reconnecting.

If an event handler has been registered with a reconnecting connection, it is called with this reason code when reconnection attempts begin.

Completion Code

MQCC_WARNING

Programmer response

Let WebSphere MQ continue with its next reconnection attempt, change the interval before the reconnection, or stop the reconnection. Change any application state that depends on the reconnection.

Note: Reconnection might start while the application is in the middle of an MQI call.

2545 (09F1) (RC2545): MQRC_RECONNECTED:**Explanation**

The connection reconnected successfully and all handles are reinstated.

If reconnection succeeds, an event handler registered with the connection is called with this reason code.

Completion Code

MQCC_OK

Programmer response

Set any application state that depends on the reconnection.

Note: Reconnection might finish while the application is in the middle of an MQI call.

2546 (09F2) (RC2546): MQRC_RECONNECT_QMID_MISMATCH:**Explanation**

A reconnectable connection specified MQCNO_RECONNECT_Q_MGR and the connection attempted to reconnect to a different queue manager.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the configuration for a reconnectable client resolves to a single queue manager.

If the application does not require reconnection to exactly the same queue manager, use the MQCONNX option MQCNO_RECONNECT.

2547 (09F3) (RC2547): MQRC_RECONNECT_INCOMPATIBLE:**Explanation**

An MQI option is incompatible with reconnectable connections.

This error indicates that the option relies on information in a queue manager that is lost during reconnection. For example, the option MQPMO_LOGICAL_ORDER, requires the queue manager to remember information about logical message ordering that is lost during reconnection.

Completion Code

MQCC_FAILED

Programmer response

Modify your application to remove the incompatible option, or do not allow the application to be reconnectable.

2548 (09F4) (RC2548): MQRC_RECONNECT_FAILED:**Explanation**

After reconnecting, an error occurred while reinstating the handles for a reconnectable connection.

For example, an attempt to reopen a queue that had been open when the connection broke, failed.

Completion Code

MQCC_FAILED

Programmer response

Investigate the cause of the error in the error logs. Consider using the MQSTAT API to find further details of the failure.

2549 (09F5) (RC2549): MQRC_CALL_INTERRUPTED:

Explanation

MQPUT, MQPUT1, or MQCMIT was interrupted and reconnection processing cannot reestablish a definite outcome.

This reason code is returned to a client that is using a reconnectable connection if the connection is broken between sending the request to the queue manager and receiving the response, and if the outcome is not certain. For example, an interrupted MQPUT of a persistent message outside sync point might or might not have stored the message. Alternatively an interrupted MQPUT1 of a persistent message or message with default persistence (which could be persistent) outside sync point might or might not have stored the message. The timing of the failure affects whether the message remains on the queue or not. If MQCMIT was interrupted the transaction might or might not have been committed.

Completion Code

MQCC_FAILED

Programmer response

Repeat the call following reconnection, but be aware that in some cases, repeating the call might be misleading.

The application design determines the appropriate recovery action. In many cases, getting and putting persistent messages inside sync point resolves indeterminate outcomes. Where persistent messages need to be processed outside sync point, it might be necessary to establish whether the interrupted operation succeeded before the interruption and repeating it if it did not.

2550 (09F6) (RC2550): MQRC_NO_SUBS_MATCHED:

Explanation

An MQPUT or MQPUT1 call was successful but no subscriptions matched the topic.

Completion Code

MQCC_WARNING

Programmer response

No response is required, unless this reason code was not expected by the application that put the message.

2551 (09F7) (RC2551): MQRC_SELECTION_NOT_AVAILABLE:

Explanation

An MQSUB call subscribed to publications using a SelectionString. WebSphere MQ is unable to accept the call because it does not follow the rules for specifying selection strings, which are documented in

 Message selector syntax (*WebSphere MQ V7.1 Programming Guide*). It is possible that the selection string is acceptable to an extended message selection provider, however no extended message selection provider was available to validate the selection string. If a subscription is being created, the MQSUB fails; otherwise MQSUB completes with a warning.

An MQPUT or MQPUT1 call published a message and at least one subscriber had a content filter but WebSphere MQ could not determine whether the publication should be delivered to the subscriber (for example, because no extended message selection provider was available to validate the selection string). The MQPUT or MQPUT1 call will fail with MQRC_SELECTION_NOT_AVAILABLE and no subscribers will receive the publication.

Completion Code

MQCC_WARNING or MQCC_FAILED

Programmer response

If it was intended that the selection string should be handled by the extended message selection provider, ensure that the extended message selection provider is correctly configured and running. If extended message selection was not intended, see  Message selector syntax (*WebSphere MQ V7.1 Programming Guide*) and ensure that you have correctly followed the rules for specifying selection strings.

If a subscription is being resumed, the subscription will not be delivered any messages until a extended message selection provider is available and a message matches the SelectionString of the resumed subscription.

2552 (09F8) (RC2552): MQRC_CHANNEL_SSL_WARNING:

Explanation

An SSL security event has occurred. This is not fatal to an SSL connection but is likely to be of interest to an administrator.

Completion code

MQCC_WARNING

Programmer response

None. This reason code is only used to identify the corresponding event message.

2553 (09F9) (RC2553): MQRC_OCSP_URL_ERROR:

Explanation

The OCSPResponderURL field does not contain a correctly formatted HTTP URL.

Completion code

MQCC_FAILED

Programmer response

Check and correct the OCSPResponderURL. If you do not intend to access an OCSP responder, set the AuthInfoType of the authentication information object to MQAIT_CRL_LDAP.

2554 (09FA) (RC2554): MQRC_CONTENT_ERROR:

Explanation

There are 2 explanations for reason code 2554:

1. An MQPUT call was issued with a message where the content could not be parsed to determine whether the message should be delivered to a subscriber with an extended message selector. No subscribers will receive the publication.
2. MQRC_CONTENT_ERROR can be returned from MQSUB and MQSUBRQ if a selection string selecting on the content of the message was specified.

Completion Code

MQCC_FAILED

Programmer response

There are 2 programmer responses for reason code 2554 because there are two causes:

1. If reason code 2554 was issued because of reason 1 then check for error messages from the extended message selection provider and ensure that the message content is well formed before retrying the operation.
2. If reason code 2554 was issued because of reason 2 then because the error occurred at the time that the retained message was published, either a system administrator must clear the retained queue, or you cannot specify a selection string selecting on the content.

2555 (09FB) (RC2555): MQRC_RECONNECT_Q_MGR_REQD:

Explanation

The MQCNO_RECONNECT_Q_MGR option is required.

An option, such MQMO_MATCH_MSG_TOKEN in an MQGET call or opening a durable subscription, was specified in the client program that requires re-connection to the same queue manager.

Completion Code

MQCC_FAILED

Programmer response

Change the MQCONNX call to use MQCNO_RECONNECT_Q_MGR, or modify the client program not to use the conflicting option.

2556 (09FC) (RC2556): MQRC_RECONNECT_TIMED_OUT:

Explanation

A reconnection attempt timed out.

The failure might occur in any MQI verb if a connection is configured to reconnect. You can customize the timeout in the MQClient.ini file

Completion Code

MQCC_FAILED

Programmer response

Look at the error logs to find out why reconnection did not complete within the time limit.

2557 (09FD) (RC2557): MQRC_PUBLISH_EXIT_ERROR:

Explanation

A publish exit function returned an invalid response code, or failed in some other way. This can be returned from the MQPUT, MQPUT1, MQSUB and MQSUBRQ function calls. This reason code does not occur on WebSphere MQ for z/OS.

Completion Code

MQCC_FAILED

Programmer response

Check the publish exit logic to ensure that the exit is returning valid values in the ExitResponse field of the  MQPSXP - Publish exit data structure (*WebSphere MQ V7.1 Reference*) structure. Consult the WebSphere MQ error log files and FFST records for more details about the problem.

2558 (09FE) (RC2558): MQRC_COMMINFO_ERROR:

Explanation

The configuration of either the name of the COMMINFO object or the object itself is incorrect.

Completion Code

MQCC_FAILED

Programmer response

Check the configuration of the TOPIC and COMMINFO objects and retry the operation.

2560 (0A00) (RC2560): MQRC_MULTICAST_ONLY:

Explanation

An attempt was made to use a topic which is defined as multicast only in a non-multicast way. Possible causes for this error are:

1. An MQPUT1 call was issued to the topic
2. An MQOPEN call was issued using the MQOO_NO_MULTICAST option
3. An MQSUB call was issued using the MQSO_NO_MULTICAST option
4. The application is connected directly through bindings, that is, there is no client connection
5. The application is being run from a release prior to version 7.1

Completion Code

MQCC_FAILED

Programmer response

Either change the topic definition to enable non-multicast, or change the application.

2561 (0A01) (RC2561): MQRC_DATA_SET_NOT_AVAILABLE:

Explanation

A WebSphere MQI call or command was issued to operate on a shared queue, but the call failed because the data for the shared message has been offloaded to a shared message data set that is temporarily unavailable to the current queue manager. This can occur either because of a problem in accessing the data set or because the data set was previously found to be damaged, and is awaiting completion of recovery processing.

This return code can also occur if the shared message data set has not been defined for the queue manager being used. You might be using the wrong queue manager in the queue-sharing group.

- This reason code occurs only on z/OS.

Completion Code

MQCC_FAILED

Programmer response

The problem is temporary; wait a short while, and then retry the operation.

Use `DIS CFSTRUCT(...)` `SMDSCONN(*)` to display the status of the SMDS connection.

To start the connection if the `STATUS` is not `OPEN`, use `STA SMDSCONN(*) CFSTRUCT(...)`.

Use `DISPLAY CFSTATUS(...)` `TYPE(SMDS)` and check the status is active on the queue manager that you are using.

2562 (0A02) (RC2562): MQRC_GROUPING_NOT_ALLOWED:

Explanation

An `MQPUT` call was issued to put a grouped message to a handle which is publishing over multicast.

Completion Code

MQCC_FAILED

Programmer response

Either change the topic definition to disable multicast or change the application to not use grouped messages.

2563 (0A03) (RC2563): MQRC_GROUP_ADDRESS_ERROR:**Explanation**

An MQOPEN or MQSUB call was issued to a multicast topic which has been defined with an incorrect group address field.

Completion Code

MQCC_FAILED

Programmer response

Correct the group address field in the COMMINFO definition linked to the TOPIC object.

2564 (0A04) (RC2564): MQRC_MULTICAST_CONFIG_ERROR:**Explanation**

An MQOPEN, MQSUB or MQPUT call was issued which invoked the multicast component. The call failed because the multicast configuration is incorrect.

Completion Code

MQCC_FAILED

Programmer response

Check the multicast configuration and error logs and retry the operation.

2565 (0A05) (RC2565): MQRC_MULTICAST_INTERFACE_ERROR:**Explanation**

An MQOPEN, MQSUB or MQPUT call was made which attempted to a network interface for multicast. The interface returned an error. Possible causes for the error are:

1. The required network interface does not exist.
2. The interface is not active.
3. The interface does not support the required IP version.

Completion Code

MQCC_FAILED

Programmer response

Verify that the IP address and the system network configuration are valid. Check the multicast configuration and error logs and retry the operation.

2566 (0A06) (RC2566): MQRC_MULTICAST_SEND_ERROR:**Explanation**

An MQPUT call was made which attempted to send multicast traffic over the network. The system failed to send one or more network packets.

Completion Code

MQCC_FAILED

Programmer response

Verify that the IP address and the system network configuration are valid. Check the multicast configuration and error logs and retry the operation.

2567 (0A07) (RC2567): MQRC_MULTICAST_INTERNAL_ERROR:**Explanation**

An MQOPEN, MQSUB or MQPUT call was issued which invoked the multicast component. An internal error occurred which prevented the operation completing successfully.

Completion Code

MQCC_FAILED

Programmer response

Tell the systems administrator.

2568 (0A08) (RC2568): MQRC_CONNECTION_NOT_AVAILABLE:**Explanation**

An MQCONN or MQCONNX call was made when the queue manager was unable to provide a connection of the requested connection type on the current installation. A client connection cannot be made on a server only installation. A local connection cannot be made on a client only installation.

This error can also occur when WebSphere MQ fails an attempt to load a library from the installation that the requested queue manager is associated with.

Completion Code

MQCC_FAILED

Programmer response

Ensure that the connection type requested is applicable to the type of installation. If the connection type is applicable to the installation then consult the error log for more information about the nature of the error.

2569 (0A09) (RC2569): MQRC_SYNCPOINT_NOT_ALLOWED:

Explanation

An MQPUT or MQPUT1 call using MQPMO_SYNCPOINT was made to a topic that is defined as MCAST(ENABLED). This is not allowed.

Completion Code

MQCC_FAILED

Programmer response

Change the application to use MQPMO_NO_SYNCPOINT, or alter the topic to disable the use of Multicast and retry the operation.

2583 (0A17) (RC2583): MQRC_INSTALLATION_MISMATCH:

Explanation

The application attempted to connect to a queue manager that is not associated with the same IBM WebSphere MQ installation as the loaded libraries.

Completion Code

MQCC_FAILED

Programmer response

An application must use the libraries from the installation the queue manager is associated with. If the *AMQ_SINGLE_INSTALLATION* environment variable is set, you must ensure that the application connects only to queue managers associated with a single installation. Otherwise, if WebSphere MQ is unable to automatically locate the correct libraries, you must modify the application, or the library search path, to ensure that the correct libraries are used.

2587 (0A1B) (RC2587): MQRC_HMSG_NOT_AVAILABLE:

Explanation

On an MQGET, MQPUT, or MQPUT1 call, a message handle supplied is not valid with the installation the queue manager is associated with. The message handle was created by MQCRTMH specifying the MQHC_UNASSOCIATED_HCONN option. It can be used only with queue managers associated with the first installation used in the process.

Completion Code

MQCC_FAILED

Programmer response

To pass properties between two queue managers associated with different installations, convert the message handle retrieved using MQGET into a buffer using the MQMHBUF call. Then pass that buffer into the MQPUT or MQPUT1 call of the other queue manager. Alternatively, use the **setmqm** command to associate one of the queue managers with the installation that the other queue manager is using. Using the **setmqm** command might change the version of WebSphere MQ that the queue manager uses.

2589 (0A1D) (RC2589) MQRC_INSTALLATION_MISSING:

Explanation

On an MQCONN or MQCONNX call, an attempt was made to connect to a queue manager where the associated installation is no longer installed.

Completion Code

MQCC_FAILED

Programmer response

Associate the queue manager with a different installation using the **setmqm** command before attempting to connect to the queue manager again.

2590 (0A1E) (RC2590): MQRC_FASTPATH_NOT_AVAILABLE:

Explanation

On an MQCONNX call, the MQCNO_FASTPATH_BINDING option was specified. However, a fastpath connection to the queue manager cannot be made. This issue can occur when a non-fastpath connection to a queue manager was made in the process before this MQCONNX call.

Completion Code

MQCC_FAILED

Programmer response

Either change all MQCONNX calls within the process to be fastpath, or use the *AMQ_SINGLE_INSTALLATION* environment variable to restrict connections to a single installation, allowing the queue manager to accept fastpath and non-fastpath connections from the same process, in any order.

2591 (0A1F) (RC2591): MQRC_CIPHER_SPEC_NOT_SUITE_B:

Explanation

A client application is configured for NSA Suite B compliant operation but the CipherSpec for the client connection channel is not permitted at the configured Suite B security level. This can occur for Suite B CipherSpecs which fall outside the currently configured security level, for example if ECDHE_ECDSA_AES_128_GCM_SHA256 (which is 128-bit Suite B) is used when only the 192-bit Suite B security level is configured

For more information about which CipherSpecs are Suite B compliant, refer to "Specifying CipherSpecs" on page 774.

Completion Code

MQCC_FAILED

Programmer response

Select an appropriate CipherSpec which is permitted at the configured Suite B security level.

2592 (0A20) (RC2592): MQRC_SUITE_B_ERROR:

Explanation

The configuration of Suite B is invalid. For example, an unrecognized value was specified in the **MQSUIEB** environment variable, the **EncryptionPolicySuiteB** SSL stanza setting or the MQSCO **EncryptionPolicySuiteB** field.

Completion Code

MQCC_FAILED

Programmer response

Determine the fault in the Suite B configuration and amend.

2593 (0A21)(RC2593): MQRC_CERT_VAL_POLICY_ERROR:

Explanation

The certificate validation policy configuration is invalid. An unrecognized or unsupported value was specified in the **MQCERTVPOL** environment variable, the **CertificateValPolicy** SSL stanza setting or the MQSCO **CertificateValPolicy** field.

Completion Code

MQCC_FAILED

Programmer response

Specify a valid certificate validation policy which is supported on the current platform.

6100 (17D4) (RC6100): MQRC_REOPEN_EXCL_INPUT_ERROR:

Explanation

An open object does not have the correct **ImqObject open options** and requires one or more additional options. An implicit reopen is required but closure has been prevented.

Closure has been prevented because the queue is open for exclusive input and closure might result in the queue being accessed by another process or thread, before the queue is reopened by the process or thread that presently has access.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Set the **open options** explicitly to cover all eventualities so that implicit reopening is not required.

6101 (17D5) (RC6101): MQRC_REOPEN_INQUIRE_ERROR:

Explanation

An open object does not have the correct `ImqObject` **open options** and requires one or more additional options. An implicit reopen is required but closure has been prevented.

Closure has been prevented because one or more characteristics of the object need to be checked dynamically prior to closure, and the **open options** do not already include `MQOO_INQUIRE`.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Set the **open options** explicitly to include `MQOO_INQUIRE`.

6102 (17D6) (RC6102): MQRC_REOPEN_SAVED_CONTEXT_ERR:

Explanation

An open object does not have the correct `ImqObject` **open options** and requires one or more additional options. An implicit reopen is required but closure has been prevented.

Closure has been prevented because the queue is open with `MQOO_SAVE_ALL_CONTEXT`, and a destructive get has been performed previously. This has caused retained state information to be associated with the open queue and this information would be destroyed by closure.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Set the **open options** explicitly to cover all eventualities so that implicit reopening is not required.

6103 (17D7) (RC6103): MQRC_REOPEN_TEMPORARY_Q_ERROR:

Explanation

An open object does not have the correct `ImqObject` **open options** and requires one or more additional options. An implicit reopen is required but closure has been prevented.

Closure has been prevented because the queue is a local queue of the definition type `MQQDT_TEMPORARY_DYNAMIC`, that would be destroyed by closure.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Set the **open options** explicitly to cover all eventualities so that implicit reopening is not required.

6104 (17D8) (RC6104): MQRC_ATTRIBUTE_LOCKED:

Explanation

An attempt has been made to change the value of an attribute of an object while that object is open, or, for an ImqQueueManager object, while that object is connected. Certain attributes cannot be changed in these circumstances. Close or disconnect the object (as appropriate) before changing the attribute value.

An object might have been connected, opened, or both unexpectedly and implicitly to perform an MQINQ call. Check the attribute cross-reference table in  C++ and MQI cross-reference (*WebSphere MQ V7.1 Reference*) to determine whether any of your method invocations result in an MQINQ call.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Include MQOO_INQUIRE in the ImqObject **open options** and set them earlier.

6105 (17D9) (RC6105): MQRC_CURSOR_NOT_VALID:

Explanation

The browse cursor for an open queue has been invalidated since it was last used by an implicit reopen.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Set the ImqObject **open options** explicitly to cover all eventualities so that implicit reopening is not required.

6106 (17DA) (RC6106): MQRC_ENCODING_ERROR:

Explanation

The encoding of the (next) message item needs to be MQENC_NATIVE for pasting.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6107 (17DB) (RC6107): MQRC_STRUC_ID_ERROR:

Explanation

The structure id for the (next) message item, which is derived from the 4 characters beginning at the data pointer, is either missing or is inconsistent with the class of object into which the item is being pasted.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6108 (17DC) (RC6108): MQRC_NULL_POINTER:

Explanation

A null pointer has been supplied where a nonnull pointer is either required or implied.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6109 (17DD) (RC6109): MQRC_NO_CONNECTION_REFERENCE:

Explanation

The **connection reference** is null. A connection to an ImqQueueManager object is required.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6110 (17DE) (RC6110): MQRC_NO_BUFFER:

Explanation

No buffer is available. For an ImqCache object, one cannot be allocated, denoting an internal inconsistency in the object state that should not occur.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6111 (17DF) (RC6111): MQRC_BINARY_DATA_LENGTH_ERROR:

Explanation

The length of the binary data is inconsistent with the length of the target attribute. Zero is a correct length for all attributes.

- The correct length for an **accounting token** is MQ_ACCOUNTING_TOKEN_LENGTH.
- The correct length for an **alternate security id** is MQ_SECURITY_ID_LENGTH.
- The correct length for a **correlation id** is MQ_CORREL_ID_LENGTH.
- The correct length for a **facility token** is MQ_FACILITY_LENGTH.
- The correct length for a **group id** is MQ_GROUP_ID_LENGTH.
- The correct length for a **message id** is MQ_MSG_ID_LENGTH.
- The correct length for an **instance id** is MQ_OBJECT_INSTANCE_ID_LENGTH.
- The correct length for a **transaction instance id** is MQ_TRAN_INSTANCE_ID_LENGTH.
- The correct length for a **message token** is MQ_MSG_TOKEN_LENGTH.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6112 (17E0) (RC6112): MQRC_BUFFER_NOT_AUTOMATIC:

Explanation

A user-defined (and managed) buffer cannot be resized. A user-defined buffer can only be replaced or withdrawn. A buffer must be automatic (system-managed) before it can be resized.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

6113 (17E1) (RC6113): MQRC_INSUFFICIENT_BUFFER:

Explanation

There is insufficient buffer space available after the data pointer to accommodate the request. This might be because the buffer cannot be resized.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6114 (17E2) (RC6114): MQRC_INSUFFICIENT_DATA:

Explanation

There is insufficient data after the data pointer to accommodate the request.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6115 (17E3) (RC6115): MQRC_DATA_TRUNCATED:

Explanation

Data has been truncated when copying from one buffer to another. This might be because the target buffer cannot be resized, or because there is a problem addressing one or other buffer, or because a buffer is being downsized with a smaller replacement.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6116 (17E4) (RC6116): MQRC_ZERO_LENGTH:

Explanation

A zero length has been supplied where a positive length is either required or implied.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6117 (17E5) (RC6117): MQRC_NEGATIVE_LENGTH:

Explanation

A negative length has been supplied where a zero or positive length is required.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6118 (17E6) (RC6118): MQRC_NEGATIVE_OFFSET:

Explanation

A negative offset has been supplied where a zero or positive offset is required.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6119 (17E7) (RC6119): MQRC_INCONSISTENT_FORMAT:

Explanation

The format of the (next) message item is inconsistent with the class of object into which the item is being pasted.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6120 (17E8) (RC6120): MQRC_INCONSISTENT_OBJECT_STATE:

Explanation

There is an inconsistency between this object, which is open, and the referenced ImqQueueManager object, which is not connected.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6121 (17E9) (RC6121): MQRC_CONTEXT_OBJECT_NOT_VALID:

Explanation

The ImqPutMessageOptions **context reference** does not reference a valid ImqQueue object. The object has been previously destroyed.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6122 (17EA) (RC6122): MQRC_CONTEXT_OPEN_ERROR:

Explanation

The `ImqPutMessageOptions` **context reference** references an `ImqQueue` object that could not be opened to establish a context. This might be because the `ImqQueue` object has inappropriate **open options**. Inspect the referenced object **reason code** to establish the cause.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6123 (17EB) (RC6123): MQRC_STRUC_LENGTH_ERROR:

Explanation

The length of a data structure is inconsistent with its content. For an `MQRMH`, the length is insufficient to contain the fixed fields and all offset data.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

6124 (17EC) (RC6124): MQRC_NOT_CONNECTED:

Explanation

A method failed because a required connection to a queue manager was not available, and a connection cannot be established implicitly because the `IMQ_IMPL_CONN` flag of the `ImqQueueManager` **behavior** class attribute is `FALSE`.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Establish a connection to a queue manager and retry.

6125 (17ED) (RC6125): MQRC_NOT_OPEN:

Explanation

A method failed because an object was not open, and opening cannot be accomplished implicitly because the `IMQ_IMPL_OPEN` flag of the `ImqObject` **behavior** class attribute is `FALSE`.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Open the object and retry.

6126 (17EE) (RC6126): MQRC_DISTRIBUTION_LIST_EMPTY:

Explanation

An ImqDistributionList failed to open because there are no ImqQueue objects referenced.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Establish at least one ImqQueue object in which the **distribution list reference** addresses the ImqDistributionList object, and retry.

6127 (17EF) (RC6127): MQRC_INCONSISTENT_OPEN_OPTIONS:

Explanation

A method failed because the object is open, and the ImqObject **open options** are inconsistent with the required operation. The object cannot be reopened implicitly because the IMQ_IMPL_OPEN flag of the ImqObject **behavior** class attribute is false.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Open the object with appropriate ImqObject **open options** and retry.

6128 (17FO) (RC6128): MQRC_WRONG_VERSION:

Explanation

A method failed because a version number specified or encountered is either incorrect or not supported.

For the ImqCICSBridgeHeader class, the problem is with the **version** attribute.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

If you are specifying a version number, use one that is supported by the class. If you are receiving message data from another program, ensure that both programs are using consistent and supported version numbers.

6129 (17F1) (RC6129): MQRC_REFERENCE_ERROR:

Explanation

An object reference is invalid.

There is a problem with the address of a referenced object. At the time of use, the address of the object is nonnull, but is invalid and cannot be used for its intended purpose.

This reason code occurs in the WebSphere MQ C++ environment.

Completion Code

MQCC_FAILED

Programmer response

Check that the referenced object is neither deleted nor out of scope, or remove the reference by supplying a null address value.

PCF reason codes

Reason codes might be returned by a broker in response to a command message in PCF format, depending on the parameters used in that message.

For more information about PCFs, see "Introduction to Programmable Command Formats" on page 6, "Automating administration tasks" on page 5, and "Using Programmable Command Formats" on page 7.

The following is a list of PCF reason codes, in numeric order, providing detailed information to help you understand them, including:

- An explanation of the circumstances that have caused the code to be raised
 - The associated completion code
 - Suggested programmer actions in response to the code
- "3001 (0BB9) (RC3001): MQRCCF_CFH_TYPE_ERROR" on page 1597
 - "3002 (0BBA) (RC3002): MQRCCF_CFH_LENGTH_ERROR" on page 1597
 - "3003 (0BBB) (RC3003): MQRCCF_CFH_VERSION_ERROR" on page 1598
 - "3004 (0BBC) (RC3004): MQRCCF_CFH_MSG_SEQ_NUMBER_ERR" on page 1598
 - "3005 (0BBD) (RC3005): MQRCCF_CFH_CONTROL_ERROR" on page 1598
 - "3006 (0BBE) (RC3006): MQRCCF_CFH_PARM_COUNT_ERROR" on page 1598
 - "3007 (0BBF) (RC3007): MQRCCF_CFH_COMMAND_ERROR" on page 1599
 - "3008 (0BC0) (RC3008): MQRCCF_COMMAND_FAILED" on page 1599
 - "3009 (0BC1) (RC3009): MQRCCF_CFIN_LENGTH_ERROR" on page 1599
 - "3010 (0BC2) (RC3010): MQRCCF_CFST_LENGTH_ERROR" on page 1599
 - "3011 (0BC3) (RC3011): MQRCCF_CFST_STRING_LENGTH_ERR" on page 1600
 - "3012 (0BC4) (RC3012): MQRCCF_FORCE_VALUE_ERROR" on page 1600
 - "3013 (0BC5) (RC3013): MQRCCF_STRUCTURE_TYPE_ERROR" on page 1600
 - "3014 (0BC6) (RC3014): MQRCCF_CFIN_PARM_ID_ERROR" on page 1600

"3015 (0BC7) (RC3015): MQRCCF_CFST_PARM_ID_ERROR" on page 1601
"3016 (0BC8) (RC3016): MQRCCF_MSG_LENGTH_ERROR" on page 1601
"3017 (0BC9) (RC3017): MQRCCF_CFIN_DUPLICATE_PARM" on page 1601
"3018 (0BCA) (RC3018): MQRCCF_CFST_DUPLICATE_PARM" on page 1601
"3019 (0BCB) (RC3019): MQRCCF_PARM_COUNT_TOO_SMALL" on page 1602
"3020 (0BCC) (RC3020): MQRCCF_PARM_COUNT_TOO_BIG" on page 1602
"3021 (0BCD) (RC3021): MQRCCF_Q_ALREADY_IN_CELL" on page 1602
"3022 (0BCE) (RC3022): MQRCCF_Q_TYPE_ERROR" on page 1602
"3023 (0BCF) (RC3023): MQRCCF_MD_FORMAT_ERROR" on page 1603
"3024 (0BD0) (RC3024): MQRCCF_CFSL_LENGTH_ERROR" on page 1603
"3025 (0BD1) (RC3025): MQRCCF_REPLACE_VALUE_ERROR" on page 1603
"3026 (0BD2) (RC3026): MQRCCF_CFIL_DUPLICATE_VALUE" on page 1603
"3027 (0BD3) (RC3027): MQRCCF_CFIL_COUNT_ERROR" on page 1604
"3028 (0BD4) (RC3028): MQRCCF_CFIL_LENGTH_ERROR" on page 1604
"3029 (0BD5) (RC3029): MQRCCF_MODE_VALUE_ERROR" on page 1604
"3029 (0BD5) (RC3029): MQRCCF_QUIESCE_VALUE_ERROR" on page 1604
"3030 (0BD6) (RC3030): MQRCCF_MSG_SEQ_NUMBER_ERROR" on page 1604
"3031 (0BD7) (RC3031): MQRCCF_PING_DATA_COUNT_ERROR" on page 1605
"3032 (0BD8) (RC3032): MQRCCF_PING_DATA_COMPARE_ERROR" on page 1605
"3033 (0BD9) (RC3033): MQRCCF_CFSL_PARM_ID_ERROR" on page 1605
"3034 (0BDA) (RC3034): MQRCCF_CHANNEL_TYPE_ERROR" on page 1605
"3035 (0BDB) (RC3035): MQRCCF_PARM_SEQUENCE_ERROR" on page 1606
"3036 (0BDC) (RC3036): MQRCCF_XMIT_PROTOCOL_TYPE_ERR" on page 1606
"3037 (0BDD) (RC3037): MQRCCF_BATCH_SIZE_ERROR" on page 1606
"3038 (0BDE) (RC3038): MQRCCF_DISC_INT_ERROR" on page 1606
"3039 (0BDF) (RC3039): MQRCCF_SHORT_RETRY_ERROR" on page 1607
"3040 (0BE0) (RC3040): MQRCCF_SHORT_TIMER_ERROR" on page 1607
"3041 (0BE1) (RC3041): MQRCCF_LONG_RETRY_ERROR" on page 1607
"3042 (0BE2) (RC3042): MQRCCF_LONG_TIMER_ERROR" on page 1607
"3043 (0BE3) (RC3043): MQRCCF_SEQ_NUMBER_WRAP_ERROR" on page 1608
"3044 (0BE4) (RC3044): MQRCCF_MAX_MSG_LENGTH_ERROR" on page 1608
"3045 (0BE5) (RC3045): MQRCCF_PUT_AUTH_ERROR" on page 1608
"3046 (0BE6) (RC3046): MQRCCF_PURGE_VALUE_ERROR" on page 1608
"3047 (0BE7) (RC3047): MQRCCF_CFIL_PARM_ID_ERROR" on page 1609
"3048 (0BE8) (RC3048): MQRCCF_MSG_TRUNCATED" on page 1609
"3049 (0BE9) (RC3049): MQRCCF_CCSID_ERROR" on page 1609
"3050 (0BEA) (RC3050): MQRCCF_ENCODING_ERROR" on page 1610
"3052 (0BEC) (RC3052): MQRCCF_DATA_CONV_VALUE_ERROR" on page 1610
"3053 (0BED) (RC3053): MQRCCF_INDOUBT_VALUE_ERROR" on page 1610
"3054 (0BEE) (RC3054): MQRCCF_ESCAPE_TYPE_ERROR" on page 1610
"3062 (0BF6) (RC3062): MQRCCF_CHANNEL_TABLE_ERROR" on page 1611
"3063 (0BF7) (RC3063): MQRCCF_MCA_TYPE_ERROR" on page 1611
"3064 (0BF8) (RC3064): MQRCCF_CHL_INST_TYPE_ERROR" on page 1611
"3065 (0BF9) (RC3065): MQRCCF_CHL_STATUS_NOT_FOUND" on page 1611
"3066 (0BEA) (RC3066): MQRCCF_CFSL_DUPLICATE_PARM" on page 1612

"3067 (0BFB) (RC3067): MQRCCF_CFSL_TOTAL_LENGTH_ERROR" on page 1612
"3068 (0BFC) (RC3068): MQRCCF_CFSL_COUNT_ERROR" on page 1612
"3069 (0BFD) (RC3069): MQRCCF_CFSL_STRING_LENGTH_ERR" on page 1612
"3070 (0BFE) (RC3070): MQRCCF_BROKER_DELETED" on page 1613
"3071 (0BFF) (RC3071): MQRCCF_STREAM_ERROR" on page 1613
"3072 (0C00) (RC3072): MQRCCF_TOPIC_ERROR" on page 1613
"3073 (0C01) (RC3073): MQRCCF_NOT_REGISTERED" on page 1613
"3074 (0C02) (RC3074): MQRCCF_Q_MGR_NAME_ERROR" on page 1614
"3075 (0C03) (RC3075): MQRCCF_INCORRECT_STREAM" on page 1614
"3076 (0C04) (RC3076): MQRCCF_Q_NAME_ERROR" on page 1614
"3077 (0C05) (RC3077): MQRCCF_NO_RETAINED_MSG" on page 1615
"3078 (0C06) (RC3078): MQRCCF_DUPLICATE_IDENTITY" on page 1615
"3079 (0C07) (RC3079): MQRCCF_INCORRECT_Q" on page 1615
"3080 (0C08) (RC3080): MQRCCF_CORREL_ID_ERROR" on page 1616
"3081 (0C09) (RC3081): MQRCCF_NOT_AUTHORIZED" on page 1616
"3082 (0C0A) (RC3082): MQRCCF_UNKNOWN_STREAM" on page 1616
"3083 (0C0B) (RC3083): MQRCCF_REG_OPTIONS_ERROR" on page 1617
"3084 (0C0C) (RC3084): MQRCCF_PUB_OPTIONS_ERROR" on page 1617
"3085 (0C0D) (RC3085): MQRCCF_UNKNOWN_BROKER" on page 1617
"3086 (0C0E) (RC3086): MQRCCF_Q_MGR_CCSD_ERROR" on page 1617
"3087 (0C0F) (RC3087): MQRCCF_DEL_OPTIONS_ERROR" on page 1618
"3088 (0C10) (RC3088): MQRCCF_CLUSTER_NAME_CONFLICT" on page 1618
"3089 (0C11) (RC3089): MQRCCF_REPOS_NAME_CONFLICT" on page 1618
"3090 (0C12) (RC3090): MQRCCF_CLUSTER_Q_USAGE_ERROR" on page 1618
"3091 (0C13) (RC3091): MQRCCF_ACTION_VALUE_ERROR" on page 1619
"3092 (0C14) (RC3092): MQRCCF_COMMS_LIBRARY_ERROR" on page 1619
"3093 (0C15) (RC3093): MQRCCF_NETBIOS_NAME_ERROR" on page 1619
"3094 (0C16) (RC3094): MQRCCF_BROKER_COMMAND_FAILED" on page 1620
"3095 (0C17) (RC3095): MQRCCF_CFST_CONFLICTING_PARM" on page 1620
"3096 (0C18) (RC3096): MQRCCF_PATH_NOT_VALID" on page 1620
"3097 (0C19) (RC3097): MQRCCF_PARM_SYNTAX_ERROR" on page 1620
"3098 (0C1A) (RC3098): MQRCCF_PWD_LENGTH_ERROR" on page 1621
"3150 (0C4E) (RC3150): MQRCCF_FILTER_ERROR" on page 1621
"3151 (0C4F) (RC3151): MQRCCF_WRONG_USER" on page 1621
"3152 (0C50) (RC3152): MQRCCF_DUPLICATE_SUBSCRIPTION" on page 1621
"3153 (0C51) (RC3153): MQRCCF_SUB_NAME_ERROR" on page 1622
"3154 (0C52) (RC3154): MQRCCF_SUB_IDENTITY_ERROR" on page 1622
"3155 (0C53) (RC3155): MQRCCF_SUBSCRIPTION_IN_USE" on page 1622
"3156 (0C54) (RC3156): MQRCCF_SUBSCRIPTION_LOCKED" on page 1622
"3157 (0C55) (RC3157): MQRCCF_ALREADY_JOINED" on page 1623
"3160 (0C58) (RC3160): MQRCCF_OBJECT_IN_USE" on page 1623
"3161 (0C59) (RC3161): MQRCCF_UNKNOWN_FILE_NAME" on page 1623
"3162 (0C5A) (RC3162): MQRCCF_FILE_NOT_AVAILABLE" on page 1623
"3163 (0C5B) (RC3163): MQRCCF_DISC_RETRY_ERROR" on page 1624
"3164 (0C5C) (RC3164): MQRCCF_ALLOC_RETRY_ERROR" on page 1624

"3165 (0C5D) (RC3165): MQRCCF_ALLOC_SLOW_TIMER_ERROR" on page 1624
"3166 (0C5E) (RC3166): MQRCCF_ALLOC_FAST_TIMER_ERROR" on page 1624
"3167 (0C5F) (RC3167): MQRCCF_PORT_NUMBER_ERROR" on page 1625
"3168 (0C60) (RC3168): MQRCCF_CHL_SYSTEM_NOT_ACTIVE" on page 1625
"3169 (0C61) (RC3169): MQRCCF_ENTITY_NAME_MISSING" on page 1625
"3170 (0C62) (RC3170): MQRCCF_PROFILE_NAME_ERROR" on page 1625
"3171 (0C63) (RC3171): MQRCCF_AUTH_VALUE_ERROR" on page 1626
"3172 (0C64) (RC3172): MQRCCF_AUTH_VALUE_MISSING" on page 1626
"3173 (0C65) (RC3173): MQRCCF_OBJECT_TYPE_MISSING" on page 1626
"3174 (0C66) (RC3174): MQRCCF_CONNECTION_ID_ERROR" on page 1626
"3175 (0C67) (RC3175): MQRCCF_LOG_TYPE_ERROR" on page 1627
"3176 (0C68) (RC3176): MQRCCF_PROGRAM_NOT_AVAILABLE" on page 1627
"3177 (0C69) (RC3177): MQRCCF_PROGRAM_AUTH_FAILED" on page 1627
"3200 (0C80) (RC3200): MQRCCF_NONE_FOUND" on page 1627
"3201 (0C81) (RC3201): MQRCCF_SECURITY_SWITCH_OFF" on page 1628
"3202 (0C82) (RC3202): MQRCCF_SECURITY_REFRESH_FAILED" on page 1628
"3203 (0C83) (RC3203): MQRCCF_PARM_CONFLICT" on page 1628
"3204 (0C84) (RC3204): MQRCCF_COMMAND_INHIBITED" on page 1629
"3205 (0C85) (RC3205): MQRCCF_OBJECT_BEING_DELETED" on page 1629
"3207 (0C87) (RC3207): MQRCCF_STORAGE_CLASS_IN_USE" on page 1629
"3208 (0C88) (RC3208): MQRCCF_OBJECT_NAME_RESTRICTED" on page 1629
"3209 (0C89) (RC3209): MQRCCF_OBJECT_LIMIT_EXCEEDED" on page 1629
"3210 (0C8A) (RC3210): MQRCCF_OBJECT_OPEN_FORCE" on page 1630
"3211 (0C8B) (RC3211): MQRCCF_DISPOSITION_CONFLICT" on page 1630
"3212 (0C8C) (RC3212): MQRCCF_Q_MGR_NOT_IN_QSG" on page 1630
"3213 (0C8D) (RC3213): MQRCCF_ATTR_VALUE_FIXED" on page 1631
"3215 (0C8F) (RC3215): MQRCCF_NAMELIST_ERROR" on page 1631
"3217 (0C91) (RC3217): MQRCCF_NO_CHANNEL_INITIATOR" on page 1631
"3218 (0C93) (RC3218): MQRCCF_CHANNEL_INITIATOR_ERROR" on page 1631
"3222 (0C96) (RC3222): MQRCCF_COMMAND_LEVEL_CONFLICT" on page 1632
"3223 (0C97) (RC3223): MQRCCF_Q_ATTR_CONFLICT" on page 1632
"3224 (0C98) (RC3224): MQRCCF_EVENTS_DISABLED" on page 1632
"3225 (0C99) (RC3225): MQRCCF_COMMAND_SCOPE_ERROR" on page 1632
"3226 (0C9A) (RC3226): MQRCCF_COMMAND_REPLY_ERROR" on page 1632
"3227 (0C9B) (RC3227): MQRCCF_FUNCTION_RESTRICTED" on page 1633
"3228 (0C9C) (RC3228): MQRCCF_PARM_MISSING" on page 1633
"3229 (0C9D) (RC3229): MQRCCF_PARM_VALUE_ERROR" on page 1633
"3230 (0C9E) (RC3230): MQRCCF_COMMAND_LENGTH_ERROR" on page 1634
"3231 (0C9F) (RC3231): MQRCCF_COMMAND_ORIGIN_ERROR" on page 1634
"3232 (0CA0) (RC3232): MQRCCF_LISTENER_CONFLICT" on page 1634
"3233 (0CA1) (RC3233): MQRCCF_LISTENER_STARTED" on page 1634
"3234 (0CA2) (RC3234): MQRCCF_LISTENER_STOPPED" on page 1635
"3235 (0CA3) (RC3235): MQRCCF_CHANNEL_ERROR" on page 1635
"3236 (0CA4) (RC3236): MQRCCF_CF_STRUC_ERROR" on page 1635
"3237 (0CA5) (RC3237): MQRCCF_UNKNOWN_USER_ID" on page 1636

"3238 (0CA6) (RC3238): MQRCCF_UNEXPECTED_ERROR" on page 1636
"3239 (0CA7) (RC3239): MQRCCF_NO_XCF_PARTNER" on page 1636
"3240 (0CA8) (RC3240): MQRCCF_CFGR_PARM_ID_ERROR" on page 1636
"3241 (0CA9) (RC3241): MQRCCF_CFIF_LENGTH_ERROR" on page 1636
"3242 (0CAA) (RC3242): MQRCCF_CFIF_OPERATOR_ERROR" on page 1637
"3243 (0CAB) (RC3243): MQRCCF_CFIF_PARM_ID_ERROR" on page 1637
"3244 (0CAC) (RC3244): MQRCCF_CFSF_FILTER_VAL_LEN_ERR" on page 1637
"3245 (0CAD) (RC3245): MQRCCF_CFSF_LENGTH_ERROR" on page 1637
"3246 (0CAE) (RC3246): MQRCCF_CFSF_OPERATOR_ERROR" on page 1638
"3247 (0CAF) (RC3247): MQRCCF_CFSF_PARM_ID_ERROR" on page 1638
"3248 (0CB0) (RC3248): MQRCCF_TOO_MANY_FILTERS" on page 1638
"3249 (0CB1) (RC3249): MQRCCF_LISTENER_RUNNING" on page 1638
"3250 (0CB2) (RC3250): MQRCCF_LSTR_STATUS_NOT_FOUND" on page 1639
"3251 (0CB3) (RC3251): MQRCCF_SERVICE_RUNNING" on page 1639
"3252 (0CB4) (RC3252): MQRCCF_SERV_STATUS_NOT_FOUND" on page 1639
"3253 (0CB5) (RC3253): MQRCCF_SERVICE_STOPPED" on page 1639
"3254 (0CB6) (RC3254): MQRCCF_CFBS_DUPLICATE_PARM" on page 1639
"3255 (0CB7) (RC3255): MQRCCF_CFBS_LENGTH_ERROR" on page 1640
"3256 (0CB8) (RC3256): MQRCCF_CFBS_PARM_ID_ERROR" on page 1640
"3257 (0CB9) (RC3257): MQRCCF_CFBS_STRING_LENGTH_ERR" on page 1640
"3258 (0CBA) (RC3258): MQRCCF_CFGR_LENGTH_ERROR" on page 1640
"3259 (0CBB) (RC3259): MQRCCF_CFGR_PARM_COUNT_ERROR" on page 1641
"3260 (0CBC) (RC3260): MQRCCF_CONN_NOT_STOPPED" on page 1641
"3261 (0CBD) (RC3261): MQRCCF_SERVICE_REQUEST_PENDING" on page 1641
"3262 (0CBE) (RC3262): MQRCCF_NO_START_CMD" on page 1641
"3263 (0CBF) (RC3263): MQRCCF_NO_STOP_CMD" on page 1641
"3264 (0CC0) (RC3264): MQRCCF_CFBF_LENGTH_ERROR" on page 1642
"3265 (0CC1) (RC3265): MQRCCF_CFBF_PARM_ID_ERROR" on page 1642
"3266 (0CC2) (RC3266): MQRCCF_CFBF_FILTER_VAL_LEN_ERR" on page 1642
"3267 (0CC3) (RC3267): MQRCCF_CFBF_OPERATOR_ERROR" on page 1642
"3268 (0CC4) (RC3268): MQRCCF_LISTENER_STILL_ACTIVE" on page 1643
"3269 (0CC5) (RC3269): MQRCCF_DEF_XMIT_Q_CLUS_ERROR" on page 1643
"3300 (0CE4) (RC3300): MQRCCF_TOPICSTR_ALREADY_EXISTS" on page 1643
"3301 (0CE5) (RC3301): MQRCCF_SHARING_CONVS_ERROR" on page 1643
"3302 (0CE6) (RC3302): MQRCCF_SHARING_CONVS_TYPE" on page 1643
"3303 (0CE7) (RC3303): MQRCCF_SECURITY_CASE_CONFLICT" on page 1644
"3305 (0CE9) (RC3305): MQRCCF_TOPIC_TYPE_ERROR" on page 1644
"3306 (0CEA) (RC3306): MQRCCF_MAX_INSTANCES_ERROR" on page 1644
"3307 (0CEB) (RC3307): MQRCCF_MAX_INSTS_PER_CLNT_ERR" on page 1644
"3308 (0CEC) (RC3308): MQRCCF_TOPIC_STRING_NOT_FOUND" on page 1644
"3309 (0CED) (RC3309): MQRCCF_SUBSCRIPTION_POINT_ERR" on page 1645
"3311 (0CEF) (RC2432): MQRCCF_SUB_ALREADY_EXISTS" on page 1645
"3317 (0CF5) (RC3317): MQRCCF_INVALID_DESTINATION" on page 1645
"3318 (0CF6) (RC3318): MQRCCF_PUBSUB_INHIBITED" on page 1645
"3326 (0CFE) (RC3326): MQRCCF_CHLAUTH_TYPE_ERROR" on page 1646

"3327 (0CFF) (RC3327): MQRCCF_CHLAUTH_ACTION_ERROR" on page 1646
"3335 (0D07) (RC3335): MQRCCF_CHLAUTH_USRSRC_ERROR" on page 1646
"3336 (0D08) (RC3336): MQRCCF_WRONG_CHLAUTH_TYPE" on page 1646
"3337 (0D09) (RC3337): MQRCCF_CHLAUTH_ALREADY_EXISTS" on page 1646
"3338 (0D0A) (RC3338): MQRCCF_CHLAUTH_NOT_FOUND" on page 1647
"3339 (0D0B) (RC3339): MQRCCF_WRONG_CHLAUTH_ACTION" on page 1647
"3340 (0D0C) (RC3340): MQRCCF_WRONG_CHLAUTH_USERSRC" on page 1647
"3341 (0D0D) (RC3341): MQRCCF_CHLAUTH_WARN_ERROR" on page 1647
"3342 (0D0E) (RC3342): MQRCCF_WRONG_CHLAUTH_MATCH" on page 1647
"3343 (0D0F) (RC3343): MQRCCF_IPADDR_RANGE_CONFLICT" on page 1648
"3344 (0D10) (RC3344): MQRCCF_CHLAUTH_MAX_EXCEEDED" on page 1648
"3345 (0D11) (RC3345): MQRCCF_IPADDR_ERROR" on page 1648
"3346 (0D12) (RC3346): MQRCCF_IPADDR_RANGE_ERROR" on page 1648
"3347 (0D13) (RC3347): MQRCCF_PROFILE_NAME_MISSING" on page 1649
"3348 (0D14) (RC3348): MQRCCF_CHLAUTH_CLNTUSER_ERROR" on page 1649
"3349 (0D15) (RC3349): MQRCCF_CHLAUTH_NAME_ERROR" on page 1649
"3353 (0D19) (RC3353): MQRCCF_SUITE_B_ERROR" on page 1650
"3364 (0D24) (RC3364): MQRCCF_CERT_VAL_POLICY_ERROR" on page 1650
"4001 (0FA1) (RC4001): MQRCCF_OBJECT_ALREADY_EXISTS" on page 1650
"4002 (0FA2) (RC4002): MQRCCF_OBJECT_WRONG_TYPE" on page 1650
"4003 (0FA3) (RC4003): MQRCCF_LIKE_OBJECT_WRONG_TYPE" on page 1651
"4004 (0FA4) (RC4004): MQRCCF_OBJECT_OPEN" on page 1651
"4005 (0FA5) (RC4005): MQRCCF_ATTR_VALUE_ERROR" on page 1651
"4006 (0FA6) (RC4006): MQRCCF_UNKNOWN_Q_MGR" on page 1651
"4007 (0FA7) (RC4007): MQRCCF_Q_WRONG_TYPE" on page 1652
"4008 (0FA8) (RC4008): MQRCCF_OBJECT_NAME_ERROR" on page 1652
"4009 (0FA9) (RC4009): MQRCCF_ALLOCATE_FAILED" on page 1652
"4010 (0FAA) (RC4010): MQRCCF_HOST_NOT_AVAILABLE" on page 1652
"4011 (0FAB) (RC4011): MQRCCF_CONFIGURATION_ERROR" on page 1653
"4012 (0FAC) (RC4012): MQRCCF_CONNECTION_REFUSED" on page 1653
"4013 (0FAD) (RC4013): MQRCCF_ENTRY_ERROR" on page 1653
"4014 (0FAE) (RC4014): MQRCCF_SEND_FAILED" on page 1654
"4015 (0FAF) (RC4015): MQRCCF_RECEIVED_DATA_ERROR" on page 1654
"4016 (0FB0) (RC4016): MQRCCF_RECEIVE_FAILED" on page 1654
"4017 (0FB1) (RC4017): MQRCCF_CONNECTION_CLOSED" on page 1654
"4018 (0FB2) (RC4018): MQRCCF_NO_STORAGE" on page 1655
"4019 (0FB3) (RC4019): MQRCCF_NO_COMMS_MANAGER" on page 1655
"4020 (0FB4) (RC4020): MQRCCF_LISTENER_NOT_STARTED" on page 1655
"4024 (0FB8) (RC4024): MQRCCF_BIND_FAILED" on page 1655
"4025 (0FB9) (RC4025): MQRCCF_CHANNEL_INDOUBT" on page 1656
"4026 (0FBA) (RC4026): MQRCCF_MQCONN_FAILED" on page 1656
"4027 (0FBB) (RC4027): MQRCCF_MQOPEN_FAILED" on page 1656
"4028 (0FBC) (RC4028): MQRCCF_MQGET_FAILED" on page 1656
"4029 (0FBD) (RC4029): MQRCCF_MQPUT_FAILED" on page 1656
"4030 (0FBE) (RC4030): MQRCCF_PING_ERROR" on page 1657

"4031 (0FBF) (RC4031): MQRCCF_CHANNEL_IN_USE" on page 1657
"4032 (0FC0) (RC4032): MQRCCF_CHANNEL_NOT_FOUND" on page 1657
"4033 (0FC1) (RC4033): MQRCCF_UNKNOWN_REMOTE_CHANNEL" on page 1657
"4034 (0FC2) (RC4034): MQRCCF_REMOTE_QM_UNAVAILABLE" on page 1658
"4035 (0FC3) (RC4035): MQRCCF_REMOTE_QM_TERMINATING" on page 1658
"4036 (0FC4) (RC4036): MQRCCF_MQINQ_FAILED" on page 1658
"4037 (0FC5) (RC4037): MQRCCF_NOT_XMIT_Q" on page 1658
"4038 (0FC6) (RC4038): MQRCCF_CHANNEL_DISABLED" on page 1659
"4039 (0FC7) (RC4039): MQRCCF_USER_EXIT_NOT_AVAILABLE" on page 1659
"4040 (0FC8) (RC4040): MQRCCF_COMMIT_FAILED" on page 1659
"4041 (0FC9) (RC4041): MQRCCF_WRONG_CHANNEL_TYPE" on page 1659
"4042 (0FCA) (RC4042): MQRCCF_CHANNEL_ALREADY_EXISTS" on page 1660
"4043 (0FCB) (RC4043): MQRCCF_DATA_TOO_LARGE" on page 1660
"4044 (0FCC) (RC4044): MQRCCF_CHANNEL_NAME_ERROR" on page 1660
"4045 (0FCD) (RC4045): MQRCCF_XMIT_Q_NAME_ERROR" on page 1660
"4047 (0FCF) (RC4047): MQRCCF_MCA_NAME_ERROR" on page 1661
"4048 (0FD0) (RC4048): MQRCCF_SEND_EXIT_NAME_ERROR" on page 1661
"4049 (0FD1) (RC4049): MQRCCF_SEC_EXIT_NAME_ERROR" on page 1661
"4050 (0FD2) (RC4050): MQRCCF_MSG_EXIT_NAME_ERROR" on page 1661
"4051 (0FD3) (RC4051): MQRCCF_RCV_EXIT_NAME_ERROR" on page 1662
"4052 (0FD4) (RC4052): MQRCCF_XMIT_Q_NAME_WRONG_TYPE" on page 1662
"4053 (0FD5) (RC4053): MQRCCF_MCA_NAME_WRONG_TYPE" on page 1662
"4054 (0FD6) (RC4054): MQRCCF_DISC_INT_WRONG_TYPE" on page 1662
"4055 (0FD7) (RC4055): MQRCCF_SHORT_RETRY_WRONG_TYPE" on page 1663
"4056 (0FD8) (RC4056): MQRCCF_SHORT_TIMER_WRONG_TYPE" on page 1663
"4057 (0FD9) (RC4057): MQRCCF_LONG_RETRY_WRONG_TYPE" on page 1663
"4058 (0FDA) (RC4058): MQRCCF_LONG_TIMER_WRONG_TYPE" on page 1663
"4059 (0FDB) (RC4059): MQRCCF_PUT_AUTH_WRONG_TYPE" on page 1664
"4061 (0FDD) (RC4061): MQRCCF_MISSING_CONN_NAME" on page 1664
"4062 (0FDE) (RC4062): MQRCCF_CONN_NAME_ERROR" on page 1664
"4063 (0FDF) (RC4063): MQRCCF_MQSET_FAILED" on page 1664
"4064 (0FE0) (RC4064): MQRCCF_CHANNEL_NOT_ACTIVE" on page 1664
"4065 (0FE1) (RC4065): MQRCCF_TERMINATED_BY_SEC_EXIT" on page 1665
"4067 (0FE3) (RC4067): MQRCCF_DYNAMIC_Q_SCOPE_ERROR" on page 1665
"4068 (0FE4) (RC4068): MQRCCF_CELL_DIR_NOT_AVAILABLE" on page 1665
"4069 (0FE5) (RC4069): MQRCCF_MR_COUNT_ERROR" on page 1665
"4070 (0FE6) (RC4070): MQRCCF_MR_COUNT_WRONG_TYPE" on page 1666
"4071 (0FE7) (RC4071): MQRCCF_MR_EXIT_NAME_ERROR" on page 1666
"4072 (0FE8) (RC4072): MQRCCF_MR_EXIT_NAME_WRONG_TYPE" on page 1666
"4073 (0FE9) (RC4073): MQRCCF_MR_INTERVAL_ERROR" on page 1666
"4074 (0FEA) (RC4074): MQRCCF_MR_INTERVAL_WRONG_TYPE" on page 1667
"4075 (0FEB) (RC4075): MQRCCF_NPM_SPEED_ERROR" on page 1667
"4076 (0FEC) (RC4076): MQRCCF_NPM_SPEED_WRONG_TYPE" on page 1667
"4077 (0FED) (RC4077): MQRCCF_HB_INTERVAL_ERROR" on page 1667
"4078 (0FEE) (RC4078): MQRCCF_HB_INTERVAL_WRONG_TYPE" on page 1668

"4079 (0FEF) (RC4079): MQRCCF_CHAD_ERROR" on page 1668
"4080 (0FF0) (RC4080): MQRCCF_CHAD_WRONG_TYPE" on page 1668
"4081 (0FF1) (RC4081): MQRCCF_CHAD_EVENT_ERROR" on page 1668
"4082 (0FF2) (RC4082): MQRCCF_CHAD_EVENT_WRONG_TYPE" on page 1669
"4083 (0FF3) (RC4083): MQRCCF_CHAD_EXIT_ERROR" on page 1669
"4084 (0FF4) (RC4084): MQRCCF_CHAD_EXIT_WRONG_TYPE" on page 1669
"4085 (0FF5) (RC4085): MQRCCF_SUPPRESSED_BY_EXIT" on page 1669
"4086 (0FF6) (RC4086): MQRCCF_BATCH_INT_ERROR" on page 1670
"4087 (0FF7) (RC4087): MQRCCF_BATCH_INT_WRONG_TYPE" on page 1670
"4088 (0FF8) (RC4088): MQRCCF_NET_PRIORITY_ERROR" on page 1670
"4089 (0FF9) (RC4089): MQRCCF_NET_PRIORITY_WRONG_TYPE" on page 1670
"4090 (0FFA) (RC4090): MQRCCF_CHANNEL_CLOSED" on page 1670
"4092 (0FFC) (RC4092): MQRCCF_SSL_CIPHER_SPEC_ERROR" on page 1671
"4093 (0FFD) (RC4093): MQRCCF_SSL_PEER_NAME_ERROR" on page 1671
"4094 (0FFE) (RC4094): MQRCCF_SSL_CLIENT_AUTH_ERROR" on page 1671
"4095 (0FFF) (RC4095): MQRCCF_RETAINED_NOT_SUPPORTED" on page 1671

Related concepts:



IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

Related reference:



Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)

"API completion and reason codes" on page 1380

"Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes" on page 1672

"WCF custom channel exceptions" on page 1676

3001 (0BB9) (RC3001): MQRCCF_CFH_TYPE_ERROR

Explanation

Type not valid.

The MQCFH *Type* field value was not valid.

Programmer response

Specify a valid type.

3002 (0BBA) (RC3002): MQRCCF_CFH_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFH *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3003 (0BBB) (RC3003): MQRCCF_CFH_VERSION_ERROR

Explanation

Structure version number is not valid.

The MQCFH *Version* field value was not valid.

Note that z/OS requires MQCFH_VERSION_3

Programmer response

Specify a valid structure version number.

3004 (0BBC) (RC3004): MQRCCF_CFH_MSG_SEQ_NUMBER_ERR

Explanation

Message sequence number not valid.

The MQCFH *MsgSeqNumber* field value was not valid.

Programmer response

Specify a valid message sequence number.

3005 (0BBD) (RC3005): MQRCCF_CFH_CONTROL_ERROR

Explanation

Control option not valid.

The MQCFH *Control* field value was not valid.

Programmer response

Specify a valid control option.

3006 (0BBE) (RC3006): MQRCCF_CFH_PARM_COUNT_ERROR

Explanation

Parameter count not valid.

The MQCFH *ParameterCount* field value was not valid.

Programmer response

Specify a valid parameter count.

3007 (0BBF) (RC3007): MQRCCF_CFH_COMMAND_ERROR

Explanation

Command identifier not valid.

The MQCFH *Command* field value was not valid.

Programmer response

Specify a valid command identifier.

3008 (0BC0) (RC3008): MQRCCF_COMMAND_FAILED

Explanation

Command failed.

The command has failed.

Programmer response

Refer to the previous error messages for this command.

3009 (0BC1) (RC3009): MQRCCF_CFIN_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFIN or MQCFIN64 *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3010 (0BC2) (RC3010): MQRCCF_CFST_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFST *StrucLength* field value was not valid. The value was not a multiple of four or was inconsistent with the MQCFST *StringLength* field value.

Programmer response

Specify a valid structure length.

3011 (0BC3) (RC3011): MQRCCF_CFST_STRING_LENGTH_ERR

Explanation

String length not valid.

The MQCFST *StringLength* field value was not valid. The value was negative or greater than the maximum permitted length of the parameter specified in the *Parameter* field.

Programmer response

Specify a valid string length for the parameter.

3012 (0BC4) (RC3012): MQRCCF_FORCE_VALUE_ERROR

Explanation

Force value not valid.

The force value specified was not valid.

Programmer response

Specify a valid force value.

3013 (0BC5) (RC3013): MQRCCF_STRUCTURE_TYPE_ERROR

Explanation

Structure type not valid.

The structure *Type* value was not valid.

Programmer response

Specify a valid structure type.

3014 (0BC6) (RC3014): MQRCCF_CFIN_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFIN or MQCFIN64 *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3015 (0BC7) (RC3015): MQRCCF_CFST_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFST *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3016 (0BC8) (RC3016): MQRCCF_MSG_LENGTH_ERROR

Explanation

Message length not valid.

The message data length was inconsistent with the length implied by the parameters in the message, or a positional parameter was out of sequence.

Programmer response

Specify a valid message length, and check that positional parameters are in the correct sequence.

3017 (0BC9) (RC3017): MQRCCF_CFIN_DUPLICATE_PARM

Explanation

Duplicate parameter.

Two MQCFIN or MQCFIN64 or MQCFIL or MQCFIL64 structures, or any two of those types of structure, with the same parameter identifier were present.

Programmer response

Check for and remove duplicate parameters.

3018 (0BCA) (RC3018): MQRCCF_CFST_DUPLICATE_PARM

Explanation

Duplicate parameter.

Two MQCFST structures, or an MQCFSL followed by an MQCFST structure, with the same parameter identifier were present.

Programmer response

Check for and remove duplicate parameters.

3019 (0BCB) (RC3019): MQRCCF_PARM_COUNT_TOO_SMALL

Explanation

Parameter count too small.

The MQCFH *ParameterCount* field value was less than the minimum required for the command.

Programmer response

Specify a parameter count that is valid for the command.

3020 (0BCC) (RC3020): MQRCCF_PARM_COUNT_TOO_BIG

Explanation

Parameter count too big.

The MQCFH *ParameterCount* field value was more than the maximum for the command.

Programmer response

Specify a parameter count that is valid for the command.

3021 (0BCD) (RC3021): MQRCCF_Q_ALREADY_IN_CELL

Explanation

Queue already exists in cell.

An attempt was made to define a queue with cell scope, or to change the scope of an existing queue from queue-manager scope to cell scope, but a queue with that name already existed in the cell.

Programmer response

Do one of the following:

- Delete the existing queue and retry the operation.
- Change the scope of the existing queue from cell to queue-manager and retry the operation.
- Create the new queue with a different name.

3022 (0BCE) (RC3022): MQRCCF_Q_TYPE_ERROR

Explanation

Queue type not valid.

The *QType* value was not valid.

Programmer response

Specify a valid queue type.

3023 (0BCF) (RC3023): MQRCCF_MD_FORMAT_ERROR

Explanation

Format not valid.

The MQMD *Format* field value was not MQFMT_ADMIN.

Programmer response

Specify the valid format.

3024 (0BD0) (RC3024): MQRCCF_CFSL_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFSL *StrucLength* field value was not valid. The value was not a multiple of four or was inconsistent with the MQCFSL *StringLength* field value.

Programmer response

Specify a valid structure length.

3025 (0BD1) (RC3025): MQRCCF_REPLACE_VALUE_ERROR

Explanation

Replace value not valid.

The *Replace* value was not valid.

Programmer response

Specify a valid replace value.

3026 (0BD2) (RC3026): MQRCCF_CFIL_DUPLICATE_VALUE

Explanation

Duplicate parameter value.

In the MQCFIL or MQCFIL64 structure, there was a duplicate parameter value in the list.

Programmer response

Check for and remove duplicate parameter values.

3027 (0BD3) (RC3027): MQRCCF_CFIL_COUNT_ERROR

Explanation

Count of parameter values not valid.

The MQCFIL or MQCFIL64 *Count* field value was not valid. The value was negative or greater than the maximum permitted for the parameter specified in the *Parameter* field.

Programmer response

Specify a valid count for the parameter.

3028 (0BD4) (RC3028): MQRCCF_CFIL_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFIL or MQCFIL64 *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3029 (0BD5) (RC3029): MQRCCF_MODE_VALUE_ERROR

Explanation

Mode value not valid.

The *Mode* value was not valid.

Programmer response

Specify a valid mode value.

3029 (0BD5) (RC3029): MQRCCF_QUIESCE_VALUE_ERROR

Explanation

Former name for MQRCCF_MODE_VALUE_ERROR.

3030 (0BD6) (RC3030): MQRCCF_MSG_SEQ_NUMBER_ERROR

Explanation

Message sequence number not valid.

The message sequence number parameter value was not valid.

Programmer response

Specify a valid message sequence number.

3031 (0BD7) (RC3031): MQRCCF_PING_DATA_COUNT_ERROR

Explanation

Data count not valid.

The Ping Channel *DataCount* value was not valid.

Programmer response

Specify a valid data count value.

3032 (0BD8) (RC3032): MQRCCF_PING_DATA_COMPARE_ERROR

Explanation

Ping Channel command failed.

The Ping Channel command failed with a data compare error. The data offset that failed is returned in the message (with parameter identifier MQIACF_ERROR_OFFSET).

Programmer response

Consult your systems administrator.

3033 (0BD9) (RC3033): MQRCCF_CFSL_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFSL *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3034 (0BDA) (RC3034): MQRCCF_CHANNEL_TYPE_ERROR

Explanation

Channel type not valid.

The *ChannelType* specified was not valid, or did not match the type of an existing channel being copied, changed or replaced, or the command and the specified disposition cannot be used with that type of channel.

Programmer response

Specify a valid channel name, type, or disposition.

3035 (0BDB) (RC3035): MQRCCF_PARM_SEQUENCE_ERROR

Explanation

Parameter sequence not valid.

The sequence of parameters is not valid for this command.

Programmer response

Specify the positional parameters in a valid sequence for the command.

3036 (0BDC) (RC3036): MQRCCF_XMIT_PROTOCOL_TYPE_ERR

Explanation

Transmission protocol type not valid.

The *TransportType* value was not valid.

Programmer response

Specify a valid transmission protocol type.

3037 (0BDD) (RC3037): MQRCCF_BATCH_SIZE_ERROR

Explanation

Batch size not valid.

The batch size specified was not valid.

Programmer response

Specify a valid batch size value.

3038 (0BDE) (RC3038): MQRCCF_DISC_INT_ERROR

Explanation

Disconnection interval not valid.

The disconnection interval specified was not valid.

Programmer response

Specify a valid disconnection interval.

3039 (0BDF) (RC3039): MQRCCF_SHORT_RETRY_ERROR

Explanation

Short retry count not valid.

The *ShortRetryCount* value was not valid.

Programmer response

Specify a valid short retry count value.

3040 (0BE0) (RC3040): MQRCCF_SHORT_TIMER_ERROR

Explanation

Short timer value not valid.

The *ShortRetryInterval* value was not valid.

Programmer response

Specify a valid short timer value.

3041 (0BE1) (RC3041): MQRCCF_LONG_RETRY_ERROR

Explanation

Long retry count not valid.

The long retry count value specified was not valid.

Programmer response

Specify a valid long retry count value.

3042 (0BE2) (RC3042): MQRCCF_LONG_TIMER_ERROR

Explanation

Long timer not valid.

The long timer (long retry wait interval) value specified was not valid.

Programmer response

Specify a valid long timer value.

3043 (0BE3) (RC3043): MQRCCF_SEQ_NUMBER_WRAP_ERROR

Explanation

Sequence wrap number not valid.

The *SeqNumberWrap* value was not valid.

Programmer response

Specify a valid sequence wrap number.

3044 (0BE4) (RC3044): MQRCCF_MAX_MSG_LENGTH_ERROR

Explanation

Maximum message length not valid.

The maximum message length value specified was not valid.

Programmer response

Specify a valid maximum message length.

3045 (0BE5) (RC3045): MQRCCF_PUT_AUTH_ERROR

Explanation

Put authority value not valid.

The *PutAuthority* value was not valid.

Programmer response

Specify a valid authority value.

3046 (0BE6) (RC3046): MQRCCF_PURGE_VALUE_ERROR

Explanation

Purge value not valid.

The *Purge* value was not valid.

Programmer response

Specify a valid purge value.

3047 (0BE7) (RC3047): MQRCCF_CFIL_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFIL or MQCFIL64 *Parameter* field value was not valid, or specifies a parameter that cannot be filtered, or that is also specified as a parameter to select a subset of objects.

Programmer response

Specify a valid parameter identifier.

3048 (0BE8) (RC3048): MQRCCF_MSG_TRUNCATED

Explanation

Message truncated.

The command server received a message that is larger than its maximum valid message size.

Programmer response

Check the message contents are correct.

3049 (0BE9) (RC3049): MQRCCF_CCSID_ERROR

Explanation

Coded character-set identifier error.

In a command message, one of the following occurred:

- The *CodedCharSetId* field in the message descriptor of the command does not match the coded character-set identifier of the queue manager at which the command is being processed, or
- The *CodedCharSetId* field in a string parameter structure within the message text of the command is not
 - MQCCSI_DEFAULT, or
 - the coded character-set identifier of the queue manager at which the command is being processed, as in the *CodedCharSetId* field in the message descriptor.

The error response message contains the correct value.

This reason can also occur if a ping cannot be performed because the coded character-set identifiers are not compatible. In this case the correct value is not returned.

Programmer response

Construct the command with the correct coded character-set identifier, and specify this in the message descriptor when sending the command. For ping, use a suitable coded character-set identifier.

3050 (0BEA) (RC3050): MQRCCF_ENCODING_ERROR

Explanation

Encoding error.

The *Encoding* field in the message descriptor of the command does not match that required for the platform at which the command is being processed.

Programmer response

Construct the command with the correct encoding, and specify this in the message descriptor when sending the command.

3052 (0BEC) (RC3052): MQRCCF_DATA_CONV_VALUE_ERROR

Explanation

Data conversion value not valid.

The value specified for *DataConversion* is not valid.

Programmer response

Specify a valid value.

3053 (0BED) (RC3053): MQRCCF_INDOUBT_VALUE_ERROR

Explanation

In-doubt value not valid.

The value specified for *InDoubt* is not valid.

Programmer response

Specify a valid value.

3054 (0BEE) (RC3054): MQRCCF_ESCAPE_TYPE_ERROR

Explanation

Escape type not valid.

The value specified for *EscapeType* is not valid.

Programmer response

Specify a valid value.

3062 (0BF6) (RC3062): MQRCCF_CHANNEL_TABLE_ERROR

Explanation

Channel table value not valid.

The *ChannelTable* specified was not valid, or was not appropriate for the channel type specified on an Inquire Channel or Inquire Channel Names command.

Programmer response

Specify a valid channel table value.

3063 (0BF7) (RC3063): MQRCCF_MCA_TYPE_ERROR

Explanation

Message channel agent type not valid.

The *MCAType* value specified was not valid.

Programmer response

Specify a valid value.

3064 (0BF8) (RC3064): MQRCCF_CHL_INST_TYPE_ERROR

Explanation

Channel instance type not valid.

The *ChannelInstanceType* specified was not valid.

Programmer response

Specify a valid channel instance type.

3065 (0BF9) (RC3065): MQRCCF_CHL_STATUS_NOT_FOUND

Explanation

Channel status not found.

For Inquire Channel Status, no channel status is available for the specified channel. This might indicate that the channel has not been used.

Programmer response

None, unless this is unexpected, in which case consult your systems administrator.

3066 (0BFA) (RC3066): MQRCCF_CFSL_DUPLICATE_PARM

Explanation

Duplicate parameter.

Two MQCFSL structures, or an MQCFST followed by an MQCFSL structure, with the same parameter identifier were present.

Programmer response

Check for and remove duplicate parameters.

3067 (0BFB) (RC3067): MQRCCF_CFSL_TOTAL_LENGTH_ERROR

Explanation

Total string length error.

The total length of the strings (not including trailing blanks) in a MQCFSL structure exceeds the maximum allowable for the parameter.

Programmer response

Check that the structure has been specified correctly, and if so reduce the number of strings.

3068 (0BFC) (RC3068): MQRCCF_CFSL_COUNT_ERROR

Explanation

Count of parameter values not valid.

The MQCFSL *Count* field value was not valid. The value was negative or greater than the maximum permitted for the parameter specified in the *Parameter* field.

Programmer response

Specify a valid count for the parameter.

3069 (0BFD) (RC3069): MQRCCF_CFSL_STRING_LENGTH_ERR

Explanation

String length not valid.

The MQCFSL *StringLength* field value was not valid. The value was negative or greater than the maximum permitted length of the parameter specified in the *Parameter* field.

Programmer response

Specify a valid string length for the parameter.

3070 (0BFE) (RC3070): MQRCCF_BROKER_DELETED

Explanation

Broker has been deleted.

When a broker is deleted using the *dltmqbrk* command, all broker queues created by the broker are deleted. Before this can be done the queues are emptied of all command messages; any that are found are placed on the dead-letter queue with this reason code.

Programmer response

Process the command messages that were placed on the dead-letter queue.

3071 (0BFF) (RC3071): MQRCCF_STREAM_ERROR

Explanation

Stream name is not valid.

The stream name parameter is not valid. Stream names must obey the same naming rules as for WebSphere MQ queues.

Programmer response

Retry the command with a valid stream name parameter.

3072 (0C00) (RC3072): MQRCCF_TOPIC_ERROR

Explanation

Topic name is invalid.

A command has been sent to the broker containing a topic name that is not valid. Note that wildcard topic names are not allowed for *Register Publisher* and *Publish* commands.

Programmer response

Retry the command with a valid topic name parameter. Up to 256 characters of the topic name in question are returned with the error response message. If the topic name contains a null character, this is assumed to terminate the string and is not considered to be part of it. A zero length topic name is not valid, as is one that contains an escape sequence that is not valid.

3073 (0C01) (RC3073): MQRCCF_NOT_REGISTERED

Explanation

Subscriber or publisher is not registered.

A *Deregister* command has been issued to remove registrations for a topic, or topics, for which the publisher or subscriber is not registered. If multiple topics were specified on the command, it fails with a completion code of MQCC_WARNING if the publisher or subscriber was registered for some, but not all, of the topics specified. This error code is also returned to a subscriber issuing a *Request Update* command for a topic for which he does not have a subscription.

Programmer response

Investigate why the publisher or subscriber is not registered. In the case of a subscriber, the subscriptions might have expired, or been removed automatically by the broker if the subscriber is no longer authorized.

3074 (0C02) (RC3074): MQRCCF_Q_MGR_NAME_ERROR

Explanation

An invalid or unknown queue manager name has been supplied.

A queue manager name has been supplied as part of a publisher or subscriber identity. This might have been supplied as an explicit parameter or in the *ReplyToQMGr* field in the message descriptor of the command. Either the queue manager name is not valid, or in the case of a subscriber identity, the subscriber's queue could not be resolved because the remote queue manager is not known to the broker queue manager.

Programmer response

Retry the command with a valid queue manager name. If appropriate, the broker includes a further error reason code within the error response message. If one is supplied, follow the guidance for that reason code in "Reason codes" on page 1379 to resolve the problem.

3075 (0C03) (RC3075): MQRCCF_INCORRECT_STREAM

Explanation

Stream name does not match the stream queue it was sent to.

A command has been sent to a stream queue that specified a different stream name parameter.

Programmer response

Retry the command either by sending it to the correct stream queue or by modifying the command so that the stream name parameter matches.

3076 (0C04) (RC3076): MQRCCF_Q_NAME_ERROR

Explanation

An invalid or unknown queue name has been supplied.

A queue name has been supplied as part of a publisher or subscriber identity. This might have been supplied as an explicit parameter or in the *ReplyToQ* field in the message descriptor of the command. Either the queue name is not valid, or in the case of a subscriber identity, the broker has failed to open the queue.

Programmer response

Retry the command with a valid queue name. If appropriate, the broker includes a further error reason code within the error response message. If one is supplied, follow the guidance for that reason code in "Reason codes" on page 1379 to resolve the problem.

3077 (0C05) (RC3077): MQRCCF_NO_RETAINED_MSG

Explanation

No retained message exists for the topic specified.

A *Request Update* command has been issued to request the retained message associated with the specified topic. No retained message exists for that topic.

Programmer response

If the topic or topics in question should have retained messages, the publishers of these topics might not be publishing with the correct publication options to cause their publications to be retained.

3078 (0C06) (RC3078): MQRCCF_DUPLICATE_IDENTITY

Explanation

Publisher or subscriber identity already assigned to another user ID.

Each publisher and subscriber has a unique identity consisting of a queue manager name, a queue name, and optionally a correlation identifier. Associated with each identity is the user ID under which that publisher or subscriber first registered. A specific identity can be assigned only to one user ID at a time. While the identity is registered with the broker all commands wanting to use it must specify the correct user ID. When a publisher or a subscriber no longer has any registrations with the broker the identity can be used by another user ID.

Programmer response

Either retry the command using a different identity or remove all registrations associated with the identity so that it can be used by a different user ID. The user ID to which the identity is currently assigned is returned within the error response message. A *Deregister* command could be issued to remove these registrations. If the user ID in question cannot be used to execute such a command, you need to have the necessary authority to open the SYSTEM.BROKER.CONTROL.QUEUE using the MQOO_ALTERNATE_USER_AUTHORITY option.

3079 (0C07) (RC3079): MQRCCF_INCORRECT_Q

Explanation

Command sent to wrong broker queue.

The command is a valid broker command but the queue it has been sent to is incorrect. *Publish* and *Delete Publication* commands need to be sent to the stream queue, all other commands need to be sent to the SYSTEM.BROKER.CONTROL.QUEUE.

Programmer response

Retry the command by sending it to the correct queue.

3080 (0C08) (RC3080): MQRCCF_CORREL_ID_ERROR

Explanation

Correlation identifier used as part of an identity is all binary zeros.

Each publisher and subscriber is identified by a queue manager name, a queue name, and optionally a correlation identifier. The correlation identifier is typically used to allow multiple subscribers to share the same subscriber queue. In this instance a publisher or subscriber has indicated within the Registration or Publication options supplied on the command that their identity does include a correlation identifier, but a valid identifier has not been supplied. The <RegOpt>CorrelAsId</RegOpt> has been specified, but the correlation identifier of the message is nulls.

Programmer response

Change the program to retry the command ensuring that the correlation identifier supplied in the message descriptor of the command message is not all binary zeros.

3081 (0C09) (RC3081): MQRCCF_NOT_AUTHORIZED

Explanation

Subscriber has insufficient authority.

To receive publications a subscriber application needs both browse authority for the stream queue that it is subscribing to, and put authority for the queue that publications are to be sent to. Subscriptions are rejected if the subscriber does not have both authorities. In addition to having browse authority for the stream queue, a subscriber would also require *altusr* authority for the stream queue to subscribe to certain topics that the broker itself publishes information on. These topics start with the MQ/SA/ prefix.

Programmer response

Ensure that the subscriber has the necessary authorities and reissue the request. The problem might occur because the subscriber's user ID is not known to the broker. This can be identified if a further error reason code of MQRC_UNKNOWN_ENTITY is returned within the error response message.

3082 (0C0A) (RC3082): MQRCCF_UNKNOWN_STREAM

Explanation

Stream is not known by the broker or could not be created.

A command message has been put to the SYSTEM.BROKER.CONTROL.QUEUE for an unknown stream. This error code is also returned if dynamic stream creation is enabled and the broker failed to create a stream queue for the new stream using the SYSTEM.BROKER.MODEL.STREAM queue.

Programmer response

Retry the command for a stream that the broker supports. If the broker should support the stream, either define the stream queue manually, or correct the problem that prevented the broker from creating the stream queue itself.

3083 (0C0B) (RC3083): MQRCCF_REG_OPTIONS_ERROR

Explanation

Invalid registration options have been supplied.

The registration options (between <RegOpt> and </RegOpt>) provided on a command are not valid.

Programmer response

Retry the command with a valid combination of options.

3084 (0C0C) (RC3084): MQRCCF_PUB_OPTIONS_ERROR

Explanation

Invalid publication options have been supplied.

The publication options provided on a Publish command are not valid.

Programmer response

Retry the command with a valid combination of options.

3085 (0C0D) (RC3085): MQRCCF_UNKNOWN_BROKER

Explanation

Command received from an unknown broker.

Within a multi-broker network, related brokers pass subscriptions and publications between each other as a series of command messages. One such command message has been received from a broker that is not, or is no longer, related to the detecting broker.

Programmer response

This situation can occur if the broker network is not quiesced while topology changes are made to the network.

If you are removing a broker from the topology when the queue manager is inactive, your changes are propagated at queue manager restart.

If you are removing a broker from the topology when the queue manager is active, make sure the channels are also active, so that your changes are immediately propagated.

3086 (0C0E) (RC3086): MQRCCF_Q_MGR_CCSID_ERROR

Explanation

Queue manager coded character set identifier error.

The coded character set value for the queue manager was not valid.

Programmer response

Specify a valid value.

3087 (0C0F) (RC3087): MQRCCF_DEL_OPTIONS_ERROR

Explanation

Invalid delete options have been supplied.

The options provided with a *Delete Publication* command are not valid.

Programmer response

Retry the command with a valid combination of options.

3088 (0C10) (RC3088): MQRCCF_CLUSTER_NAME_CONFLICT

Explanation

ClusterName and *ClusterNameList* attributes conflict.

The command was rejected because it would have resulted in the *ClusterName* attribute and the *ClusterNameList* attribute both being nonblank. At least one of these attributes must be blank.

Programmer response

If the command specified one of these attributes only, you must also specify the other one, but with a value of blanks. If the command specified both attributes, ensure that one of them has a value of blanks.

3089 (0C11) (RC3089): MQRCCF_REPOS_NAME_CONFLICT

Explanation

RepositoryName and *RepositoryNameList* attributes conflict.

Either:

- The command was rejected because it would have resulted in the *RepositoryName* and *RepositoryNameList* attributes both being nonblank. At least one of these attributes must be blank.
- For a Reset Queue Manager Cluster command, the queue manager does not provide a full repository management service for the specified cluster. That is, the *RepositoryName* attribute of the queue manager is not the specified cluster name, or the namelist specified by the *RepositoryNameList* attribute does not contain the cluster name.

Programmer response

Reissue the command with the correct values or on the correct queue manager.

3090 (0C12) (RC3090): MQRCCF_CLUSTER_Q_USAGE_ERROR

Explanation

Queue cannot be a cluster queue.

The command was rejected because it would have resulted in a cluster queue also being a transmission queue, which is not permitted, or because the queue in question cannot be a cluster queue.

Programmer response

Ensure that the command specifies either:

- The *Usage* parameter with a value of MQUS_NORMAL, or
- The *ClusterName* and *ClusterNameList* parameters with values of blanks.
- A *QName* parameter with a value that is not one of these reserved queues:
 - SYSTEM.CHANNEL.INITQ
 - SYSTEM.CHANNEL.SYNCQ
 - SYSTEM.CLUSTER.COMMAND.QUEUE
 - SYSTEM.CLUSTER.REPOSITORY.QUEUE
 - SYSTEM.COMMAND.INPUT
 - SYSTEM.QSG.CHANNEL.SYNCQ
 - SYSTEM.QSG.TRANSMIT.QUEUE

3091 (0C13) (RC3091): MQRCCF_ACTION_VALUE_ERROR

Explanation

Action value not valid.

The value specified for *Action* is not valid. There is only one valid value.

Programmer response

Specify MQACT_FORCE_REMOVE as the value of the *Action* parameter.

3092 (0C14) (RC3092): MQRCCF_COMMS_LIBRARY_ERROR

Explanation

Library for requested communications protocol could not be loaded.

The library needed for the requested communications protocol could not be loaded.

Programmer response

Install the library for the required communications protocol, or specify a communications protocol that has already been installed.

3093 (0C15) (RC3093): MQRCCF_NETBIOS_NAME_ERROR

Explanation

NetBIOS listener name not defined.

The NetBIOS listener name is not defined.

Programmer response

Add a local name to the configuration file and retry the operation.

3094 (0C16) (RC3094): MQRCCF_BROKER_COMMAND_FAILED

Explanation

The broker command failed to complete.

A broker command was issued but it failed to complete.

Programmer response

Diagnose the problem using the provided information and issue a corrected command.

For more information, look at the IBM WebSphere MQ error logs.

3095 (0C17) (RC3095): MQRCCF_CFST_CONFLICTING_PARM

Explanation

Conflicting parameters.

The command was rejected because the parameter identified in the error response was in conflict with another parameter in the command.

Programmer response

Consult the description of the parameter identified to ascertain the nature of the conflict, and the correct command.

3096 (0C18) (RC3096): MQRCCF_PATH_NOT_VALID

Explanation

Path not valid.

The path specified was not valid.

Programmer response

Specify a valid path.

3097 (0C19) (RC3097): MQRCCF_PARM_SYNTAX_ERROR

Explanation

Syntax error found in parameter.

The parameter specified contained a syntax error.

Programmer response

Check the syntax for this parameter.

3098 (0C1A) (RC3098): MQRCCF_PWD_LENGTH_ERROR

Explanation

Password length error.

The password string length is rounded up by to the nearest eight bytes. This rounding causes the total length of the *SSLCryptoHardware* string to exceed its maximum.

Programmer response

Decrease the size of the password, or of earlier fields in the *SSLCryptoHardware* string.

3150 (0C4E) (RC3150): MQRCCF_FILTER_ERROR

Explanation

Filter not valid. This could be because either:

1. In an inquire command message, the specification of a filter is not valid.
2. In a publish/subscribe command message, the content-based filter expression supplied in the publish/subscribe command message contains invalid syntax, and cannot be used.

Programmer response

1. Correct the specification of the filter parameter structure in the inquire command message.
2. Correct the syntax of the filter expression in the publish/subscribe command message. The filter expression is the value of the *Filter* tag in the *psc* folder in the MQRFH2 structure. See the *Websphere MQ Integrator V2 Programming Guide* for details of valid syntax.

3151 (0C4F) (RC3151): MQRCCF_WRONG_USER

Explanation

Wrong user.

A publish/subscribe command message cannot be executed on behalf of the requesting user because the subscription that it would update is already owned by a different user. A subscription can be updated or deregistered only by the user that originally registered the subscription.

Programmer response

Ensure that applications that need to issue commands against existing subscriptions are running under the user identifier that originally registered the subscription. Alternatively, use different subscriptions for different users.

3152 (0C50) (RC3152): MQRCCF_DUPLICATE_SUBSCRIPTION

Explanation

The subscription already exists.

A matching subscription already exists.

Programmer response

Either modify the new subscription properties to distinguish it from the existing subscription or deregister the existing subscription. Then reissue the command.

3153 (0C51) (RC3153): MQRCCF_SUB_NAME_ERROR

Explanation

The subscription name parameter is in error.

Either the subscription name is of an invalid format or a matching subscription already exists with no subscription name.

Programmer response

Either correct the subscription name or remove it from the command and reissue the command.

3154 (0C52) (RC3154): MQRCCF_SUB_IDENTITY_ERROR

Explanation

The subscription identity parameter is in error.

Either the supplied value exceeds the maximum length allowed or the subscription identity is not currently a member of the subscription's identity set and a Join registration option was not specified.

Programmer response

Either correct the identity value or specify a Join registration option to add this identity to the identity set for this subscription.

3155 (0C53) (RC3155): MQRCCF_SUBSCRIPTION_IN_USE

Explanation

The subscription is in use.

An attempt to modify or deregister a subscription was attempted by a member of the identity set when they were not the only member of this set.

Programmer response

Reissue the command when you are the only member of the identity set. To avoid the identity set check and force the modification or deregistration remove the subscription identity from the command message and reissue the command.

3156 (0C54) (RC3156): MQRCCF_SUBSCRIPTION_LOCKED

Explanation

The subscription is locked.

The subscription is currently exclusively locked by another identity.

Programmer response

Wait for this identity to release the exclusive lock.

3157 (0C55) (RC3157): MQRCCF_ALREADY_JOINED

Explanation

The identity already has an entry for this subscription.

A Join registration option was specified but the subscriber identity was already a member of the subscription's identity set.

Programmer response

None. The command completed, this reason code is a warning.

3160 (0C58) (RC3160): MQRCCF_OBJECT_IN_USE

Explanation

Object in use by another command.

A modification of an object was attempted while the object was being modified by another command.

Programmer response

Retry the command.

3161 (0C59) (RC3161): MQRCCF_UNKNOWN_FILE_NAME

Explanation

File not defined to CICS.

A file name parameter identifies a file that is not defined to CICS.

Programmer response

Provide a valid file name or create a CSD definition for the required file.

3162 (0C5A) (RC3162): MQRCCF_FILE_NOT_AVAILABLE

Explanation

File not available to CICS.

A file name parameter identifies a file that is defined to CICS, but is not available.

Programmer response

Check that the CSD definition for the file is correct and enabled.

3163 (0C5B) (RC3163): MQRCCF_DISC_RETRY_ERROR

Explanation

Disconnection retry count not valid.

The *DiscRetryCount* value was not valid.

Programmer response

Specify a valid count.

3164 (0C5C) (RC3164): MQRCCF_ALLOC_RETRY_ERROR

Explanation

Allocation retry count not valid.

The *AllocRetryCount* value was not valid.

Programmer response

Specify a valid count.

3165 (0C5D) (RC3165): MQRCCF_ALLOC_SLOW_TIMER_ERROR

Explanation

Allocation slow retry timer value not valid.

The *AllocRetrySlowTimer* value was not valid.

Programmer response

Specify a valid timer value.

3166 (0C5E) (RC3166): MQRCCF_ALLOC_FAST_TIMER_ERROR

Explanation

Allocation fast retry timer value not valid.

The *AllocRetryFastTimer* value was not valid.

Programmer response

Specify a valid value.

3167 (0C5F) (RC3167): MQRCCF_PORT_NUMBER_ERROR

Explanation

Port number value not valid.

The *PortNumber* value was not valid.

Programmer response

Specify a valid port number value.

3168 (0C60) (RC3168): MQRCCF_CHL_SYSTEM_NOT_ACTIVE

Explanation

Channel system is not active.

An attempt was made to start a channel while the channel system was inactive.

Programmer response

Activate the channel system before starting a channel.

3169 (0C61) (RC3169): MQRCCF_ENTITY_NAME_MISSING

Explanation

Entity name required but missing.

A parameter specifying entity names must be supplied.

Programmer response

Specify the required parameter.

3170 (0C62) (RC3170): MQRCCF_PROFILE_NAME_ERROR

Explanation

Profile name not valid.

A profile name is not valid. Profile names might include wildcard characters or might be given explicitly. If you give an explicit profile name, then the object identified by the profile name must exist. This error might also occur if you specify more than one double asterisk in a profile name.

Programmer response

Specify a valid name.

3171 (0C63) (RC3171): MQRCCF_AUTH_VALUE_ERROR

Explanation

Authorization value not valid.

A value for the *AuthorizationList* or *AuthorityRemove* or *AuthorityAdd* parameter was not valid.

Programmer response

Specify a valid value.

3172 (0C64) (RC3172): MQRCCF_AUTH_VALUE_MISSING

Explanation

Authorization value required but missing.

A parameter specifying authorization values must be supplied.

Programmer response

Specify the required parameter.

3173 (0C65) (RC3173): MQRCCF_OBJECT_TYPE_MISSING

Explanation

Object type value required but missing.

A parameter specifying the object type must be supplied.

Programmer response

Specify the required parameter.

3174 (0C66) (RC3174): MQRCCF_CONNECTION_ID_ERROR

Explanation

Error in connection id parameter.

The *ConnectionId* specified was not valid.

Programmer response

Specify a valid connection id.

3175 (0C67) (RC3175): MQRCCF_LOG_TYPE_ERROR

Explanation

Log type not valid.

The log type value specified was not valid.

Programmer response

Specify a valid log type value.

3176 (0C68) (RC3176): MQRCCF_PROGRAM_NOT_AVAILABLE

Explanation

Program not available.

A request to start or stop a service failed because the request to start the program failed. This could be because the program could not be found at the specified location, or that insufficient system resources are available currently to start it.

Programmer response

Check that the correct name is specified in the definition of the service, and that the program is in the appropriate libraries, before retrying the request.

3177 (0C69) (RC3177): MQRCCF_PROGRAM_AUTH_FAILED

Explanation

Program not available.

A request to start or stop a service failed because the user does not have sufficient access authority to start the program at the specified location.

Programmer response

Correct the program name and location, and the user's authority, before retrying the request.

3200 (0C80) (RC3200): MQRCCF_NONE_FOUND

Explanation

No items found matching request criteria.

An Inquire command found no items that matched the specified name and satisfied any other criteria requested.

3201 (0C81) (RC3201): MQRCCF_SECURITY_SWITCH_OFF

Explanation

Security refresh or reverification not processed, security switch set OFF.

Either

- a Reverify Security command was issued, but the subsystem security switch is off, so there are no internal control tables to flag for reverification; or
- a Refresh Security command was issued, but the security switch for the requested class or the subsystem security switch is off.

The switch in question might be returned in the message (with parameter identifier MQIACF_SECURITY_SWITCH).

3202 (0C82) (RC3202): MQRCCF_SECURITY_REFRESH_FAILED

Explanation

Security refresh did not take place.

A SAF RACROUTE REQUEST=STAT call to your external security manager (ESM) returned a non-zero return code. In consequence, the requested security refresh could not be done. The security item affected might be returned in the message (with parameter identifier MQIACF_SECURITY_ITEM).

Possible causes of this problem are:

- The class is not installed
- The class is not active
- The external security manager (ESM) is not active
- The RACF z/OS router table is incorrect

Programmer response

For information about resolving the problem, see the explanations of messages CSQH003I and CSQH004I.

3203 (0C83) (RC3203): MQRCCF_PARM_CONFLICT

Explanation

Incompatible parameters or parameter values.

The parameters or parameter values for a command are incompatible. One of the following occurred:

- A parameter was not specified that is required by another parameter or parameter value.
- A parameter or parameter value was specified that is not allowed with some other parameter or parameter value.
- The values for two specified parameters were not both blank or non-blank.
- The values for two specified parameters were incompatible.

The parameters in question might be returned in the message (with parameter identifiers MQIACF_PARAMETER_ID).

Programmer response

Reissue the command with correct parameters and values.

3204 (0C84) (RC3204): MQRCCF_COMMAND_INHIBITED

Explanation

Commands not allowed at present time.

The queue manager cannot accept commands at the present time, because it is restarting or terminating, or because the command server is not running.

3205 (0C85) (RC3205): MQRCCF_OBJECT_BEING_DELETED

Explanation

Object is being deleted.

The object specified on a command is in the process of being deleted, so the command is ignored.

3207 (0C87) (RC3207): MQRCCF_STORAGE_CLASS_IN_USE

Explanation

Storage class is active or queue is in use.

The command for a local queue involved a change to the *StorageClass* value, but there are messages on the queue, or other threads have the queue open.

Programmer response

Remove the messages from the queue, or wait until any other threads have closed the queue.

3208 (0C88) (RC3208): MQRCCF_OBJECT_NAME_RESTRICTED

Explanation

Incompatible object name and type.

The command used a reserved object name with an incorrect object type or subtype. The object is only allowed to be of a predetermined type, as listed in the explanation of message CSQM108I.

3209 (0C89) (RC3209): MQRCCF_OBJECT_LIMIT_EXCEEDED

Explanation

Local queue limit exceeded.

The command failed because no more local queues could be defined. There is an implementation limit of 524 287 for the total number of local queues that can exist. For shared queues, there is a limit of 512 queues in a single coupling facility structure.

Programmer response

Delete any existing queues that are no longer required.

3210 (0C8A) (RC3210): MQRCCF_OBJECT_OPEN_FORCE

Explanation

Object is in use, but could be changed specifying *Force* as MQFC_YES.

The object specified is in use. This could be because it is open through the API, or for certain parameter changes, because there are messages currently on the queue. The requested changes can be made by specifying *Force* as MQFC_YES on a Change command.

Programmer response

Wait until the object is not in use. Alternatively specify *Force* as MQFC_YES for a change command.

3211 (0C8B) (RC3211): MQRCCF_DISPOSITION_CONFLICT

Explanation

Parameters are incompatible with disposition.

The parameters or parameter values for a command are incompatible with the disposition of an object. One of the following occurred:

- A value specified for the object name or other parameter is not allowed for a local queue with a disposition that is shared or a model queue used to create a dynamic queue that is shared.
- A value specified for a parameter is not allowed for an object with such disposition.
- A value specified for a parameter must be non-blank for an object with such disposition.
- The *CommandScope* and *QSGDisposition* or *ChannelDisposition* parameter values are incompatible.
- The action requested for a channel cannot be performed because it has the wrong disposition.

The parameter and disposition in question may be returned in the message (with parameter identifiers MQIACF_PARAMETER_ID and MQIA_QSG_DISP).

Programmer response

Reissue the command with correct parameters and values.

3212 (0C8C) (RC3212): MQRCCF_Q_MGR_NOT_IN_QSG

Explanation

Queue manager is not in a queue-sharing group.

The command or its parameters are not allowed when the queue manager is not in a queue-sharing group. The parameter in question might be returned in the message (with parameter identifier MQIACF_PARAMETER_ID).

Programmer response

Reissue the command correctly.

3213 (0C8D) (RC3213): MQRCCF_ATTR_VALUE_FIXED

Explanation

Parameter value cannot be changed.

The value for a parameter cannot be changed. The parameter in question might be returned in the message (with parameter identifier MQIACF_PARAMETER_ID).

Programmer response

To change the parameter, the object must be deleted and then created again with the new value.

3215 (0C8F) (RC3215): MQRCCF_NAMELIST_ERROR

Explanation

Namelist is empty or wrong type.

A namelist used to specify a list of clusters has no names in it or does not have type MQNT_CLUSTER or MQNT_NONE.

Programmer response

Reissue the command specifying a namelist that is not empty and has a suitable type.

3217 (0C91) (RC3217): MQRCCF_NO_CHANNEL_INITIATOR

Explanation

Channel initiator not active.

The command requires the channel initiator to be started.

3218 (0C93) (RC3218): MQRCCF_CHANNEL_INITIATOR_ERROR

Explanation

Channel initiator cannot be started, or no suitable channel initiator is available.

This might occur because of the following reasons:

- The channel initiator cannot be started because:
 - It is already active.
 - There are insufficient system resources.
 - The queue manager was shutting down.
- The shared channel cannot be started because there was no suitable channel initiator available for any active queue manager in the queue-sharing group. This could be because:
 - No channel initiators are running.
 - The channel initiators that are running are too busy to allow any channel, or a channel of the particular type, to be started.

3222 (0C96) (RC3222): MQRCCF_COMMAND_LEVEL_CONFLICT

Explanation

Incompatible queue manager command levels.

Changing the *CFLevel* parameter of a CF structure, or deleting a CF structure, requires that all queue managers in the queue-sharing group have a command level of at least 530. Some of the queue managers have a lower level.

3223 (0C97) (RC3223): MQRCCF_Q_ATTR_CONFLICT

Explanation

Queue attributes are incompatible.

The queues involved in a Move Queue command have different values for one or more of these attributes: *DefinitionType*, *HardenGetBackout*, *Usage*. Messages cannot be moved safely if these attributes differ.

3224 (0C98) (RC3224): MQRCCF_EVENTS_DISABLED

Explanation

Events not enabled.

The command required performance or configuration events to be enabled.

Programmer response

Use the Change Queue manager command to enable the events if required.

3225 (0C99) (RC3225): MQRCCF_COMMAND_SCOPE_ERROR

Explanation

Queue-sharing group error.

While processing a command that used the *CommandScope* parameter, an error occurred while trying to send data to the coupling facility.

Programmer response

Notify your system programmer.

3226 (0C9A) (RC3226): MQRCCF_COMMAND_REPLY_ERROR

Explanation

Error saving command reply information.

While processing a command that used the *CommandScope* parameter, or a command for the channel initiator, an error occurred while trying to save information about the command.

Programmer response

The most likely cause is insufficient storage. If the problem persists, you may need to restart the queue manager after making more storage available.

3227 (0C9B) (RC3227): MQRCCF_FUNCTION_RESTRICTED

Explanation

Restricted command or parameter value used.

The command, or the value specified for one of its parameters, is not allowed because the installation and customization options chosen do not allow all functions to be used. The parameter in question might be returned in the message (with parameter identifier MQIACF_PARAMETER_ID).

3228 (0C9C) (RC3228): MQRCCF_PARM_MISSING

Explanation

Required parameter not specified.

The command did not specify a parameter or parameter value that was required. It might be for one of the following reasons:

- A parameter that is always required.
- A parameter that is one of a set of two or more alternative required parameters.
- A parameter that is required because some other parameter was specified.
- A parameter that is a list of values which has too few values.

The parameter in question might be returned in the message (with parameter identifier MQIACF_PARAMETER_ID).

Programmer response

Reissue the command with correct parameters and values.

3229 (0C9D) (RC3229): MQRCCF_PARM_VALUE_ERROR

Explanation

Parameter value invalid.

The value specified for a parameter was not acceptable. It might be for one of the following reasons:

- Outside the acceptable numeric range for the parameter.
- Not one of a list of acceptable values for the parameter.
- Using characters that are invalid for the parameter.
- Completely blank, when such is not allowed for the parameter.
- A filter value that is invalid for the parameter being filtered.

The parameter in question might be returned in the message (with parameter identifier MQIACF_PARAMETER_ID).

Programmer response

Reissue the command with correct parameters and values.

3230 (0C9E) (RC3230): MQRCCF_COMMAND_LENGTH_ERROR

Explanation

Command exceeds allowable length.

The command is so large that its internal form has exceeded the maximum length allowed. The size of the internal form of the command is affected by both the length, and the complexity of the command.

3231 (0C9F) (RC3231): MQRCCF_COMMAND_ORIGIN_ERROR

Explanation

Command issued incorrectly.

The command cannot be issued using command server. This is an internal error.

Programmer response

Notify your system programmer.

3232 (0CA0) (RC3232): MQRCCF_LISTENER_CONFLICT

Explanation

Address conflict for listener.

A listener was already active for a port and IP address combination that conflicted with the *Port* and *IPAddress* values specified by a Start Channel Listener or Stop Channel Listener command. The *Port* and *IPAddress* value combination specified must match a combination for which the listener is active. It cannot be a superset or a subset of that combination.

Programmer response

Reissue the command with correct values, if required.

3233 (0CA1) (RC3233): MQRCCF_LISTENER_STARTED

Explanation

Listener is started.

An attempt was made to start a listener, but it is already active for the requested *TransportType*, *InboundDisposition*, *Port*, and *IPAddress* values. The requested parameter values might be returned in the message, if applicable (with parameter identifiers MQIACH_XMIT_PROTOCOL_TYPE, MQIACH_INBOUND_DISP, MQIACH_PORT_NUMBER, MQCACH_IP_ADDRESS).

3234 (0CA2) (RC3234): MQRCCF_LISTENER_STOPPED

Explanation

Listener is stopped.

An attempt was made to stop a listener, but it is not active or already stopping for the requested *TransportType*, *InboundDisposition*, *Port*, and *IPAddress* values. The requested parameter values might be returned in the message, if applicable (with parameter identifiers MQIACH_XMIT_PROTOCOL_TYPE, MQIACH_INBOUND_DISP, MQIACH_PORT_NUMBER, MQCACH_IP_ADDRESS).

3235 (0CA3) (RC3235): MQRCCF_CHANNEL_ERROR

Explanation

Channel command failed.

A channel command failed because of an error in the channel definition, or at the remote end of the channel, or in the communications system. An error identifier value *nnn* may be returned in the message (with parameter identifier MQIACF_ERROR_ID).

Programmer response

For information about the error, see the explanation of the corresponding error message. Error *nnn* generally corresponds to message CSQX*nnn*, although there are some exceptions. For more information, see  Distributed queuing message codes (*WebSphere MQ V7.1 Reference*).

3236 (0CA4) (RC3236): MQRCCF_CF_STRUC_ERROR

Explanation

CF structure error.

A command could not be processed because of a coupling facility or CF structure error. It might be:

- A Backup CF Structure or Recover CF Structure command when the status of the CF structure is unsuitable. In this case, the CF structure status might be returned in the message together with the CF structure name (with parameter identifiers MQIACF_CF_STRUC_STATUS and MQCA_CF_STRUC_NAME).
- A command could not access an object because of an error in the coupling facility information, or because a CF structure has failed. In this case, the name of the object involved might be returned in the message (with parameter identifier MQCA_Q_NAME, for example).
- A command involving a shared channel could not access the channel status or synchronization key information.

Programmer response

In the case of a Backup CF Structure or Recover CF Structure command, take action appropriate to the CF structure status reported.

In other cases, check for error messages on the console log that might relate to the problem. Check whether the coupling facility structure has failed and check that Db2 is available.

3237 (0CA5) (RC3237): MQRCCF_UNKNOWN_USER_ID

Explanation

User identifier not found.

A user identifier specified in a Reverify Security command was not valid because there was no entry found for it in the internal control table. This could be because the identifier was entered incorrectly in the command, or because it was not in the table (for example, because it had timed-out). The user identifier in question might be returned in the message (with parameter identifier MQCACF_USER_IDENTIFIER).

3238 (0CA6) (RC3238): MQRCCF_UNEXPECTED_ERROR

Explanation

Unexpected or severe error.

An unexpected or severe error or other failure occurred. A code associated with the error might be returned in the message (with parameter identifier MQIACF_ERROR_ID).

Programmer response

Notify your system programmer.

3239 (0CA7) (RC3239): MQRCCF_NO_XCF_PARTNER

Explanation

MQ is not connected to the XCF partner.

The command involving the IMS Bridge cannot be processed because MQ is not connected to the XCF partner. The group and member names of the XCF partner in question might be returned in the message (with parameter identifiers MQCA_XCF_GROUP_NAME and MQCA_XCF_MEMBER_NAME).

3240 (0CA8) (RC3240): MQRCCF_CFGR_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFGR *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3241 (0CA9) (RC3241): MQRCCF_CFIF_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFIF *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3242 (0CAA) (RC3242): MQRCCF_CFIF_OPERATOR_ERROR**Explanation**

Parameter count not valid.

The MQCFIF *Operator* field value was not valid.

Programmer response

Specify a valid operator value.

3243 (0CAB) (RC3243): MQRCCF_CFIF_PARM_ID_ERROR**Explanation**

Parameter identifier is not valid.

The MQCFIF *Parameter* field value was not valid, or specifies a parameter that cannot be filtered, or that is also specified as a parameter to select a subset of objects.

Programmer response

Specify a valid parameter identifier.

3244 (0CAC) (RC3244): MQRCCF_CFSF_FILTER_VAL_LEN_ERR**Explanation**

Filter value length not valid.

The MQCFSF *FilterValueLength* field value was not valid.

Programmer response

Specify a valid length.

3245 (0CAD) (RC3245): MQRCCF_CFSF_LENGTH_ERROR**Explanation**

Structure length not valid.

The MQCFSF *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3246 (0CAE) (RC3246): MQRCCF_CFSF_OPERATOR_ERROR

Explanation

Parameter count not valid.

The MQCFSF *Operator* field value was not valid.

Programmer response

Specify a valid operator value.

3247 (0CAF) (RC3247): MQRCCF_CFSF_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFSF *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3248 (0CB0) (RC3248): MQRCCF_TOO_MANY_FILTERS

Explanation

Too many filters.

The command contained more than the maximum permitted number of filter structures.

Programmer response

Specify the command correctly.

3249 (0CB1) (RC3249): MQRCCF_LISTENER_RUNNING

Explanation

Listener is running.

An attempt was made to perform an operation on a listener, but it is currently active.

Programmer response

Stop the listener if required.

3250 (0CB2) (RC3250): MQRCCF_LSTR_STATUS_NOT_FOUND

Explanation

Listener status not found.

For Inquire Listener Status, no listener status is available for the specified listener. This might indicate that the listener has not been used.

Programmer response

None, unless this is unexpected, in which case consult your systems administrator.

3251 (0CB3) (RC3251): MQRCCF_SERVICE_RUNNING

Explanation

Service is running.

An attempt was made to perform an operation on a service, but it is currently active.

Programmer response

Stop the service if required.

3252 (0CB4) (RC3252): MQRCCF_SERV_STATUS_NOT_FOUND

Explanation

Service status not found.

For Inquire Service Status, no service status is available for the specified service. This might indicate that the service has not been used.

Programmer response

None, unless this is unexpected, in which case consult your systems administrator.

3253 (0CB5) (RC3253): MQRCCF_SERVICE_STOPPED

Explanation

Service is stopped.

An attempt was made to stop a service, but it is not active or already stopping.

3254 (0CB6) (RC3254): MQRCCF_CFBS_DUPLICATE_PARM

Explanation

Duplicate parameter.

Two MQCFBS structures with the same parameter identifier were present.

Programmer response

Check for and remove duplicate parameters.

3255 (0CB7) (RC3255): MQRCCF_CFBS_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFBS *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3256 (0CB8) (RC3256): MQRCCF_CFBS_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFBS *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3257 (0CB9) (RC3257): MQRCCF_CFBS_STRING_LENGTH_ERR

Explanation

String length not valid.

The MQCFBS *StringLength* field value was not valid. The value was negative or greater than the maximum permitted length of the parameter specified in the *Parameter* field.

Programmer response

Specify a valid string length for the parameter.

3258 (0CBA) (RC3258): MQRCCF_CFGR_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFGR *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3259 (0CBB) (RC3259): MQRCCF_CFGR_PARM_COUNT_ERROR

Explanation

Parameter count not valid.

The MQCFGR *ParameterCount* field value was not valid. The value was negative or greater than the maximum permitted for the parameter identifier specified in the *Parameter* field.

Programmer response

Specify a valid count for the parameter.

3260 (0CBC) (RC3260): MQRCCF_CONN_NOT_STOPPED

Explanation

Connection not stopped.

The Stop Connection command could not be executed, so the connection was not stopped.

3261 (0CBD) (RC3261): MQRCCF_SERVICE_REQUEST_PENDING

Explanation

A Suspend or Resume Queue Manager command was issued, or a Refresh Security command, but such a command is currently in progress.

Programmer response

Wait until the current request completes, then reissue the command if necessary.

3262 (0CBE) (RC3262): MQRCCF_NO_START_CMD

Explanation

No start command.

The service cannot be started because no start command is specified in the service definition.

Programmer response

Correct the definition of the service.

3263 (0CBF) (RC3263): MQRCCF_NO_STOP_CMD

Explanation

No stop command.

The service cannot be stopped because no stop command is specified in the service definition.

Programmer response

Correct the definition of the service.

3264 (0CC0) (RC3264): MQRCCF_CFBF_LENGTH_ERROR

Explanation

Structure length not valid.

The MQCFBF *StrucLength* field value was not valid.

Programmer response

Specify a valid structure length.

3265 (0CC1) (RC3265): MQRCCF_CFBF_PARM_ID_ERROR

Explanation

Parameter identifier is not valid.

The MQCFBF *Parameter* field value was not valid.

Programmer response

Specify a valid parameter identifier.

3266 (0CC2) (RC3266): MQRCCF_CFBF_FILTER_VAL_LEN_ERR

Explanation

Filter value length not valid.

The MQCFBF *FilterValueLength* field value was not valid.

Programmer response

Specify a valid length.

3267 (0CC3) (RC3267): MQRCCF_CFBF_OPERATOR_ERROR

Explanation

Parameter count not valid.

The MQCFBF *Operator* field value was not valid.

Programmer response

Specify a valid operator value.

3268 (0CC4) (RC3268): MQRCCF_LISTENER_STILL_ACTIVE

Explanation

Listener still active.

An attempt was made to stop a listener, but it failed and the listener is still active. For example, the listener might still have active channels.

Programmer response

Wait for the active connections to the listener to complete before trying the request again.

3269 (0CC5) (RC3269): MQRCCF_DEF_XMIT_Q_CLUS_ERROR

Explanation

The specified queue is not allowed to be used as the default transmission queue because it is reserved for use exclusively by clustering.

Programmer response

Change the value of the Default Transmission Queue, and try the command again.

3300 (0CE4) (RC3300): MQRCCF_TOPICSTR_ALREADY_EXISTS

Explanation

The topic string specified already exists in another topic object.

Programmer response

Verify that the topic string used is correct.

3301 (0CE5) (RC3301): MQRCCF_SHARING_CONVS_ERROR

Explanation

An invalid value has been given for SharingConversations parameter in the Channel definition

Programmer response

Correct the value used in the PCF SharingConversations (MQCFIN) parameter; see  [Change, Copy, and Create Channel \(WebSphere MQ V7.1 Reference\)](#) for more information.

3302 (0CE6) (RC3302): MQRCCF_SHARING_CONVS_TYPE

Explanation

SharingConversations parameter is not allowed for this channel type.

Programmer response

See  [Change, Copy, and Create Channel \(WebSphere MQ V7.1 Reference\)](#) to ensure that the channel type is compatible with the SharingConversations parameter.

3303 (0CE7) (RC3303): MQRCCF_SECURITY_CASE_CONFLICT

Explanation

A Refresh Security PCF command was issued, but the case currently in use differs from the system setting and if refreshed would result in the set of classes using different case settings.

Programmer response

Check that the class used is set up correctly and that the system setting is correct. If a change in case setting is required, issue the REFRESH SECURITY(*) command to change all classes.

3305 (0CE9) (RC3305): MQRCCF_TOPIC_TYPE_ERROR

Explanation

An Inquire or Delete Topic PCF command was issued with an invalid TopicType parameter.

Programmer response

Correct the TopicType parameter and reissue the command. For more details on the TopicType, see

 [Change, Copy, and Create Topic \(WebSphere MQ V7.1 Reference\)](#).

3306 (0CEA) (RC3306): MQRCCF_MAX_INSTANCES_ERROR

Explanation

An invalid value was given for the maximum number of simultaneous instances of a server-connection channel (MaxInstances) for the channel definition.

Programmer response

See  [Change, Copy, and Create Channel \(WebSphere MQ V7.1 Reference\)](#) for more information and correct the PCF application.

3307 (0CEB) (RC3307): MQRCCF_MAX_INSTS_PER_CLNT_ERR

Explanation

An invalid value was given for the MaxInstancesPerClient property.

Programmer response

See  [Change, Copy, and Create Channel \(WebSphere MQ V7.1 Reference\)](#) for the range of values and correct the application.

3308 (0CEC) (RC3308): MQRCCF_TOPIC_STRING_NOT_FOUND

Explanation

When processing an Inquire Topic Status command, the topic string specified did not match any topic nodes in the topic tree.

Programmer response

Verify the topic string is correct.

3309 (0CED) (RC3309): MQRCCF_SUBSCRIPTION_POINT_ERR

Explanation

The Subscription point was not valid. Valid subscription points are the topic strings of the topic objects listed in the SYSTEM.QPUBSUB.SUBPOINT.NAMELIST.

Programmer response

Use a subscription point that matches the topic string of a topic object listed in the SYSTEM.QPUBSUB.SUBPOINT.NAMELIST (or remove the subscription point parameter and this uses the default subscription point)

3311 (0CEF) (RC2432): MQRCCF_SUB_ALREADY_EXISTS

Explanation

When processing a Copy or Create Subscription command, the target *Subscription* identifier exists.

Programmer response

If you are trying to copy an existing subscription, ensure that the *ToSubscriptionName* parameter contains a unique value. If you are trying to create a Subscription ensure that the combination of the *SubName* parameter, and *TopicObject* parameter or *TopicString* parameter are unique.

3317 (0CF5) (RC3317): MQRCCF_INVALID_DESTINATION

Explanation

The Subscription or Topic object used in a Change, Copy, Create or Delete PCF command is invalid.

Programmer response

Investigate and correct the required parameters for the specific command you are using. For more details, see  [Change, Copy, and Create Subscription \(WebSphere MQ V7.1 Reference\)](#).

3318 (0CF6) (RC3318): MQRCCF_PUBSUB_INHIBITED

Explanation

MQSUB, MQOPEN, MQPUT and MQPUT1 calls are currently inhibited for all publish/subscribe topics, either by means of the queue manager attribute PSMODE or because processing of publish/subscribe state at queue manager start-up has failed, or has not yet completed.

Completion Code

MQCC_FAILED

Programmer response

If this queue manager does not intentionally inhibit publish/subscribe, investigate any error messages that describe the failure at queue manager start-up, or wait until start-up processing completes. You can use the DISPLAY PUBSUB command to check the status of the publish/subscribe engine to ensure it is ready for use, and additionally on z/OS you will receive an information message CSQM076I.

3326 (0CFE) (RC3326): MQRCCF_CHLAUTH_TYPE_ERROR**Explanation**

Channel authentication record type not valid.

The **type** parameter specified on the **set** command was not valid.

Programmer response

Specify a valid type.

3327 (0CFF) (RC3327): MQRCCF_CHLAUTH_ACTION_ERROR**Explanation**

Channel authentication record action not valid.

The **action** parameter specified on the **set** command was not valid.

Programmer response

Specify a valid action.

3335 (0D07) (RC3335): MQRCCF_CHLAUTH_USRSRC_ERROR**Explanation**

Channel authentication record user source not valid.

The **user source** parameter specified on the **set** command was not valid.

Programmer response

Specify a valid user source.

3336 (0D08) (RC3336): MQRCCF_WRONG_CHLAUTH_TYPE**Explanation**

Parameter not allowed for this channel authentication record type.

The parameter is not allowed for the type of channel authentication record being set. Refer to the description of the parameter in error to determine the types of record for which this parameter is valid.

Programmer response

Remove the parameter.

3337 (0D09) (RC3337): MQRCCF_CHLAUTH_ALREADY_EXISTS**Explanation**

Channel authentication record already exists

An attempt was made to add a channel authentication record, but it already exists.

Programmer response

Specify action as MQACT_REPLACE.

3338 (0D0A) (RC3338): MQRCCF_CHLAUTH_NOT_FOUND

Explanation

Channel authentication record not found.

The specified channel authentication record does not exist.

Programmer response

Specify a channel authentication record that exists.

3339 (0D0B) (RC3339): MQRCCF_WRONG_CHLAUTH_ACTION

Explanation

Parameter not allowed for this action on a channel authentication record.

The parameter is not allowed for the action being applied to a channel authentication record. Refer to the description of the parameter in error to determine the actions for which this parameter is valid.

Programmer response

Remove the parameter.

3340 (0D0C) (RC3340): MQRCCF_WRONG_CHLAUTH_USERSRC

Explanation

Parameter not allowed for this channel authentication record user source value.

The parameter is not allowed for a channel authentication record with the value that the **user source** field contains. Refer to the description of the parameter in error to determine the values of user source for which this parameter is valid.

Programmer response

Remove the parameter.

3341 (0D0D) (RC3341): MQRCCF_CHLAUTH_WARN_ERROR

Explanation

Channel authentication record **warn** value not valid.

The **warn** parameter specified on the **set** command was not valid.

Programmer response

Specify a valid value for **warn**.

3342 (0D0E) (RC3342): MQRCCF_WRONG_CHLAUTH_MATCH

Explanation

Parameter not allowed for this channel authentication record **match** value.

The parameter is not allowed for an **inquire channel authentication record** command with the value that the **match** field contains. Refer to the description of the parameter in error to find the values of **match** for which this parameter is valid.

Programmer response

Remove the parameter.

3343 (0D0F) (RC3343): MQRCCF_IPADDR_RANGE_CONFLICT

Explanation

A channel authentication record contained an IP address with a range that overlapped an existing range. A range must be a superset or subset of any existing ranges for the same channel profile name, or completely separate.

Programmer response

Specify a range that is a superset or subset of an existing range, or is completely separate to all existing ranges.

3344 (0D10) (RC3344): MQRCCF_CHLAUTH_MAX_EXCEEDED

Explanation

A channel authentication record was set taking the total number of entries for that type on a single channel profile over the maximum number allowed.

Programmer response

Remove some channel authentication records to make room.

3345 (0D11) (RC3345): MQRCCF_IPADDR_ERROR

Explanation

A channel authentication record contained an invalid IP address, or invalid wildcard pattern to match against IP addresses.

Programmer response

Specify a valid IP address or pattern.

Related reference:

 [Generic IP addresses \(WebSphere MQ V7.1 Reference\)](#)

3346 (0D12) (RC3346): MQRCCF_IPADDR_RANGE_ERROR

Explanation

A channel authentication record contained an IP address with a range that was invalid, for example, the lower number is higher than or equal to the upper number of the range.

Programmer response

Specify a valid range in the IP address.

3347 (0D13) (RC3347): MQRCCF_PROFILE_NAME_MISSING

Explanation

Profile name missing.

A profile name was required for the command but none was specified.

Programmer response

Specify a valid profile name.

3348 (0D14) (RC3348): MQRCCF_CHLAUTH_CLNTUSER_ERROR

Explanation

Channel authentication record **client user** value not valid.

The **client user** value contains a wildcard character, which is not allowed.

Programmer response

Specify a valid value for the client user field.

3349 (0D15) (RC3349): MQRCCF_CHLAUTH_NAME_ERROR

Explanation

Channel authentication record channel name not valid.

When a channel authentication record specifies an IP address to block, the **channel name** value must be a single asterisk (*).

Programmer response

Enter a single asterisk in the channel name.

3350 (0D16) (RC3350): MQRCCF_CHLAUTH_RUNCHECK_ERROR

Runcheck command is using generic values.

Explanation

An Inquire Channel Authentication Record command using MQMATCH_RUNCHECK was issued, but one or more of the input fields on the command were provided with generic values, which is not allowed.

Programmer response

Enter non-generic values for channel name, address, one of the client user ID or remote queue manager and SSL Peer Name if used.

3353 (0D19) (RC3353): MQRCCF_SUITE_B_ERROR

Invalid values have been specified.

Explanation

An invalid combination of values has been specified for the `MQIA_SUITE_B_STRENGTH` parameter.

Programmer response

Review the combination entered and retry with appropriate values.

3364 (0D24) (RC3364): MQRCCF_CERT_VAL_POLICY_ERROR

The certificate validation policy is invalid.

Explanation

An invalid certificate validation policy value was specified for the `MQIA_CERT_VAL_POLICY` attribute. The specified value is unknown or is not supported on the current platform.

Programmer response

Review the value specified and try again with an appropriate certificate validation policy.

4001 (0FA1) (RC4001): MQRCCF_OBJECT_ALREADY_EXISTS

Explanation

Object already exists.

An attempt was made to create an object, but the object already existed and the *Replace* parameter was not specified as `MQRP_YES`.

Programmer response

Specify *Replace* as `MQRP_YES`, or use a different name for the object to be created.

4002 (0FA2) (RC4002): MQRCCF_OBJECT_WRONG_TYPE

Explanation

Object has wrong type or disposition.

An object already exists with the same name but a different subtype or disposition from that specified by the command.

Programmer response

Ensure that the specified object is the same subtype and disposition.

4003 (0FA3) (RC4003): MQRCCF_LIKE_OBJECT_WRONG_TYPE

Explanation

New and existing objects have different subtype.

An attempt was made to create an object based on the definition of an existing object, but the new and existing objects had different subtypes.

Programmer response

Ensure that the new object has the same subtype as the one on which it is based.

4004 (0FA4) (RC4004): MQRCCF_OBJECT_OPEN

Explanation

Object is open.

An attempt was made to operate on an object that was in use.

Programmer response

Wait until the object is not in use, and then retry the operation. Alternatively specify *Force* as MQFC_YES for a change command.

4005 (0FA5) (RC4005): MQRCCF_ATTR_VALUE_ERROR

Explanation

Attribute value not valid or repeated.

One or more of the attribute values specified are not valid or are repeated. The error response message contains the failing attribute selectors (with parameter identifier MQIACF_PARAMETER_ID).

Programmer response

Specify the attribute values correctly.

4006 (0FA6) (RC4006): MQRCCF_UNKNOWN_Q_MGR

Explanation

Queue manager not known.

The queue manager specified was not known.

Programmer response

Specify the name of the queue manager to which the command is sent, or blank.

4007 (0FA7) (RC4007): MQRCCF_Q_WRONG_TYPE

Explanation

Action not valid for the queue of specified type.

An attempt was made to perform an action on a queue of the wrong type.

Programmer response

Specify a queue of the correct type.

4008 (0FA8) (RC4008): MQRCCF_OBJECT_NAME_ERROR

Explanation

Name not valid.

An object or other name name was specified using characters that were not valid.

Programmer response

Specify only valid characters for the name.

4009 (0FA9) (RC4009): MQRCCF_ALLOCATE_FAILED

Explanation

Allocation failed.

An attempt to allocate a conversation to a remote system failed. The error might be due to an entry in the channel definition that is not valid, or it might be that the listening program at the remote system is not running.

Programmer response

Ensure that the channel definition is correct, and start the listening program if necessary. If the error persists, consult your systems administrator.

4010 (0FAA) (RC4010): MQRCCF_HOST_NOT_AVAILABLE

Explanation

Remote system not available.

An attempt to allocate a conversation to a remote system was unsuccessful. The error might be transitory, and the allocate might succeed later. This reason can occur if the listening program at the remote system is not running.

Programmer response

Ensure that the listening program is running, and retry the operation.

4011 (0FAB) (RC4011): MQRCCF_CONFIGURATION_ERROR

Explanation

Configuration error.

There was a configuration error in the channel definition or communication subsystem, and allocation of a conversation was not possible. This might be caused by one of the following:

- For LU 6.2, either the *ModeName* or the *TpName* is incorrect. The *ModeName* must match that on the remote system, and the *TpName* must be specified. (On IBM i, these are held in the communications Side Object.)
- For LU 6.2, the session might not be established.
- For TCP, the *ConnectionName* in the channel definition cannot be resolved to a network address. This might be because the name has not been correctly specified, or because the name server is not available.
- The requested communications protocol might not be supported on the platform.

Programmer response

Identify the error and take appropriate action.

4012 (0FAC) (RC4012): MQRCCF_CONNECTION_REFUSED

Explanation

Connection refused.

The attempt to establish a connection to a remote system was rejected. The remote system might not be configured to allow a connection from this system.

- For LU 6.2 either the user ID or the password supplied to the remote system is incorrect.
- For TCP the remote system might not recognize the local system as valid, or the TCP listener program might not be started.

Programmer response

Correct the error or restart the listener program.

4013 (0FAD) (RC4013): MQRCCF_ENTRY_ERROR

Explanation

Connection name not valid.

The connection name in the channel definition could not be resolved into a network address. Either the name server does not contain the entry, or the name server was not available.

Programmer response

Ensure that the connection name is correctly specified and that the name server is available.

4014 (0FAE) (RC4014): MQRCCF_SEND_FAILED

Explanation

Send failed.

An error occurred while sending data to a remote system. This might be caused by a communications failure.

Programmer response

Consult your systems administrator.

4015 (0FAF) (RC4015): MQRCCF_RECEIVED_DATA_ERROR

Explanation

Received data error.

An error occurred while receiving data from a remote system. This might be caused by a communications failure.

Programmer response

Consult your systems administrator.

4016 (0FB0) (RC4016): MQRCCF_RECEIVE_FAILED

Explanation

Receive failed.

The receive operation failed.

Programmer response

Correct the error and retry the operation.

4017 (0FB1) (RC4017): MQRCCF_CONNECTION_CLOSED

Explanation

Connection closed.

An error occurred while receiving data from a remote system. The connection to the remote system has unexpectedly terminated.

Programmer response

Contact your systems administrator.

4018 (0FB2) (RC4018): MQRCCF_NO_STORAGE

Explanation

Not enough storage available.

Insufficient storage is available.

Programmer response

Consult your systems administrator.

4019 (0FB3) (RC4019): MQRCCF_NO_COMMS_MANAGER

Explanation

Communications manager not available.

The communications subsystem is not available.

Programmer response

Ensure that the communications subsystem has been started.

4020 (0FB4) (RC4020): MQRCCF_LISTENER_NOT_STARTED

Explanation

Listener not started.

The listener program could not be started. Either the communications subsystem has not been started, or the number of current channels using the communications subsystem is the maximum allowed, or there are too many jobs waiting in the queue.

Programmer response

Ensure the communications subsystem is started or retry the operation later. Increase the number of current channels allowed, if appropriate.

4024 (0FB8) (RC4024): MQRCCF_BIND_FAILED

Explanation

Bind failed.

The bind to a remote system during session negotiation has failed.

Programmer response

Consult your systems administrator.

4025 (0FB9) (RC4025): MQRCCF_CHANNEL_INDOUBT

Explanation

Channel in-doubt.

The requested operation cannot complete because the channel is in doubt.

Programmer response

Examine the status of the channel, and either restart a channel to resolve the in-doubt state, or resolve the channel.

4026 (0FBA) (RC4026): MQRCCF_MQCONN_FAILED

Explanation

MQCONN call failed.

Programmer response

Check whether the queue manager is active.

4027 (0FBB) (RC4027): MQRCCF_MQOPEN_FAILED

Explanation

MQOPEN call failed.

Programmer response

Check whether the queue manager is active, and the queues involved are correctly set up.

4028 (0FBC) (RC4028): MQRCCF_MQGET_FAILED

Explanation

MQGET call failed.

Programmer response

Check whether the queue manager is active, and the queues involved are correctly set up, and enabled for MQGET.

4029 (0FBD) (RC4029): MQRCCF_MQPUT_FAILED

Explanation

MQPUT call failed.

Programmer response

Check whether the queue manager is active, and the queues involved are correctly set up, and not inhibited for puts.

4030 (0FBE) (RC4030): MQRCCF_PING_ERROR

Explanation

Ping error.

A ping operation can only be issued for a sender or server channel. If the local channel is a receiver channel, you must issue the ping from a remote queue manager.

Programmer response

Reissue the ping request for a different channel of the correct type, or for a receiver channel from a different queue manager.

4031 (0FBF) (RC4031): MQRCCF_CHANNEL_IN_USE

Explanation

Channel in use.

An attempt was made to perform an operation on a channel, but the channel is currently active.

Programmer response

Stop the channel or wait for it to terminate.

4032 (0FC0) (RC4032): MQRCCF_CHANNEL_NOT_FOUND

Explanation

Channel not found.

The channel specified does not exist.

Programmer response

Specify the name of a channel which exists.

4033 (0FC1) (RC4033): MQRCCF_UNKNOWN_REMOTE_CHANNEL

Explanation

Remote channel not known.

There is no definition of the referenced channel at the remote system.

Programmer response

Ensure that the local channel is correctly defined. If it is, add an appropriate channel definition at the remote system.

4034 (0FC2) (RC4034): MQRCCF_REMOTE_QM_UNAVAILABLE

Explanation

Remote queue manager not available.

The channel cannot be started because the remote queue manager is not available.

Programmer response

Start the remote queue manager.

4035 (0FC3) (RC4035): MQRCCF_REMOTE_QM_TERMINATING

Explanation

Remote queue manager terminating.

The channel is ending because the remote queue manager is terminating.

Programmer response

Restart the remote queue manager.

4036 (0FC4) (RC4036): MQRCCF_MQINQ_FAILED

Explanation

MQINQ call failed.

Programmer response

Check whether the queue manager is active.

4037 (0FC5) (RC4037): MQRCCF_NOT_XMIT_Q

Explanation

Queue is not a transmission queue.

The queue specified in the channel definition is not a transmission queue, or is in use.

Programmer response

Ensure that the queue is specified correctly in the channel definition, and that it is correctly defined to the queue manager.

4038 (0FC6) (RC4038): MQRCCF_CHANNEL_DISABLED

Explanation

Channel disabled.

An attempt was made to use a channel, but the channel was disabled (that is, stopped).

Programmer response

Start the channel.

4039 (0FC7) (RC4039): MQRCCF_USER_EXIT_NOT_AVAILABLE

Explanation

User exit not available.

The channel was terminated because the user exit specified does not exist.

Programmer response

Ensure that the user exit is correctly specified and the program is available.

4040 (0FC8) (RC4040): MQRCCF_COMMIT_FAILED

Explanation

Commit failed.

An error was received when an attempt was made to commit a unit of work.

Programmer response

Consult your systems administrator.

4041 (0FC9) (RC4041): MQRCCF_WRONG_CHANNEL_TYPE

Explanation

Parameter not allowed for this channel type.

The parameter is not allowed for the type of channel being created, copied, or changed. Refer to the description of the parameter in error to determine the types of channel for which the parameter is valid

Programmer response

Remove the parameter.

4042 (0FCA) (RC4042): MQRCCF_CHANNEL_ALREADY_EXISTS

Explanation

Channel already exists.

An attempt was made to create a channel but the channel already existed and *Replace* was not specified as MQRP_YES.

Programmer response

Specify *Replace* as MQRP_YES or use a different name for the channel to be created.

4043 (0FCB) (RC4043): MQRCCF_DATA_TOO_LARGE

Explanation

Data too large.

The data to be sent exceeds the maximum that can be supported for the command.

Programmer response

Reduce the size of the data.

4044 (0FCC) (RC4044): MQRCCF_CHANNEL_NAME_ERROR

Explanation

Channel name error.

The *ChannelName* parameter contained characters that are not allowed for channel names.

Programmer response

Specify a valid name.

4045 (0FCD) (RC4045): MQRCCF_XMIT_Q_NAME_ERROR

Explanation

Transmission queue name error.

The *XmitQName* parameter contains characters that are not allowed for queue names. This reason code also occurs if the parameter is not present when a sender or server channel is being created, and no default value is available.

Programmer response

Specify a valid name, or add the parameter.

4047 (0FCF) (RC4047): MQRCCF_MCA_NAME_ERROR

Explanation

Message channel agent name error.

The *MCAName* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4048 (0FD0) (RC4048): MQRCCF_SEND_EXIT_NAME_ERROR

Explanation

Channel send exit name error.

The *SendExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4049 (0FD1) (RC4049): MQRCCF_SEC_EXIT_NAME_ERROR

Explanation

Channel security exit name error.

The *SecurityExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4050 (0FD2) (RC4050): MQRCCF_MSG_EXIT_NAME_ERROR

Explanation

Channel message exit name error.

The *MsgExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4051 (0FD3) (RC4051): MQRCCF_RCV_EXIT_NAME_ERROR

Explanation

Channel receive exit name error.

The *ReceiveExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4052 (0FD4) (RC4052): MQRCCF_XMIT_Q_NAME_WRONG_TYPE

Explanation

Transmission queue name not allowed for this channel type.

The *XmitQName* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4053 (0FD5) (RC4053): MQRCCF_MCA_NAME_WRONG_TYPE

Explanation

Message channel agent name not allowed for this channel type.

The *MCAName* parameter is only allowed for sender, server or requester channel types.

Programmer response

Remove the parameter.

4054 (0FD6) (RC4054): MQRCCF_DISC_INT_WRONG_TYPE

Explanation

Disconnection interval not allowed for this channel type.

The *DiscInterval* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4055 (0FD7) (RC4055): MQRCCF_SHORT_RETRY_WRONG_TYPE

Explanation

Short retry parameter not allowed for this channel type.

The *ShortRetryCount* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4056 (0FD8) (RC4056): MQRCCF_SHORT_TIMER_WRONG_TYPE

Explanation

Short timer parameter not allowed for this channel type.

The *ShortRetryInterval* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4057 (0FD9) (RC4057): MQRCCF_LONG_RETRY_WRONG_TYPE

Explanation

Long retry parameter not allowed for this channel type.

The *LongRetryCount* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4058 (0FDA) (RC4058): MQRCCF_LONG_TIMER_WRONG_TYPE

Explanation

Long timer parameter not allowed for this channel type.

The *LongRetryInterval* parameter is only allowed for sender or server channel types.

Programmer response

Remove the parameter.

4059 (0FDB) (RC4059): MQRCCF_PUT_AUTH_WRONG_TYPE

Explanation

Put authority parameter not allowed for this channel type.

The *PutAuthority* parameter is only allowed for receiver or requester channel types.

Programmer response

Remove the parameter.

4061 (0FDD) (RC4061): MQRCCF_MISSING_CONN_NAME

Explanation

Connection name parameter required but missing.

The *ConnectionName* parameter is required for sender or requester channel types, but is not present.

Programmer response

Add the parameter.

4062 (0FDE) (RC4062): MQRCCF_CONN_NAME_ERROR

Explanation

Error in connection name parameter.

The *ConnectionName* parameter contains one or more blanks at the start of the name.

Programmer response

Specify a valid connection name.

4063 (0FDF) (RC4063): MQRCCF_MQSET_FAILED

Explanation

MQSET call failed.

Programmer response

Check whether the queue manager is active.

4064 (0FE0) (RC4064): MQRCCF_CHANNEL_NOT_ACTIVE

Explanation

Channel not active.

An attempt was made to stop a channel, but the channel was already stopped.

Programmer response

No action is required.

4065 (0FE1) (RC4065): MQRCCF_TERMINATED_BY_SEC_EXIT

Explanation

Channel terminated by security exit.

A channel security exit terminated the channel.

Programmer response

Check that the channel is attempting to connect to the correct queue manager, and if so that the security exit is specified correctly, and is working correctly, at both ends.

4067 (0FE3) (RC4067): MQRCCF_DYNAMIC_Q_SCOPE_ERROR

Explanation

Dynamic queue scope error.

The *Scope* attribute of the queue is to be MQSCO_CELL, but this is not allowed for a dynamic queue.

Programmer response

Predefine the queue if it is to have cell scope.

4068 (0FE4) (RC4068): MQRCCF_CELL_DIR_NOT_AVAILABLE

Explanation

Cell directory is not available.

The *Scope* attribute of the queue is to be MQSCO_CELL, but no name service supporting a cell directory has been configured.

Programmer response

Configure the queue manager with a suitable name service.

4069 (0FE5) (RC4069): MQRCCF_MR_COUNT_ERROR

Explanation

Message retry count not valid.

The *MsgRetryCount* value was not valid.

Programmer response

Specify a value in the range 0-999 999 999.

4070 (0FE6) (RC4070): MQRCCF_MR_COUNT_WRONG_TYPE

Explanation

Message-retry count parameter not allowed for this channel type.

The *MsgRetryCount* parameter is allowed only for receiver and requester channels.

Programmer response

Remove the parameter.

4071 (0FE7) (RC4071): MQRCCF_MR_EXIT_NAME_ERROR

Explanation

Channel message-retry exit name error.

The *MsgRetryExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4072 (0FE8) (RC4072): MQRCCF_MR_EXIT_NAME_WRONG_TYPE

Explanation

Message-retry exit parameter not allowed for this channel type.

The *MsgRetryExit* parameter is allowed only for receiver and requester channels.

Programmer response

Remove the parameter.

4073 (0FE9) (RC4073): MQRCCF_MR_INTERVAL_ERROR

Explanation

Message retry interval not valid.

The *MsgRetryInterval* value was not valid.

Programmer response

Specify a value in the range 0-999 999 999.

4074 (0FEA) (RC4074): MQRCCF_MR_INTERVAL_WRONG_TYPE

Explanation

Message-retry interval parameter not allowed for this channel type.

The *MsgRetryInterval* parameter is allowed only for receiver and requester channels.

Programmer response

Remove the parameter.

4075 (0FEB) (RC4075): MQRCCF_NPM_SPEED_ERROR

Explanation

Nonpersistent message speed not valid.

The *NonPersistentMsgSpeed* value was not valid.

Programmer response

Specify MQNPMS_NORMAL or MQNPMS_FAST.

4076 (0FEC) (RC4076): MQRCCF_NPM_SPEED_WRONG_TYPE

Explanation

Nonpersistent message speed parameter not allowed for this channel type.

The *NonPersistentMsgSpeed* parameter is allowed only for sender, receiver, server, requester, cluster sender, and cluster receiver channels.

Programmer response

Remove the parameter.

4077 (0FED) (RC4077): MQRCCF_HB_INTERVAL_ERROR

Explanation

Heartbeat interval not valid.

The *HeartbeatInterval* value was not valid.

Programmer response

Specify a value in the range 0-999 999.

4078 (0FEE) (RC4078): MQRCCF_HB_INTERVAL_WRONG_TYPE

Explanation

Heartbeat interval parameter not allowed for this channel type.

The *HeartbeatInterval* parameter is allowed only for receiver and requester channels.

Programmer response

Remove the parameter.

4079 (0FEF) (RC4079): MQRCCF_CHAD_ERROR

Explanation

Channel automatic definition error.

The *ChannelAutoDef* value was not valid.

Programmer response

Specify MQCHAD_ENABLED or MQCHAD_DISABLED.

4080 (0FF0) (RC4080): MQRCCF_CHAD_WRONG_TYPE

Explanation

Channel automatic definition parameter not allowed for this channel type.

The *ChannelAutoDef* parameter is allowed only for receiver and server-connection channels.

Programmer response

Remove the parameter.

4081 (0FF1) (RC4081): MQRCCF_CHAD_EVENT_ERROR

Explanation

Channel automatic definition event error.

The *ChannelAutoDefEvent* value was not valid.

Programmer response

Specify MQEVR_ENABLED or MQEVR_DISABLED.

4082 (0FF2) (RC4082): MQRCCF_CHAD_EVENT_WRONG_TYPE

Explanation

Channel automatic definition event parameter not allowed for this channel type.

The *ChannelAutoDefEvent* parameter is allowed only for receiver and server-connection channels.

Programmer response

Remove the parameter.

4083 (0FF3) (RC4083): MQRCCF_CHAD_EXIT_ERROR

Explanation

Channel automatic definition exit name error.

The *ChannelAutoDefExit* value contained characters that are not allowed for program names on the platform in question.

Programmer response

Specify a valid name.

4084 (0FF4) (RC4084): MQRCCF_CHAD_EXIT_WRONG_TYPE

Explanation

Channel automatic definition exit parameter not allowed for this channel type.

The *ChannelAutoDefExit* parameter is allowed only for receiver and server-connection channels.

Programmer response

Remove the parameter.

4085 (0FF5) (RC4085): MQRCCF_SUPPRESSED_BY_EXIT

Explanation

Action suppressed by exit program.

An attempt was made to define a channel automatically, but this was inhibited by the channel automatic definition exit. The *AuxErrorDataInt1* parameter contains the feedback code from the exit indicating why it inhibited the channel definition.

Programmer response

Examine the value of the *AuxErrorDataInt1* parameter, and take any action that is appropriate.

4086 (0FF6) (RC4086): MQRCCF_BATCH_INT_ERROR

Explanation

Batch interval not valid.

The batch interval specified was not valid.

Programmer response

Specify a valid batch interval value.

4087 (0FF7) (RC4087): MQRCCF_BATCH_INT_WRONG_TYPE

Explanation

Batch interval parameter not allowed for this channel type.

The *BatchInterval* parameter is allowed only for sender and server channels.

Programmer response

Remove the parameter.

4088 (0FF8) (RC4088): MQRCCF_NET_PRIORITY_ERROR

Explanation

Network priority value is not valid.

Programmer response

Specify a valid value.

4089 (0FF9) (RC4089): MQRCCF_NET_PRIORITY_WRONG_TYPE

Explanation

Network priority parameter not allowed for this channel type.

The *NetworkPriority* parameter is allowed for sender and server channels only.

Programmer response

Remove the parameter.

4090 (0FFA) (RC4090): MQRCCF_CHANNEL_CLOSED

Explanation

Channel closed.

The channel was closed prematurely. This can occur because a user stopped the channel while it was running, or a channel exit decided to close the channel.

Programmer response

Determine the reason that the channel was closed prematurely. Restart the channel if required.

4092 (0FFC) (RC4092): MQRCCF_SSL_CIPHER_SPEC_ERROR

Explanation

SSL cipher specification not valid.

The *SSLCipherSpec* specified is not valid.

Programmer response

Specify a valid cipher specification.

4093 (0FFD) (RC4093): MQRCCF_SSL_PEER_NAME_ERROR

Explanation

SSL peer name not valid.

The *SSLPeerName* specified is not valid.

Programmer response

Specify a valid peer name.

4094 (0FFE) (RC4094): MQRCCF_SSL_CLIENT_AUTH_ERROR

Explanation

SSL client authentication not valid.

The *SSLClientAuth* specified is not valid.

Programmer response

Specify a valid client authentication.

4095 (0FFF) (RC4095): MQRCCF_RETAINED_NOT_SUPPORTED

Explanation

Retained messages used on restricted stream.

An attempt has been made to use retained messages on a publish/subscribe stream defined to be restricted to JMS usage. JMS does not support the concept of retained messages and the request is rejected.

Programmer response

Either modify the application not to use retained messages, or modify the broker *JmsStreamPrefix* configuration parameter so that this stream is not treated as a JMS stream.

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes

WebSphere MQ can use Secure Sockets Layer (SSL) with the various communication protocols. Use this topic to identify the error codes that can be returned by SSL.

The table in this appendix documents the return codes, in decimal form, from the Secure Sockets Layer (SSL) that can be returned in messages from the distributed queuing component.

If the return code is not listed, or if you want more information, see [IBM Global Security Kit return codes](#).

Table 154. SSL return codes

Return code (decimal)	Explanation
1	Handle is not valid.
3	An internal error has occurred.
4	Insufficient storage is available
5	Handle is in the incorrect state.
6	Key label is not found.
7	No certificates available.
8	Certificate validation error.
9	Cryptographic processing error.
10	ASN processing error.
11	LDAP processing error.
12	An unexpected error has occurred.
102	Error detected while reading key database or SAF key ring.
103	Incorrect key database record format.
106	Incorrect key database password.
109	No certificate authority certificates.
201	No key database password supplied.
202	Error detected while opening the key database.
203	Unable to generate temporary key pair
204	Key database password is expired.
302	Connection is active.
401	Certificate is expired or is not valid yet.
402	No SSL cipher specifications.
403	No certificate received from partner.
405	Certificate format is not supported.
406	Error while reading or writing data.
407	Key label does not exist.
408	Key database password is not correct.
410	SSL message format is incorrect.
411	Message authentication code is incorrect.
412	SSL protocol or certificate type is not supported.

Table 154. SSL return codes (continued)

Return code (decimal)	Explanation
413	Certificate signature is incorrect.
414	Certificate is not valid.
415	SSL protocol violation.
416	Permission denied.
417	Self-signed certificate cannot be validated.
420	Socket closed by remote partner.
421	SSL V2 cipher is not valid.
422	SSL V3 cipher is not valid.
427	LDAP is not available.
428	Key entry does not contain a private key.
429	SSL V2 header is not valid.
431	Certificate is revoked.
432	Session renegotiation is not allowed.
433	Key exceeds allowable export size.
434	Certificate key is not compatible with cipher suite.
435	Certificate authority is unknown.
436	Certificate revocation list cannot be processed.
437	Connection closed.
438	Internal error reported by remote partner.
439	Unknown alert received from remote partner.
501	Buffer size is not valid.
502	Socket request would block.
503	Socket read request would block.
504	Socket write request would block.
505	Record overflow.
601	Protocol is not SSL V3 or TLS V1.
602	Function identifier is not valid.
701	Attribute identifier is not valid.
702	The attribute has a negative length, which is invalid.
703	The enumeration value is invalid for the specified enumeration type.
704	Invalid parameter list for replacing the SID cache routines.
705	The value is not a valid number.
706	Conflicting parameters were set for additional certificate validation
707	The AES cryptographic algorithm is not supported.
708	The PEERID does not have the correct length.
1501	GSK_SC_OK
1502	GSK_SC_CANCEL
1601	The trace started successfully.
1602	The trace stopped successfully.

Table 154. SSL return codes (continued)

Return code (decimal)	Explanation
1603	No trace file was previously started so it cannot be stopped.
1604	Trace file already started so it cannot be started again.
1605	Trace file cannot be opened. The first parameter of <code>gsk_start_trace()</code> must be a valid full path filename.

In some cases, the secure sockets library reports a certificate validation error in an AMQ9633 error message. Table 2 lists the certificate validation errors that can be returned in messages from the distributed queuing component.

Table 155. Certificate validation errors

Return code (decimal)	Explanation
575001	Internal error
575002	ASN error due to a malformed certificate
575003	Cryptographic error
575004	Key database error
575005	Directory error
575006	Invalid implementation library
575008	No appropriate validator
575009	The root CA is not trusted
575010	No certificate chain was built
575011	Digital signature algorithm mismatch
575012	Digital signature mismatch
575013	X.509 version does not allow Key IDs
575014	X.509 version does not allow extensions
575015	Unknown X.509 certificate version
575016	The certificate validity range is invalid
575017	The certificate is not yet valid
575018	The certificate has expired
575019	The certificate contains unknown critical extensions
575020	The certificate contains duplicate extensions
575021	The issuers directory name does not match the issuer's issuer
575022	The Authority Key ID serial number value does not match the serial number of the issuer
575023	The Authority Key ID and Subject Key ID do not match
575024	Unrecognized issuer alternative name
575025	The certificate Basic Constraints forbid use as a CA
575026	The certificate has a non-zero Basic Constraints path length but is not a CA
575027	The certificate Basic Constraints maximum path length was exceeded
575028	The certificate is not permitted to sign other certificates
575029	The certificate is not signed by a CA
575030	Unrecognized Subject Alternative Name

Table 155. Certificate validation errors (continued)

Return code (decimal)	Explanation
575031	The certificate chain is invalid
575032	The certificate is revoked
575033	Unrecognized CRL distribution point
575034	Name chaining failed
575035	Certificate is not in a chain
575036	The CRL is not yet valid
575037	The CRL has expired
575038	The certificate version does not allow critical extensions
575039	Unknown CRL distribution points
575040	No CRLs for CRL distribution points
575041	Indirect CRLs are not supported
575042	Missing issuing CRL distribution point name
575043	Distribution points do not match
575044	No available CRL data source
575045	CA Subject name is null
575046	Distinguished names do not chain
575047	Missing Subject Alternative Name
575048	Unique ID mismatch
575049	Name not permitted
575050	Name excluded
575051	CA certificate is missing Critical Basic Constraints
575052	Name constraints are not critical
575053	Name constraints minimum subtree value if set is not zero
575054	Name constraints maximum subtree value if set is not allowed
575055	Unsupported name constraint
575056	Empty policy constraints
575057	Bad certificate policies
575058	Certificate policies not acceptable
575059	Bad acceptable certificate policies
575060	Certificate policy mappings are critical
575061	Revocation status could not be determined
575062	Extended key usage error
575063	Unknown OCSP version
575064	Unknown OCSP response
575065	Bad OCSP key usage extension
575066	Bad OCSP nonce
575067	Missing OCSP nonce
575068	No OCSP client available

Related concepts:

 IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

Related reference:

 Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)

“API completion and reason codes” on page 1380

“PCF reason codes” on page 1590

“WCF custom channel exceptions”

WCF custom channel exceptions

Diagnostic messages are listed in this topic in numeric order, grouped according to the part of the WCF custom channel from which they originate.

Reading a message

For each message, this information is provided:

- The message identifier, in two parts:
 1. The characters "WCFCH" which identify the message as being from the WCF custom channel for WebSphere MQ
 2. A four-digit decimal code followed by the character 'E'
- The text of the message.
- An explanation of the message giving further information.
- The response required from the user. In some cases, particularly for information messages, the response required might be "none".

Message variables

Some messages display text or numbers that vary according to the circumstances causing the message to occur; these circumstances are known as *message variables*. The message variables are indicated as {0}, {1}, and so on.

In some cases a message might have variables in the Explanation or Response. Find the values of the message variables by looking in the error log. The complete message, including the Explanation and the Response, is recorded there.

The following message types are described:

“WCFCH0001E-0100E: General/State messages” on page 1677

“WCFCH0101E-0200E: URI Properties messages” on page 1678

“WCFCH0201E-0300E: Factory/Listener messages” on page 1680

“WCFCH0301E-0400E: Channel messages” on page 1681

“WCFCH0401E-0500E: Binding messages” on page 1683

“WCFCH0501E-0600E: Binding properties messages” on page 1684

“WCFCH0601E-0700E: Async operations messages” on page 1684

Related concepts:

 IBM WebSphere MQ for z/OS messages, completion, and reason codes (*WebSphere MQ V7.1 Reference*)

Related reference:

 Diagnostic messages: AMQ4000-9999 (*WebSphere MQ V7.1 Reference*)

“API completion and reason codes” on page 1380

“PCF reason codes” on page 1590

“Secure Sockets Layer (SSL) and Transport Layer Security (TLS) return codes” on page 1672

“WCF custom channel exceptions” on page 1676

WCFCH0001E-0100E: General/State messages

Use the following information to understand WCFCH0001E-0100E general/state messages.

WCFCH0001E

An object cannot be opened because its state is '{0}'.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the  WebSphere MQ support Web page, or the  IBM Support Assistant web page, to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0002E

An object cannot be closed because its state is '{0}'.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the  WebSphere MQ support Web page, or the  IBM Support Assistant web page, to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0003E

An object cannot be used because its state is '{0}'.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the  WebSphere MQ support Web page, or the  IBM Support Assistant web page, to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0004E

The specified 'Timeout' value '{0}' is out of range.

Explanation

The value is out of range, it must be greater than or equal to 'TimeSpan.Zero'.

Response

Specify a value which is in range or, to disable Timeout, specify a 'TimeSpan.MaxValue' value.

WCFCH0005E

The operation did not complete within the specified time of '{0}' for endpoint address '{1}'.

Explanation

A timeout occurred.

Response

Investigate the cause for the timeout.

WCFCH0006E

The parameter '{0}' is not of the expected type '{1}'

Explanation

A parameter with an unexpected type has been passed to a method call.

Response

Review the exception stack trace for further information.

WCFCH0007E

The parameter '{0}' must not be null.

Explanation

A method has been called with a required parameter set to a null value.

Response

Modify the application to provide a value for this parameter.

WCFCH0008E

An error occurred while processing an operation for endpoint address '{0}'.

Explanation

The operation failed to complete.

Response

Review the linked exceptions and stack trace for further information.

WCFCH0101E-0200E: URI Properties messages

Use the following information to understand WCFCH0101E-0200E URI properties messages.

WCFCH0101E

The endpoint URI must start with the valid character string '{0}'.

Explanation

The endpoint URI is incorrect, it must start with a valid character string.

Response

Specify an endpoint URI which starts with a valid character string.

WCFCH0102E

The endpoint URI must contain a '{0}' parameter with a value.

Explanation

The endpoint URI is incorrect, a parameter and its value are missing.

Response

Specify an endpoint URI with a value for this parameter.

WCFCH0103E

The endpoint URI must contain a '{0}' parameter with a value of '{1}'.

Explanation

The endpoint URI is incorrect, the parameter must contain the correct value.

Response

Specify an endpoint URI with a correct parameter and value.

WCFCH0104E

The endpoint URI contains a '{0}' parameter with an invalid value of '{1}'.

Explanation

The endpoint URI is incorrect, a valid parameter value must be specified.

Response

Specify an endpoint URI with a correct value for this parameter.

WCFCH0105E

The endpoint URI contains a '{0}' parameter with an invalid queue or queue manager name.

Explanation

The endpoint URI is incorrect, a valid queue and queue manager name must be specified.

Response

Specify an endpoint URI with valid values for the queue and the queue manager.

WCFCH0106E

The '{0}' property is a required property and must appear as the first property in the endpoint URI.

Explanation

The endpoint URI is incorrect, a parameter is either missing or in the wrong position.

Response

Specify an endpoint URI which contains this property as the first parameter.

WCFCH0107E

The property '{1}' cannot be used when the binding property is set to '{0}'.

Explanation

The endpoint URI connectionFactory parameter is incorrect, an invalid combination of properties has been used.

Response

Specify an endpoint URI connectionFactory which contains a valid combination of properties or binding.

WCFCH0109E

Property '{1}' must also be specified when property '{0}' is specified.

Explanation

The endpoint URI connectionFactory parameter is incorrect, it contains an invalid combination of properties.

Response

Specify an endpoint URI connectionFactory which contains a valid combination of properties.

WCFCH0110E

Property '{0}' has an invalid value '{1}'.

Explanation

The endpoint URI connectionFactory parameter is incorrect, the property does not contain a valid value.

Response

Specify an endpoint URI connectionFactory which contains a valid value for the property.

WCFCH0111E

The value '{0}' is not supported for the binding mode property. XA operations are not supported.

Explanation

The endpoint URI connectionFactory parameter is incorrect, the binding mode is not supported.

Response

Specify an endpoint URI connectionFactory which contains a valid value for the binding mode.

WCFCH0112E

The endpoint URI '{0}' is badly formatted.

Explanation

The endpoint URI must follow the format described in the documentation.

Response

Review the endpoint URI to ensure that it contains a valid value.

WCFCH0201E-0300E: Factory/Listener messages

Use the following information to understand WCFCH0201E-0300E factory/listener messages.

WCFCH0201E

Channel shape '{0}' is not supported.

Explanation

The users application or the WCF service contract has requested a channel shape which is not supported.

Response

Identify and use a channel shape which is supported by the channel.

WCFCH0202E

'{0}' MessageEncodingBindingElements have been specified.

Explanation

The WCF binding configuration used by an application contains more than one message encoder.

Response

Specify no more than 1 MessageEncodingBindingElement in the binding configuration.

WCFCH0203E

The endpoint URI address for the service listener must be used exactly as provided.

Explanation

The binding information for the endpoint URI address must specify a value of 'Explicit' for the 'listenUriMode' parameter.

Response

Change the parameter value to 'Explicit'.

WCFCH0204E

SSL is not supported for managed client connections [endpoint URI: '{0}'].

Explanation

The endpoint URI specifies an SSL connection type which is only supported for unmanaged client connections.

Response

Modify the channels binding properties to specify an unmanaged client connection mode.

WCFCH0301E-0400E: Channel messages

Use the following information to understand WCFCH0301E-0400E channel messages.

WCFCH0301E

The URI scheme '{0}' is not supported.

Explanation

The requested endpoint contains a URI scheme which is not supported by the channel.

Response

Specify a valid scheme for the channel.

WCFCH0302E

The received message '{0}' was not a JMS bytes or a JMS text message.

Explanation

A message has been received but it is not of the correct type. It must be either a JMS bytes message or a JMS text message.

Response

Check the origin and contents of the message and determine the cause for it being incorrect.

WCFCH0303E

'ReplyTo' destination missing.

Explanation

A reply cannot be sent because the original request does not contain a 'ReplyTo' destination.

Response

Investigate the reason for the missing destination value.

WCFCH0304E

The connection attempt to queue manager '{0}' failed for endpoint '{1}'

Explanation

The queue manager could not be contacted at the given address.

Response

Review the linked exception for further details.

WCFCH0305E

The connection attempt to the default queue manager failed for endpoint '{0}'

Explanation

The queue manager could not be contacted at the given address.

Response

Review the linked exception for further details.

WCFCH0306E

An error occurred while attempting to receive data from endpoint '{0}'

Explanation

The operation could not be completed.

Response

Review the linked exception for further details.

WCFCH0307E

An error occurred while attempting to send data for endpoint '{0}'

Explanation

The operation could not be completed.

Response

Review the linked exception for further details.

WCFCH0308E

An error occurred while attempting to close the channel for endpoint '{0}'

Explanation

The operation could not be completed.

Response

Review the linked exception for further details.

WCFCH0309E

An error occurred while attempting to open the channel for endpoint '{0}'

Explanation

The operation could not be completed.

Response

The endpoint might be down, unavailable, or unreachable, review the linked exception for further details.

WCFCH0310E

The timeout '{0}' was exceeded while attempting to receive data from endpoint '{0}'

Explanation

The operation did not complete in the time allowed.

Response

Review the system status and configuration and increase the timeout if required.

WCFCH0311E

The timeout '{0}' was exceeded while attempting to send data for endpoint '{0}'

Explanation

The operation did not complete in the time allowed.

Response

Review the system status and configuration and increase the timeout if required.

WCFCH0312E

The timeout '{0}' was exceeded while attempting to close the channel for endpoint '{0}'

Explanation

The operation did not complete in the time allowed.

Response

Review the system status and configuration and increase the timeout if required.

WCFCH0313E

The timeout '{0}' was exceeded while attempting to open the channel for endpoint '{0}'

Explanation

The operation did not complete in the time allowed.

Response

The endpoint might be down, unavailable, or unreachable, review the system status and configuration and increase the timeout if required.

WCFCH0401E-0500E: Binding messages

Use the following information to understand WCFCH0401E-0500E binding messages.

WCFCH0401E

No context.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the IBM Support Portal for WebSphere MQ (see https://www.ibm.com/support/home/product/P439881V74305Y86/IBM_MQ), or the IBM Support Assistant (at https://www.ibm.com/support/home/product/C100515X13178X21/other_software/ibm_support_assistant), to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0402E

Channel type '{0}' is not supported.

Explanation

The users application or the WCF service contract has requested a channel shape which is not supported.

Response

Identify and use a channel shape which is supported by the channel.

WCFCH0403E

No exporter.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the IBM Support Portal for WebSphere MQ (see https://www.ibm.com/support/home/product/P439881V74305Y86/IBM_MQ), or the IBM Support Assistant (at https://www.ibm.com/support/home/product/C100515X13178X21/other_software/ibm_support_assistant), to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0404E

The WS-Addressing version '{0}' is not supported.

Explanation

The addressing version specified is not supported.

Response

Specify an addressing version which is supported.

WCFCH0405E

No importer.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the IBM Support Portal for WebSphere MQ (see https://www.ibm.com/support/home/product/P439881V74305Y86/IBM_MQ), or the IBM

Support Assistant (at https://www.ibm.com/support/home/product/C100515X13178X21/other_software/ibm_support_assistant), to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0406E

Endpoint 'Binding' value missing.

Explanation

An internal error has occurred.

Response

Use the standard facilities supplied with your system to record the problem identifier and to save any generated output files. Use either the IBM Support Portal for WebSphere MQ (see https://www.ibm.com/support/home/product/P439881V74305Y86/IBM_MQ), or the IBM Support Assistant (at https://www.ibm.com/support/home/product/C100515X13178X21/other_software/ibm_support_assistant), to see if a solution is already available. If you cannot find a match, contact your IBM support center. Do not discard these files until the problem has been resolved.

WCFCH0501E-0600E: Binding properties messages

Use the following information to understand WCFCH0501E-0600E binding properties messages.

WCFCH0501E

The binding property '{0}' has an invalid value '{1}'.

Explanation

An invalid value has been specified for a binding property.

Response

Specify a valid value for the property.

WCFCH0601E-0700E: Async operations messages

Use the following information to understand WCFCH0601E-0700E async operations messages.

WCFCH0601E

The async result parameter '{0}' object is not valid for this call.

Explanation

An invalid async result object has been provided.

Response

Specify a valid value for the parameter.

Index

Special characters

% character in RACF profiles 522

Numerics

0Cx abend 1270

5C6 abend

associated reason codes 1270

diagnostic information 1269

6C6 abend

associated reason codes 1270

diagnostic information 1269

A

abend

0C4 1181, 1270

0C7 1270

5C6 1269

6C6 1269

AICA 1298

application option of SSM entry 358

ASRA 1270

CICS 1271

IMS 1271

internal error 1269

introduction 1182

no dump taken 1334

program 1270

severe error 1269

starting after 252, 329

subsystem action 1269

U3042 361

z/OS 1271

abend code

description 1269

in dump title 1289

ABEND keyword 1329, 1334

abnormal termination 1269

abnormal termination, restarting

after 322

access

if incorrect 589

restricting by using alias queues 535

access control 432, 704, 755

API exit 760

authority to administer WebSphere

MQ 418

authority to work with WebSphere

MQ objects 422

introduction 368

user written message exit 760

user written security exit 758

access method services (AMS)

commands 299

deleting damaged BSDS 1372

new active log definition 295

renaming damaged BSDS 1372

REPRO 299

access settings 705, 709, 711

accessing CRLs

IBM i 700

Java client and JMS 703

queue manager 699

WebSphere MQ MQI client 701

Windows 700

accounting

data 1137

introduction 1114

message

format 981

message manager 1145

queue level 1146

rules for data collection 1116

sample SMF records 1146

SMF trace 1119

thread level 1146

accounting monitoring

MQI messages 985

queue messages 996

action queue manager 266

Active Directory

specifying that an MQI channel uses

SSL 398

active log

adding 1364

data set

copying 294

date 293

defining in BSDS 295

delays in off-loading 1364

deleting from BSDS 295

enlarging 295

format data sets 247

increasing the active log 295

increasing the size of the active

log 295

out of space 1364

printing 246

problem

both copies are damaged 1361

delays in off-loading 1364

dual logging lost 1359

log stopped 1360

one copy is damaged 1361

out of space 1364

read I/O errors 1362

short of space 1364

write I/O errors 1361

recording existing in BSDS 295

recovery plan, problems 1359

short of space 1364

status 293

stopped data set effect 1361, 1364

time 293

active threads 344

active units of work 344

activity report

activity report message data 935

format 927, 928

activity report (*continued*)

message data, operation specific

content 947

message descriptor 929

structure 928

activity reports

application control 891

controlling activity recording 891

queue manager control 891

requesting 891

use for 889

adding 64-bit integer items 53

adding byte string filter items 53

adding byte string items 53

adding character-string items 53

adding data items to bags 52

adding inquiry command 53

adding integer filter items 53

adding integer items 53

adding new active log 1364

adding string filter items 53

address space

abend 251, 327

canceling for WebSphere MQ 253,

330

channel initiator, dumping 1277

display list of active 1280

finding the identifier 1283

in dump formatting 1281

WebSphere MQ, dumping 1277

address space user ID 565, 568

addresses

IP

blocking 733, 734

administering by writing programs 272

administration

authority 751

description of 1, 186

introduction to 183

local, definition of 4, 184

MQAI, using 18

MQSC commands 76

PCF commands 185

remote administration, definition

of 4, 184

remote objects 110

using PCF commands 185

using the WebSphere MQ

Explorer 59

WebSphere MQ script (MQSC)

commands 184

writing Eclipse plug-ins 71

administration bag 48

administration programs 272, 273

administrative topic

changing queue attributes, commands

to use 97

deleting 98

administrative topics

copying an administrative topic

definition 98

- Adopting an MCA, information on 1310
- Advanced Message Security 439
- AIX operating system
 - creating and managing groups 468
 - MQAI support 18
 - tracing 1236, 1238
- alert monitor application, using 72
- algorithms for queue service interval events 848
- alias queues
 - command resource checking 561
 - DEFINE QALIAS command 93
 - defining alias queues 93, 174
 - remote queues as queue manager aliases 120
 - reply-to queues 120
 - restricting access using 535
 - security 535, 538
 - undelivered messages 538
 - working with alias queues 93, 174
- aliases
 - queue manager aliases 120
 - working with alias queues 93
- ALL attribute of DISPLAY SECURITY 580
- ALTER BUFFPOOL 311, 312
- alter queue attributes, security 551
- ALTER SECURITY command 577
- alternate user authority
 - introduction 424
 - server application 760
- alternate user ID
 - distributed queuing 566
 - intra-group queuing 567
- alternate user security 429
 - description 547
 - implementing 462
- alternate-user authority 518, 757
- alternative site recovery 323, 331, 335
- AMQ message 1338
- AMQA000000 work management object 190
- AMQAJRN work management object 190
- AMQAJRNMSG work management object 190
- AMQALMPX task 189
- AMQCLMAA task 189
- AMQCRC6B work management object 190
- AMQCRS6B task 189
- AMQCRSTA task 189
- AMQFCXBA task 189
- AMQFQPUB task 189
- AMQLAA0 work management object 190
- AMQPCSEA task 189
- AMQRCMLA task 189
- AMQRFCD4 work management object 190
- AMQRMPPA task 189
- AMQRRMFA task 189
- amqsail.c, sample programs 35
- amqsaiq.c, sample programs 20
- amqsaiem.c, sample programs 25
- amqsailq.c, sample programs 42
- AMQZFUMA task 189
- AMQZLAA0 task 189
- AMQZLSA0 task 189
- AMQZMGR0 task 189
- AMQZMUC0 task 189
- AMQZMUF0 task 189
- AMQZMUR0 task 189
- AMQZXMA0 task 189
- AMQZXMA0 work management object 190
- AMS (access method services)
 - commands 299
 - deleting damaged BSDS 1372
 - new active log definition 295
 - renaming damaged BSDS 1372
 - REPRO 299
- analyzing dumps 1288
- APAR 1172, 1328
 - definition 1351
 - number 1354
 - raising 1354
- API-resource security 430
 - quick reference 540
 - RESLEVEL 561
- APPC
 - security 590
- APPC channels, user IDs used 572
- application
 - running slowly 1185
 - stopped 1174
- application level security
 - comparison with link level security 435
 - introduction 439
 - providing your own 440
- application program
 - issuing commands from 272
- application programming errors, examples 1181
- application programming errors, examples of 1161
- application queues
 - creating and copying, restrict access to 518
 - defining application queues for triggering 176
- archive log
 - date 293
 - deleting 289
 - deleting information from the BSDS 297
 - discarding records 289
 - password, changing 297
 - printing 246
 - problem
 - allocation problems 1365
 - insufficient DASD for off-load 1366
 - read I/O errors during restart 1368
 - write I/O errors during off-load 1366
 - recording in BSDS 296
 - recovery plan 1365
 - restarting 288
 - time 293
- ARCHIVE LOG command 286
- archiving
 - ARCHIVE LOG command 286
 - using SET commands 288
- ARL 695
- ARM (automatic restart manager) security 590
- ARM (Automatic Restart Manager)
 - activating a policy 341
 - clusters 342
 - couple data sets 340
 - defining a policy 341
 - introduction 339
 - LU 6.2 343
 - network considerations 342
 - policy sample 340
 - queue-sharing groups 342
 - registering with 341
 - TCP/IP 342
- ASID in dump title 1289
- ASRA abend 1270
- asymmetric cryptography algorithm 369
- ATB message 1338
- ATR message 1338
- attributes
 - changing administrative topic attributes 97
 - changing local queue attributes 90, 100, 172
 - channel definition commands PUTAUT 802
 - LIKE attribute, DEFINE command 89, 98, 101
 - queue manager 85
- audit, security 586
- auditing RESLEVEL 587
- authentication 613
 - API exit 694
 - application level security service, example 439
 - digital signature 381
 - introduction 367
 - link level security service, example 437
 - SNA LU 6.2
 - conversation level authentication 456
 - session level authentication 454
 - SSL 378
 - SSPI channel exit program 487
 - user written message exit 693
 - user written security exit 692
- authentication information
 - commands 241
 - commands</ idxterm> 241
- authentication information object (AUTHINFO)
 - manipulating 703
- authentication information objects, disposition 265
- authority
 - administration 751
 - alternate-user 518, 757
 - context 758
 - context authority 518
 - authority checking (PCF)
 - HP Integrity NonStop Server 14
 - HP OpenVMS 14

- authority checking (PCF) *(continued)*
 - IBM i 14
 - UNIX systems 14
 - Windows NT 14
 - z/OS systems 14
- authority checks
 - alternate user authority 424
 - CL command in Group 2 423
 - command resource security 420
 - command security 420
 - message context 425
 - MQCLOSE call 423
 - MQCONN call 423
 - MQCONN call 423
 - MQOPEN call 423
 - MQPUT1 call 423
 - PCF command 423
 - z/OS 420
- authority profiles
 - working without 517
- Authority Revocation List (ARL) 695
- authority to administer WebSphere MQ 418
- authority to work with WebSphere MQ objects 422
- authorization service 426
 - specifying the installed 516
- authorizations
 - MQI 472, 502
 - specification tables 471, 502
- automatic definition of channels 115
- automatic restart manager (ARM)
 - security 590
- Automatic Restart Manager (ARM)
 - activating a policy 341
 - clusters 342
 - couple data sets 340
 - defining a policy 341
 - introduction 339
 - LU 6.2 343
 - network considerations 342
 - policy sample 340
 - queue-sharing groups 342
 - registering with 341
 - TCP/IP 342
- automating starting of CKTI 593

B

- backing up page sets 307
- backup
 - data 200
 - introduction 195
 - journals 200
 - media images 199
 - performance 233
 - using journals 197
- backup, of coupling facility
 - structures 315
- bags
 - adding 64-bit integer items to 53
 - adding byte string filter items to 53
 - adding byte string items to 53
 - adding character-string items to 53
 - adding data items to 52
 - adding inquiry command to 53
 - adding integer filter items to 53

- bags *(continued)*
 - adding integer items to 53
 - adding string filter items to 53
 - changing 64-bit integer items
 - within 55
 - changing byte string filter items
 - within 55
 - changing byte string items within 55
 - changing character-string items
 - within 55
 - changing information within 54
 - changing integer filter items
 - within 55
 - changing integer items within 55
 - changing string filter items within 55
 - converting 56
 - converting to PCF messages 56
 - creating 50
 - creating and deleting 50
 - deleting 50
 - inquiring within 54
 - putting 51
 - receiving 51
 - types of 48
 - using 48
- batch
 - abend 1182, 1270
 - application loop 1298
 - wait 1296
- batch connections
 - definition 530
 - RESLEVEL 563
 - security 530, 531
 - security checking 563
 - user IDs, security checking 569
- blank fields in operations and control panels 262
- blank user IDs 576
- block cipher algorithm 369
- blocking
 - IP addresses 733, 734
 - user IDs 735
- blocking access to a channel
 - for a client asserted user ID 738
 - for an SSL DN 738
 - from a remote queue manager 737
- BMP (batch message program) 358
- bootstrap data set (BSDS)
 - adding an active log 295
 - adding an archive log 296
 - both copies are damaged 1369
 - change 246
 - change log inventory utility (CSQJU003) 294
 - changing for active logs 295
 - changing for archive logs 296
 - changing log inventory utility (CSQJU003) 297
 - commands 241
 - deleting active log information 295
 - deleting archive log information 297
 - determining log inventory contents 293
 - does not agree with log 1369
 - error while opening 1368
 - errors 1368
 - I/O error 1371

- bootstrap data set (BSDS) *(continued)*
 - managing 292, 297
 - out of synchronization 1371
 - printing 246
 - recover log inventory 290
 - recovery 298
 - restoring from the archive log 299
 - single recovery 298
 - time stamps 293
 - unequal time stamps 1370
- bootstrap data sets (BSDSs) 422
- browsing queues 91
- BSDS (bootstrap data set)
 - adding an active log 295
 - adding an archive log 296
 - both copies are damaged 1369
 - change 246
 - change HLQ for the logs and BSDS 297
 - change log inventory utility (CSQJU003) 294
 - changing for active logs 295
 - changing for archive logs 296
 - changing log inventory utility (CSQJU003) 297
 - commands 241
 - deleting active log information 295
 - deleting archive log information 297
 - determining log inventory contents 293
 - does not agree with log 1369
 - error while opening 1368
 - errors 1368
 - I/O error 1371
 - managing 292, 297
 - out of synchronization 1371
 - printing 246
 - recover log inventory 290
 - recovery 298
 - restoring from the archive log 299
 - single recovery 298
 - time stamps 293
 - unequal time stamps 1370
- BSDSs 422
- buffer manager statistics 1129
- buffer pool
 - size 1130
- buffer pool size 1266
- buffer pool, associating with page sets 301
- buffer pools
 - altering 311, 312
 - commands 241
 - defining 245
 - managing 311
- buffers
 - number in buffer pool 311, 312
- building commands
 - rules for 77
- built-in formats, data conversion 121

C

- CA 371
- CA certificate
 - adding, UNIX 635, 648

- CA certificate (*continued*)
 - creating for testing
 - IBM i 623
 - extracting 646
- calls
 - mqAddByteString 53
 - mqAddByteStringFilter 53
 - mqAddInquiry 53
 - mqAddInteger 53
 - mqAddInteger64 53
 - mqAddIntegerFilter 53
 - mqAddString 53
 - mqAddStringFilter 53
 - mqBagToBuffer 56
 - mqBufferToBag 56
 - mqClearBag 55
 - mqCreateBag 50
 - mqDeleteBag 50
 - mqDeleteItem 57
 - mqExecute 58
 - mqPutBag 9
 - mqSetByteString 55
 - mqSetByteStringFilter 55
 - mqSetInteger 55
 - mqSetInteger64 55
 - mqSetIntegerFilter 55
 - mqSetString 55
 - mqSetStringFilter 55
 - mqTruncateBag 55
- canceling WebSphere MQ address
 - space 253, 330
- CARTs 249
- CBC message 1338
- ccsid.tbl, data conversion 121
- CEDF
 - introduction 1275
- CEE message 1338
- CEMT INQ TASK 1296
- certificate
 - chain 375
 - ensuring availability
 - z/OS 676
 - exporting, IBM i 626
 - importing, IBM i 627
 - untrustworthy
 - in CRL 695
 - when changes are effective
 - IBM i 630
 - UNIX 638
 - z/OS 677
- Certificate Authority
 - digital certificates 371
 - introduction 373
 - public key infrastructure (PKI) 376
 - working with Certificate Revocation Lists 695
- certificate labels, understanding the
 - requirements of 783
- certificate name filter
 - defining 681
- Certificate Name Filters (CNFs)
 - setting up on z/OS 680
 - using on z/OS 680
- certificate policy 660
 - basic and standard 664
- certificate revocation list (CRL) 701
- Certificate Revocation List (CRL)
 - accessing
 - IBM i 700
 - Java client and JMS 703
 - queue manager 699
 - WebSphere MQ Explorer 700
 - WebSphere MQ MQI client 701
 - working with 695
- certificate store
 - setting up on IBM i 620
 - stashing password
 - IBM i 622
 - Windows key repository 390
- certificates
 - expiry 375
 - untrustworthy
 - introduction 375
- certification path 375
- CF structure
 - adding 321
 - backup and recovery 315
 - disaster recovery 334
 - load balancing 316
 - managing 321
 - moving a queue to another 316
 - recovering from failure 1376
 - removing 321
- CF structures
 - commands 241
- CFRM
 - activation 1300
 - policy data set 1300
- change log inventory utility 246
- change log inventory utility (CSQJU003)
 - adding new active log 295, 1364
 - change BSDS 293, 294
 - changes for active logs 295
 - changes for archive logs 296
 - functions
 - ARCHIVE 299
 - NEWLOG 295, 296
 - time stamp in BSDS 1371
- Change Security command 577
- CHANGE SUBSYS, command of
 - IMS 353, 357
- change team 1351
- changed application 1179
- changing
 - administrative topic attributes 97
 - CCSID 121
 - local queue attributes 90, 100, 172
 - queue manager attributes 85
- changing 64-bit integer items within data
 - bags 55
- changing byte string filter items within
 - data bags 55
- changing byte string items within data
 - bags 55
- changing character-string items within
 - data bags 55
- changing information within data
 - bags 54
- changing integer filter items within data
 - bags 55
- changing integer items within data
 - bags 55
- changing key repository
 - IBM i 622
 - UNIX 637
- changing string filter items within data
 - bags 55
- Changing the HLQ for BSDS
 - BSDS 297
- Changing the HLQ for logs
 - BSDS 297
- channel
 - events
 - controlling 841
 - refuses to run 1308
 - running 1302
 - security 584, 590
 - startup negotiation errors 1308
 - switching 1313
 - user ID 566
 - channel authentication records 586
 - channel control error messages 1306
 - channel disposition 265
 - channel event
 - queue 836
 - channel exit programs
 - introduction 440
 - message exit
 - introduction 442
 - providing your own link level security 438
 - receive exit
 - introduction 442
 - providing your own link level security 438
 - security exit
 - introduction 441
 - providing your own link level security 438
 - SSPI 487
 - send exit
 - introduction 442
 - providing your own link level security 438
 - SSPI 487
 - channel exits
 - security 449, 452
 - channel initiator
 - commands 241
 - connection security 532
 - defining objects at startup 245
 - START CHANNEL commands 421
 - SYSTEM.* queue security 539
 - user IDs used 571
 - channel initiator dump
 - formatting 1287
 - channel initiator trace 1244, 1250
 - channel initiator, checking 1173
 - channel initiator, restarting with
 - ARM 342
 - channel listener
 - commands 241
 - channel name
 - meaning 1147
 - channel problems 1301
 - channel refuses to run 1308
 - channel startup negotiation errors 1308
 - channel statistics message data 1024
 - channel user ID, security 590

- channels
 - administering a remote queue manager from a local one 112
 - auto-definition of 115
 - commands 241
 - defining channels for remote administration 114
 - description of 110
 - escape command authorizations 475, 505
 - exits 449, 452
 - granting administrative access 716, 724
 - preparing channels for remote administration 113
 - remote queuing 110
 - security 447
 - starting 115
- character code sets, updating 121
- checkpoints 200
- CICS
 - abend 1182, 1270
 - abend code 1271
 - address space user ID, security checking 568
 - application loop 1298
 - connection security 531
 - execution diagnostic facility 1275
 - monitoring facility 1121
 - performance considerations 1294
 - transaction wait 1296
 - units of recovery 346
 - user IDs
 - for security 568
 - security checking 569
- CICS adapter
 - restart, what happens 345
 - security
 - authorization 593
 - for transactions 592
 - PLTPI 593
 - PLTSD 593
 - terminal user IDs 593
 - user IDs 593
 - statistics 1116
 - trace 1252
 - trace entries 1252
 - transaction services
 - security support 592
 - user IDs for security 593
- CICS bridge
 - security 594
- CICS connection
 - RESLEVEL 564
 - security
 - checking 564
 - user IDs for security 564
- CICS region
 - processor activity 1298
- cipher algorithm
 - block 369
 - stream 369
- cipher strength 369
- CipherSpec
 - alternatives for specifying 777, 800
 - introduction 380
- CipherSpec (*continued*)
 - obtaining information using WebSphere MQ Explorer 776, 799
 - specifying for WebSphere MQ MQI client 777, 800
 - working with 399
- CipherSuite
 - introduction 380
 - specifying for Java client and JMS 778, 801
- ciphertext 369
- CKBM security 593
- CKCN security 593
- CKDL security 593
- CKDP security 593
- CKQC
 - security 593
- CKRS security 593
- CKRT security 593
- CKSD security 593
- CKSQ security 593
- CKTI MCA transaction 593
- CKTI transaction
 - automating starting of 593
 - security 593
- CL commands
 - accessing WebSphere MQ objects 422
 - creating WebSphere MQ objects 170
 - Group 2
 - authority checks 423
 - definition 419
 - introduction 419, 494
 - starting a local queue manager 170
- CLASS, specifying 1246
- clearing a bag 55
- clearing a local queue 90, 172
- client asserted user ID
 - blocking channel access 738
- client channel definition files, create 247
- client channel definition table
 - specifying that an MQI channel uses SSL 398
- clients, problem determination 1313
- clients, security 573
- close options, dynamic queues 537
- cluster
 - finding a message 1303
 - keeping secure 801
 - message not on queue 1300
 - preventing queue managers joining 803
 - queue 1176
- cluster support
 - security considerations 592
- clusters
 - cluster membership, the WebSphere MQ Explorer 67
 - commands 241
 - description of 112
 - remote queuing 110
 - security 459
 - showing and hiding, WebSphere MQ Explorer 66
- clusters and ARM 342
- clusters, use of
 - security 801
- CMDSCOPE, user messages from
 - commands with 282
- CMQ modules 1272
- CNF 680
 - defining 681
- code, return 1172
- coded character sets, specifying 121
- cold start 325, 338
- collecting documentation 1352
- command
 - events
 - controlling 842
 - no response from 1184
 - problems 1183
 - queue 7
 - to take a dump 1277
 - command and response tokens 249
 - command bag 48
 - command events 867
 - command facility
 - operations and control panels 269
 - using 269
 - command files 81
 - command line, using with operations and control panels 268
 - command prefix string (CPF) 249
 - command prefix strings 248
 - command queues
 - command server status 116
 - mandatory for remote administration 113
 - command resource security 420
 - command responses 248
 - command scope 266
 - command security 420
 - implementing 462
 - command server 1184
 - displaying status 116
 - remote administration 116, 187
 - restart 275
 - sending commands to 275
 - starting 274
 - starting a command server 116
 - stopping 275
 - stopping a command server 116
 - command server reply message descriptor 277
 - command server, commands 241
 - command sets
 - MQSC commands 76
 - command string
 - entering quotes 78
 - preserving case 78
- commands
 - action queue manager 266
 - choosing a queue manager 266
 - command prefix string (CPF) 249
 - command scope 266
 - directing to another queue manager 240
 - DISPLAY 1115
 - DISPLAY CONN 1275
 - disposition 238
 - dmpmqaut 705, 709
 - dspmqaout 711
 - examples of 279

- commands (*continued*)
 - implementing resource security
 - checking 462
 - initialization 245
 - issuing 238
 - from CSQUTIL 249
 - from system-command input queue 272
 - introduction 248
 - issuing MQSC commands using an ASCII file 76
 - no reply to 279
 - operator 249
 - processor 274
 - remote queue manager 274
 - resource checking summary table 551
 - resource security profiles 560
 - rules for building 77
 - rules for using 77
 - runmqsc command, to issue MQSC commands 76
 - scope 240
 - security profiles 551
 - setmqaut 704
 - STOP QMGR 253, 330
 - synonym 78
 - target queue manager 266
 - user messages
 - from DEFINE 279
 - from DELETE 280
 - from DISPLAY QLOCAL 280
 - verifying MQSC commands 81
 - commands, PCF 185
 - complete queue manager (data and journals), Restoring a 203
 - completion code in dump title 1289
 - component identifier
 - in dump title 1289
 - list of 1344
 - component in dump title 1289
 - component-identifier keyword 1333
 - concepts and terminology 19
 - concurrent threads, limiting 1267
 - confidentiality
 - application level security service, example 439
 - cryptographic 369
 - introduction 368
 - link level security service, example 437
 - SNA LU 6.2 session level cryptography 453
 - SSL 378
 - user written security exit 780
 - configuration
 - events
 - controlling 842
 - using distributed queuing on WebSphere MQ
 - object security 430
 - object security UNIX systems 431
 - object security Windows systems 432
 - user IDs across systems 432
 - configuration events 863
 - configuring LDAP servers 697
 - configuring, Secure Sockets Layer (SSL) 584
 - connecting applications
 - using distributed queuing on distributed platforms
 - object security 430
 - object security UNIX systems 431
 - object security Windows systems 432
 - user IDs across systems 432
 - connection
 - switching 1313
 - connection security 429
 - implementing 462
 - IMS bridge 597
 - setting up 740
 - connections
 - controlling IMS 353
 - displaying 344
 - displaying details of IMS 358
 - monitoring the activity on 358
 - profiles for security 530
 - starting from IMS 354
 - stopping from IMS 353
 - to IMS, monitoring activity 358, 360
 - connectivity
 - removing queue manager 730
 - console security 568
 - console, issuing commands from 249
 - content of dump 1288
 - context authority 518, 758
 - context security 429
 - implementing 462
 - profiles 549
 - control blocks, display 1281
 - control commands 418
 - runmqsc, using interactively 79
 - control Language, IBM i 1
 - controlling
 - channel events 841
 - command events 842
 - configuration events 842
 - events 840
 - logger events 842
 - performance events 842
 - queue manager events 841
 - controlling activity recording 891
 - controlling log compression 288
 - conversion failure, problem determination 1310
 - converting bags and buffers 56
 - converting bags to PCF messages 56
 - converting PCF messages to bag form 56
 - copy a page set 247
 - copy a queue 247
 - COPY object disposition 266
 - correlation identifier 858
 - CorrelId 1300
 - CorrelId field, administration programs 277
 - CorrelId parameter 1185
 - corrupted message 1304
 - COUNT field, user messages 278
 - counting data items 57
 - couple data sets, ARM 340
 - coupling facility (CF)
 - adding a structure 321
 - backup and recovery 315
 - disaster recovery 334
 - load balancing 316
 - managing 321
 - moving a queue to another structure 316
 - recovering from failure 1376
 - removing a structure 321
 - coupling facility manager SMDS statistics 1137
 - coupling facility manager statistics 1136
 - Coupling Facility manager statistics 1136
 - Coupling Facility structure
 - accessibility 1176
 - full 1299
 - Coupling Facility structures
 - commands 241
 - commands</ idxterm> 241
 - CPF (command prefix string) 249
 - creating
 - a transmission queue 120
 - process definition 176
 - creating a local queue, sample programs 20
 - creating data bags 50
 - creating WebSphere MQ objects 170
 - CRL 695, 701
 - CRL policy 660
 - basic and standard 666
 - cross reference data thread 1145
 - cryptographic hardware
 - configuring on IBM i 631
 - configuring on UNIX 655
 - cryptology
 - algorithm 369
 - introduction 369
 - CSECT
 - in dump title 1289
 - keyword 1336
 - offset in dump title 1289
 - CSECT keyword 1328
 - CSMT log 1296
 - CSQ1LOGP (log print utility)
 - finding start RBA with 348
 - CSQ1LOGP utility 246
 - CSQ2020E message 364
 - CSQ4BCX3 452
 - CSQ4BREC sample 298, 1372
 - CSQ5PQSG utility 247
 - CSQDQ5ST 1132
 - CSQI063E message 1359
 - CSQINP1
 - input data set 245
 - security 568
 - CSQINP1 data sets
 - authority to access 422
 - MQSC commands 421
 - CSQINP2
 - input data set 245
 - security 568

CSQINP2 data sets
 authority to access 422
 MQSC commands 421
 CSQINPX
 security 590
 CSQINPX data sets
 authority to access 422
 MQSC commands 421
 CSQINPX input data set 245
 CSQJ004I message 1359
 CSQJ030E message 1360
 CSQJ100E message 1368
 CSQJ102E message 1361, 1369
 CSQJ103E message 1365
 CSQJ105E message 1361
 CSQJ106E message 1362
 CSQJ107E message 1372
 CSQJ108E message 1372
 CSQJ110E message 1364
 CSQJ111A message 1364
 CSQJ114I message 1366
 CSQJ115E message 1365
 CSQJ120E message 1370
 CSQJ122E message 1371
 CSQJ124E message 1362
 CSQJ126E message 1371
 CSQJ128E message 1366
 CSQJ232E message 1361
 CSQJU003 (change log inventory utility)
 adding new active log 295, 1364
 change BSDS 293, 294
 changes for active logs 295
 changes for archive logs 296
 functions
 ARCHIVE 299
 NEWLOG 295, 296
 time stamp in BSDS 1371
 CSQJU003 utility 246
 CSQJU004 (print log map utility)
 BSDS time stamps 293
 listing BSDS contents using 293
 CSQJU004 utility 246
 CSQJUFMT utility 247
 CSQM201 message 347, 350
 CSQP004E message 1372
 CSQP018I message 259
 CSQP019I message 259
 CSQQTRMN
 starting 361
 stopping 361
 CSQQxxx messages 1377
 CSQUCVX utility 246
 CSQUDLQH utility 247, 422
 CSQUTIL 703
 RESLEVEL 563
 security checking 563
 SYSTEM.* queue security 539
 CSQUTIL utility 421
 CSQV086E message 1336
 CSQWDMP statement 1283
 CSQYASCP, 0C4 abend during
 startup 1181
 CSV message 1338
 CTL (IMS control region) 354, 358
 CURDEPTH
 DISPLAY 1185
 MAXDEPTH 1175

CURDEPTH (*continued*)
 value of 1303
 CURDEPTH, current queue depth 89
 current queue depth, CURDEPTH 89
 customizing
 security 587

D

D RRS 1297
 DAE (dump analysis and
 elimination) 1292
 data
 activity report 928
 backup 200
 response 11
 restoring 203
 trace-route message 954
 trace-route reply message 965
 data bags
 adding 64-bit integer items to 53
 adding byte string filter items to 53
 adding byte string items to 53
 adding character-string items to 53
 adding data items to 52
 adding inquiry command to 53
 adding integer filter items to 53
 adding integer items to 53
 adding string filter items to 53
 changing 64-bit integer items
 within 55
 changing byte string filter items
 within 55
 changing byte string items within 55
 changing character-string items
 within 55
 changing information within 54
 changing integer filter items
 within 55
 changing integer items within 55
 changing string filter items within 55
 converting 56
 converting to PCF messages 56
 creating 50
 creating and deleting 50
 deleting 50
 inquiring within 54
 putting 51
 receiving 51
 types of 48
 using 48
 data conversion
 application level security 440
 built-in formats 121
 ccsid.tbl, uses for 121
 converting user-defined message
 formats 121
 default data conversion 121
 introduction 121
 updating coded character sets 121
 data conversion exit utility 246
 Data Encryption Standard (DES)
 algorithm
 SNA LU 6.2 security services 453
 Message Authentication Code (MAC)
 SNA LU 6.2 conversation level
 authentication 458

Data Encryption Standard (DES)
 (*continued*)
 Message Authentication Code (MAC)
 (*continued*)
 SNA LU 6.2 session level
 authentication 454
 data integrity
 application level security service,
 example 439
 cryptography 369
 link level security service,
 example 437
 message digests 371
 SSL 378
 data items
 counting 57
 deleting 57
 filtering 53
 querying 53
 types of 51
 data manager statistics 1129
 data sets
 page set I/O error 1372
 RACF authorization 583
 restart on losing 325, 338
 security 582
 data sets, distribution of 1267
 data types, detailed description
 activity report
 MQCFH 934
 MQEPH 933
 MQMD 929
 trace-route message
 MQCFH 959
 MQEPH 958
 MQMD 954
 trace-route reply message
 MQCFH 966
 MQMD 965
 data-sharing group
 recovering from failure 1375
 resynchronizing with WebSphere
 MQ 1375
 Db2
 abend 1270
 adding a queue-sharing group 313
 perform tasks for queue-sharing
 groups 247
 recovering from system failure 1374
 removing a queue-sharing group 314
 resynchronizing with WebSphere
 MQ 1375
 wait 1297
 DB2 DISPLAY THREAD (*) 1297
 Db2 manager statistics 1132
 DB2 tables
 CSQ45STB job 1353
 DCM 392
 DDNS
 registration time for 1311
 dead letter queue handler utility
 (CSQUDLQH) 422
 dead-letter queue 1184
 finding out its name 280
 handler utility 247
 message header 1299, 1302
 problem determination 1307

- dead-letter queue (*continued*)
 - processing 1307
- dead-letter queue, security 538
- dead-letter queues
 - defining a dead-letter queue 88, 172
- debugging
 - common programming errors 1181
 - diagnostic aids 1265
- decipherment 369
- decryption 369
- default configuration, Windows systems 1
- default data conversion 121
- default transmission queues 120
- DEFINE commands, user messages 279
- DEFINE LOG command 1364
- defining
 - a model queue 95
 - alias queue 174
 - an alias queue 93
 - application queue for triggering 176
 - dead-letter queue 172
 - initiation queue 176
 - local queue 172
 - model queue 175
- defining objects at startup 245
- DELETE commands, user messages 280
- deleting
 - a local queue 90
 - a subscription 101
 - active information log from BSDS 295
 - administrative topic 98
 - archive logs 289, 290
 - IMS tpipes 366
- deleting a local queue 172
- deleting data items 57
- deleting data bags 50
- dependent region user ID, IMS 570
- dependent region, IMS
 - controlling connections 358
 - disconnecting from 360
- DEQUEUE TMEMBER, command of IMS 365
- DES 453
- describing the problem 1331
- determining current queue depth 89
- DFH message 1338
- DFS message 1338
- DFS3611 message 1378
- DFS555I message 1378
- diagnostic aids
 - channel initiator trace 1244
 - dumps 1276
 - GTF trace 1244
 - IBM internal trace 1244
 - introduction 1265
 - line trace 1244
 - SYS1.LOGREC records 1292
 - trace 1244
 - user parameter trace 1244
- diagnostics
 - Java 1262
- dial-up support 1311
- digital certificate
 - Certificate Authority 373
 - certificate chain 375
- digital certificate (*continued*)
 - content 372
 - Distinguished Name (DN) 373
 - introduction 371
 - key repository 390
 - label on IBM i 620
 - label on UNIX 632
 - label on Windows 632
 - label on z/OS 675
 - public key infrastructure (PKI) 376
 - SSL authentication 378
 - SSL handshake 386
- Digital Certificate Manager
 - IBM i 392
- digital certificates
 - expiry 375
 - untrustworthy 375
- digital signature
 - introduction 381
 - SSL integrity 378
- directories, queue manager 518
- disabling
 - channel events 841
 - command events 842
 - configuration events 842
 - events 840
 - logger events 842
 - performance events 842
 - queue manager events 841
- disabling connectivity to queue managers 730
- disabling remote access to queue managers 739
- disaster recovery 1312
 - queue manager 331
 - queue-sharing group 323, 334, 335
- discarded messages 277
- disconnecting
 - from IMS 360
- display
 - attributes of subscriptions 99
 - default object attributes 89, 97
 - dump title 1280
 - process definitions 108, 176
 - queue 1175, 1273
 - queue manager attributes 85
 - status of command server 116
 - system status 1173
- DISPLAY CF 1299
- DISPLAY CHSTATUS
 - BYTSENT keyword 1178
- DISPLAY CMDSERV 1184
- DISPLAY CONN command 1275
- DISPLAY DQM 1172, 1174
- DISPLAY OASN command of IMS 357
- DISPLAY QLOCAL 1178
- DISPLAY QMGR COMMANDQ 1175
- DISPLAY QUEUE 1303
 - CURDEPTH attribute 1184
 - MAXDEPTH attribute 1184
- DISPLAY SECURITY 529
- DISPLAY THREAD 1172, 1297
- displaying
 - function key settings 267
 - units of recovery in CICS 347
 - units of recovery in IMS 350, 356
- disposition (object) 238
- disposition, object 265
- Distinguished Name
 - blocking channel access 738
- Distinguished Name (DN)
 - filter on z/OS 681
 - introduction 373
- distributed queuing 1299
 - defining objects at startup 245
 - MCA user ID 571
 - RESLEVEL 566
 - security checking 566
 - user IDs used 571
- distributed queuing example 179
- distributed queuing using CICS ISC SYSTEM.* queue security 539
- distributed queuing, incorrect output 1155
- distributed queuing, message not on queue 1301
- dmpmqaut command 428
- DMQ message 1338
- DMQ modules 1272
- DN 373
- DNS (Domain Name System) 343
- DOC keyword 1328, 1329, 1341
- documentation
 - problems 1341
 - required for a PMR 1352
 - sending PMR documentation 1353
 - useful in problem determination 1274
- domain controller
 - security 489
- Domain Name System (DNS) 343
- DSN message 1338
- dspmqaout command 428
- DSPMQMAUT command 428, 516
- dspmqrtrc trace command 1236
- dual logging , losing 1359
- dump
 - analyzing 1288
 - content 1288
 - display 1281
 - format 1281
 - using line mode IPCS 1283
 - using the CSQWDMP statement 1283
 - using the panels 1279
 - managing the inventory 1280
 - not taken for an abend 1334
 - options 1277
 - printing 1288
 - processing
 - using IPCS in batch 1288
 - using line mode IPCS 1283
 - using the CSQWDMP statement 1283
 - using the dump display panels 1279
 - selecting 1280, 1283
 - summary portion 1277
 - suppression 1292
 - SYS1.LOGREC data 1280
 - taking 1277
 - title 1288
 - using the z/OS dump command 1277

- dump analysis and elimination (DAE) 1292
- dump inventory, managing 1280
- dumps
 - address space 1182
 - transaction 1182
- dynamic definition of channels 115
- dynamic queues
 - close options 537
 - security 536

E

- eavesdropping 369
- EDC message 1338
- editing namelists 272
- embedded header
 - activity report 933
 - trace-route message 958
- EMCS 249
- empty a queue 247
- enabling
 - activity recording 891
 - applications for activity recording 891
 - channel events 841
 - command events 842
 - configuration events 842
 - events 840
 - logger events 842
 - performance events 842
 - Queue Depth events 856
 - differences between nonshared and shared queues 858
 - queue manager events 841
 - queue managers for activity recording 891
 - queue managers for trace-route messaging 897
 - queue service interval events 849
- encipherment 369
- encryption
 - CipherSpecs 399
 - introduction 369
 - SSL confidentiality 378
- end-to-end security 439
- ending
 - interactive MQSC commands 79
- endmqtrc trace command 1234, 1236
- environment variables
 - MQ_USER_ID 433
 - MQS_TRACE_OPTIONS 1236
- error
 - at remote sites 1306
 - logs 1314
 - message from channel control 1306
 - on channels 836
 - on event queues 837
 - response 11
- error logs
 - description of 1229
- error messages
 - non-WebSphere MQ 1173
- error messages, MQSC commands 79
- error, user data 1268
- errors, hardware 1379
- Escape PCF commands 418

- escape PCFs 185
- ESM 420
- event 832
 - attribute setting 840
 - channel 836
 - command 838
 - configuration 838
 - controlling 840
 - controlling channel 841
 - controlling command 842
 - controlling configuration 842
 - controlling logger 842
 - controlling performance 842
 - controlling queue manager 841
 - data 845
 - enabling and disabling 840
 - instrumentation example 877
 - logger 838
 - message
 - data summary 838
 - messages
 - event queues 832
 - format 844
 - lost 843
 - null 858
 - unit of work 837
 - queue depth
 - Queue Depth High 855
 - Queue Depth Low 855
 - Queue Full 855
 - queue manager 834
 - queues
 - errors 837
 - names for 832
 - transmission 843
 - triggered 843
 - unavailable 843
 - use of 832
 - reporting 832
 - service interval 846
 - shared queues (WebSphere MQ for z/OS) 837
 - statistics
 - example 1 summary 851
 - example 2 summary 853
 - resetting 846
 - timer 848
 - transmission queues, as event queues 843
 - use for 830
- event identifier 1245
- event monitor, sample programs 25
- events
 - command 867
 - logger 869
- example
 - queue-sharing group security scenario 608
 - SMF statistics records 1125
- example ARM policy 340
- example output
 - SYSUDUMP 1290
- example recovery scenarios
 - active log problems
 - both copies are damaged 1361
 - delays in off-loading 1364
 - dual logging lost 1359

- example recovery scenarios (*continued*)
 - active log problems (*continued*)
 - log stopped 1360
 - one copy is damaged 1361
 - out of space 1364
 - read I/O errors 1362
 - short of space 1364
 - write I/O errors 1361
 - archive log problems
 - allocation problems 1365
 - insufficient DASD for off-load 1366
 - read I/O errors during restart 1368
 - write I/O errors during off-load 1366
 - BSDS problems
 - both copies are damaged 1369
 - BSDS recovery 298
 - does not agree with log 1369
 - error while opening 1368
 - I/O error 1371
 - out of synchronization 1371
 - unequal time stamps 1370
 - hardware problems 1379
 - IMS problems
 - application terminates 1378
 - cannot connect to WebSphere MQ 1377
 - IMS not operational 1378
 - page set problems
 - I/O error 1372
 - page set full 1373
- examples
 - creating a transmission queue 120
 - instrumentation event 877
 - queue depth events 859
 - queue service interval events 850
- exits
 - security 416
- expiry of digital certificates 375
- exporting certificate
 - z/OS 679
- extended console support 249
- external security manager (ESM) 420
- extract information from a page set 247
- eye-catcher strings, statistics trace 1124
- EZA message 1338
- EZB message 1338
- EZY message 1338

F

- failover
 - performance 233
- failure keywords 1329, 1330
- feedback, MQSC commands 79
- FFST (First Failure Support Technology) 1320
- FFST (first-failure support technology)
 - UNIX systems 1318
 - Windows NT 1315
- filtering data items 53
- finding archive log data sets to be deleted 290
- FIPS 393, 394, 396

- format
 - type 115 SMF records 1123
 - type 116 SMF records 1138
- format a page set 247
- format of accounting and statistics messages 981
- format of activity reports 928
- format of event messages 844
- format of trace-route message 954
- format of trace-route reply message 965
- formatting
 - channel initiator dump 1287
- formatting a dump
 - using line mode IPCS 1283
 - using the CSQWDMP statement 1283
 - using the panels 1279
- free keyword format 1329
- frequent I/O 1266
- FRR keyword 1337
- full administrative access
 - to queue manager 729
- function keys
 - changing namelists 272
 - operations and control panels 267
 - showing 267
 - using 267

G

- generic profile 428
- generic profiles 515
- generic profiles, OAM 705
- global accounting interval 1120
- global definitions
 - definition 238
 - manipulating 240
- GMQADMIN security class 521, 524
- GMQNLIST security class 521
- GMQPROC security class 521
- GMQQUEUE security class 521
- GMXADMIN security class 524
- group bag 48
- group class, security 521
- GROUP object disposition 266
- GROUP objects 238
- group objects, managing 320
- grouped messages 1303
- groups
 - security 756
- GRTMQMAUT command 516
 - example 428
 - introduction 419, 494
- gsk7ikm on UNIX 631
- gsk7ikm on Windows 631
- GSKit 8 615
- GTF
 - format identifier 1247
 - formatting 1247
 - identifying WebSphere MQ control blocks 1247
 - interpreting 1248
 - specifying the job name 1245
 - starting 1245
 - user parameter trace 1245
 - USR option 1245
- GTFTRACE command 1247

H

- HALT keyword 1329
- handshake, SSL 377
- hardware errors 1379
- hash function
 - CipherSpecs 399
 - overview 371
- header
 - activity report 934
 - SMF type 115 record 1123
 - SMF type 116 record 1138, 1141
 - trace-route message 959
 - trace-route reply message 966
 - WebSphere MQ messages 982
- help
 - operations and control panels 267
- High
 - events rules 848
- high level qualifier 297
- HLQ 297
- HP-UX
 - MQAI support for 18
 - trace 1236
- HP-UX client
 - trace 1236
- HP-UXcrating and managing groups
 - security 467

I

- I/O error
 - marks active log as TRUNCATED 293
 - queues 1372
- I/O, frequent 1266
- IBM
 - software support database, searching 1327
 - support center 1327
 - change team 1351
 - dealing with 1348
 - Development Support Group 1351
 - ordering a specific PTF 1354
 - what they need to know 1350
- IBM i
 - levels supported by the WebSphere MQ Explorer 61
- IBM i Control Language 1
- IBM internal trace 1244
- IBM message 1338
- ICH message 1338
- ICHRIN03, started-task procedure table 582
- IDC message 1338
- identification
 - API exit 694
 - application level security service, example 439
 - introduction 367
 - link level security service, example 437
 - SSPI channel exit program 487
 - user written message exit 693
 - user written security exit 692

- identifier
 - component 1344
 - resource manager 1344
- identity context 425
- IEA message 1338
- IEA911E message 1335
- IEC message 1338
- IEE message 1338
- IEF message 1338
- IFASMFDP reporting program for SMF 1120
- IGAUT attribute 575
- IGQUSER attribute 575
- iKeyman
 - UNIX 631
 - Windows 631
- IKJ message 1338
- impersonation 378
- IMQ modules 1272
- IMS
 - abend 1182, 1270
 - abend code 1271
 - abend U3042 361
 - commands
 - CHANGE SUBSYS 353, 357
 - DEQUEUE TMEMBER 365
 - DISPLAY OASN 357
 - DISPLAY OASN SUBSYS 353
 - DISPLAY SUBSYS 359
 - START REGION 360
 - START SUBSYS 353
 - START TMEMBER 364
 - STOP REGION 360
 - STOP SUBSYS 353, 360
 - STOP TMEMBER 363
 - TRACE 353
 - connection security 532
 - connection status 359
 - deleting pipes 366
 - disconnecting from dependent region 360
 - in-doubt units of recovery 349
 - loop 1298
 - message flow 1304
 - OPERCMD5 security class 596
 - processor activity 1298
 - related problems 1377
 - resynchronizing the bridge 364
 - second user ID, determining 570
 - security 596
 - user IDs, security checking 570
- IMS adapter
 - connection status 359
 - control region 354
 - controlling dependent region connections 358
 - dependent regions of IMS 358
 - displaying in-doubt units of recovery 356
 - IMSID option 354
 - initializing 354
 - residual recovery entry (RRE) 357
 - restart, what happens 349
 - starting CSQQTRMN 361
 - stopping CSQQTRMN 361
 - thread 355
 - threads, displaying 356

- IMS bridge
 - application access control 597
 - Commit mode, synchronization 365
 - connection security 597
 - controlling queues 363
 - deleting messages 365
 - RACF profiles 597
 - resynchronizing 364
 - security 596
 - security checking 599
- IMS bridge, message not arriving 1303
- IMS connection
 - RESLEVEL 565
 - second user ID 565, 570
 - security checking 565
- IMS problem
 - application terminates 1378
 - cannot connect to WebSphere MQ 1377
 - IMS not operational 1378
- IMS Tpipes, commands 241
- IMS.PROCLIB library 354, 358
- in-doubt threads 344
- in-doubt units of recovery
 - CICS 345
 - IMS 350
- incident number 1350
- incorrect access 589
- incorrect output 1183
- INCORROUT keyword 1329, 1342
- indirect mode, runmqsc command 117
- Information/Access 1327
- Information/System 1327
- initialization commands 245
- initialization procedure 1180
- initiation queues
 - defining 176
- input, standard 79
- inquiring queues, sample programs 42
- inquiring within data bags 54
- installable service 426
- installation
 - security tasks 581
- installed authorization service
 - specifying 516
- installed authorization service, Specifying the 516
- instrumentation event
 - example 877
- INTEG keyword 1329
- interpreting replies to messages 278
- INTERVAL attribute of ALTER SECURITY 577
- intra-group queuing
 - RESLEVEL 566
 - security checking 566
 - user IDs for security checking 575
- introduction 19
- investigating performance 1121
- IP addresses
 - blocking 733, 734
- IPCS subcommands 1286
- IPCS VERBEXIT CSQXDPRD
 - keywords 1287
- IPPROCS attribute 1175
- ISPF, showing keys 267

- issuing
 - MQSC commands remotely 117
 - MQSC commands using an ASCII file 76
 - MQSC commands using runmqsc command 76
- issuing commands 238, 248
- IST message 1338
- items
 - counting 57
 - deleting 57
 - filtering 53
 - querying 53
- items, types of 51
- IWM message 1338
- IXC message 1338

J

- Java 422
- Java diagnostics 1262
- Java Message Service (JMS) 422
- Java tracing 1262
- JAVA_HOME on UNIX 631
- JMS 422
- job name, specifying for GTF 1245
- journals
 - backup 200
 - introduction 196
 - managing 201
 - restoring 203
 - using 197

K

- Kerberos 487
- key 369
- key database file
 - setting up 632
 - UNIX key repository 390
- key distribution problem
 - a solution 780
 - symmetric cryptography 369
- key repository
 - adding personal certificate
 - IBM i 626
 - UNIX 644
 - z/OS 679
 - adding server certificate
 - IBM i 626
 - changing
 - queue manager on IBM i 622
 - queue manager on UNIX 637
 - defining 386
 - importing personal certificate
 - z/OS 679
 - introduction 390
 - locating
 - queue manager on IBM i 622
 - queue manager on UNIX 636
 - queue manager on z/OS 676
 - WebSphere MQ MQI client on UNIX 637
 - setting up
 - IBM i 620
 - UNIX 632

- key repository (*continued*)
 - setting up (*continued*)
 - Windows 632
 - z/OS 675
 - specifying
 - WebSphere MQ MQI client on UNIX 638
 - working with
 - z/OS 676
- key ring
 - setting up 675
- KEYLIST, ISPF command 267
- KeyRepository field 386
- keyring
 - z/OS key repository 390
- keyword
 - building a string 1331
 - component-identifier 1333
 - CSECT 1336
 - modifier
 - load module 1337
 - recovery routine 1337
 - prefix 1330
 - release level 1333
 - selecting 1331
 - symptom-to-keyword
 - cross-reference 1343
 - type-of-failure
 - ABEND 1334
 - determining 1334
 - DOC 1341
 - INCORROUT 1342
 - LOOP 1338
 - MSG 1338
 - PERFM 1340
 - WAIT 1338
- keyword format
 - free 1329
 - structured database 1330
 - z/OS 1329

L

- layout
 - type 115 SMF records 1123
 - type 116 SMF records 1138
- LDAP server 701
 - configuring and updating 698
 - setting up 697
 - working with Certificate Revocation Lists 695
- LDIF (LDAP Data Interchange Format) 697
- libraries, using SAVLIB to save WebSphere MQ 233
- Lightweight Directory Access Protocol (LDAP) server 701
- LIKE attribute, DEFINE command 89, 98, 101
- limits, queue depth 861
- line trace 1244
- link level security
 - channel exit programs
 - introduction 440
 - writing your own 437
 - comparison with application level security 435

- link level security (*continued*)
 - introduction 437
 - providing your own 437
 - SNA LU 6.2 security services 453
 - SSL 386
 - SSPI channel exit program 487
- Linux
 - security 470
- listener
 - commands 241
 - starting 115
- listener, restarting with ARM 342
- listeners
 - defining listeners for remote administration 114
- load balancing
 - CF structures 316
 - page set 302
 - sample job for a CF structure 317
 - sample job for a page set 303
- load messages on a queue 247
- load module
 - in dump title 1289
 - modifier keyword 1337
- Load Module modifier keyword 1328
- local administration
 - definition of 4
 - issuing MQSC commands using an ASCII file 76
 - runmqsc command, to issue MQSC commands 76
 - using the WebSphere MQ Explorer 59
 - writing Eclipse plug-ins 71
- local administration, definition of 184
- local queues 87, 96, 172
 - changing queue attributes, commands to use 90, 100, 172
 - clearing 90, 172
 - copying a local queue definition 89, 172
 - defining 87, 172
 - defining application queues for triggering 176
 - deleting 90, 172
 - working with local queues 87, 96, 172
- local subscription
 - copying a local subscription definition 101
 - deleting 101
- local subscriptions
 - defining 99
- local topic
 - defining 96
- locating archive log data sets to be deleted 290
- locating key repository
 - IBM i
 - queue manager 622
 - queue manager on z/OS 676
 - UNIX
 - queue manager 636
 - WebSphere MQ MQI client 637
- lock manager statistics 1132
- log
 - archiving 286
- log (*continued*)
 - determining inventory contents 293
 - error 1314
 - error recovery procedures 1359
 - file, @SYSTEM 1314
 - managing 285
 - off-load, cancelling 288
 - recovering from problems
 - active log 1359
 - archive log 1365
 - recovery 289
 - restarting archive process 288
 - log buffer pools 1266
 - log data sets 422
 - log data sets, restart on losing 322
 - log manager statistics 1128
 - log print utility 246
 - log print utility (CSQ1LOGP)
 - finding start RBA with 348
 - print log records 289
 - logger
 - events
 - controlling 842
 - logger events 869
 - logging
 - change log inventory 246
 - commands 241
 - print log map 246
 - printing the log 246
 - using SET commands 288
 - logs
 - error logs 1229
 - long-running unit of work 1176
 - loop
 - batch application 1298
 - causes 1295
 - CICS application 1298
 - distinguishing from a wait 1294
 - IMS region 1298
 - TSO application 1298
 - WebSphere MQ 1298
 - LOOP keyword 1329, 1338
 - lowercase queue names
 - operations and control panels 261
 - LU 6.2
 - channels, user IDs used 572
 - LU 6.2 and ARM 343
 - LU 6.2 connection 1301
- M**
- MAC 371
- man in the middle attack
 - introduction 371
 - SNA LU 6.2 session level authentication 454
- managing
 - BSDS 292, 294
 - buffer pools 311
 - page sets 300
 - queue-sharing groups 313
 - shared queues 315
 - WebSphere MQ log 285
- managing objects for triggering 108, 176
- manipulating objects at startup 245
- manually stopping a queue manager 74
- manuals
 - problems 1341
- mapping
 - IP address to MCAUSER 739
 - MCAUSER to MCAUSER 736
 - queue manager to MCAUSER 735
 - SSL DN to MCAUSER 737
- maximum depth reached 855
- maximum line length, MQSC
 - commands 81, 185
- MAXMSGLE 1300
- MCA
 - adopting 1310
- MCA user ID, security 571, 590
- MCAUSER 433
 - setting
 - by IP address 739
 - by MCAUSER 736
 - by queue manager 735
 - by SSL DN 737
- MCAUSER parameter
 - initial value of MCAUserIdentifier field 758
- MCAUserIdentifier 433
- MCAUserIdentifier field 758
- MCPROP parameter
 - DEFINE COMMINFO 160
- media images
 - automatic media recovery failure, scenario 1358
 - introduction 199
- media recovery 1379
- member class, security 521
- message
 - contains unexpected information 1304
 - CSQV086E 1336
 - error 1306
 - grouping 1303
 - IEA911E 1335
 - incorrect queue 1304
 - not arriving on queue 1300
 - not on queue
 - cluster queue 1303
 - IMS bridge 1303
 - trigger information 1304
 - where to find more information 1338
- Message Authentication Code (MAC)
 - Data Encryption Standard (DES)
 - SNA LU 6.2 conversation level authentication 458
 - SNA LU 6.2 session level authentication 454
 - introduction 371
 - part of CipherSuite 380
- message channel agent (MCA)
 - channel exit programs 440
 - default user ID
 - role in access control 758
 - user ID in an SNA LU 6.2 attach request 458
 - use in SSL 386
- message context
 - introduction 367
 - role in access control 425

- message descriptor
 - accounting and statistics
 - messages 984
 - activity report 929
 - trace-route message 954
 - trace-route reply message 965
- message digest 371, 381
- message exit
 - introduction 442
 - providing your own link level security 438
- Message Format 160
- message length, decreasing 90, 173
- message level security 439
- message manager
 - accounting 1145
 - statistics 1129
- message monitoring 1315
- message processing program (MPP) 358
- messages
 - commands 241
 - containing unexpected information 1155
 - context
 - granting authority to pass 743
 - granting authority to set 742
 - converting user-defined message formats 121
 - discarded 277
 - granting authority to pass
 - context 743
 - granting authority to set context 742
 - interpreting replies to WebSphere MQ
 - commands 278
 - not appearing on queues 1155
 - on the system-command input queue 275
 - undelivered, security 538
 - violation, security 588
 - waiting for replies to 276
- MGCRC 248, 568
- migrating queues from non-shared to shared 319
- model queues
 - DEFINE QMODEL command 95
 - defining 95, 175
 - working with 95, 175
- model queues, security 536
- modifier keyword
 - load module 1337
 - recovery routine 1337
- module, in dump title 1289
- monitoring
 - DISPLAY commands 1115
 - IMS connection activity 358
 - message 1315
 - performance 1114
 - queue managers 830, 889
 - resource usage 1114
 - tools 1114, 1115
 - trace-route messaging 894
- monitoring performance of WebSphere MQ for Windows queues 1114
- moving queues
 - non-shared queue 302
 - shared queue 316
- MPP (message processing program)
 - connection control 358
- MQ Services
 - changing the password 70
- MQ_USER_ID 433
- mqAddByteString 53
- mqAddByteStringFilter 53
- mqAddInquiry 53
- mqAddInteger 53
- mqAddInteger64 53
- mqAddIntegerFilter 53
- mqAddString 53
- mqAddStringFilter 53
- MQADMIN security class 521, 524
- MQAI
 - concepts and terminology 19
 - examples 19
 - introduction 19
 - sample programs
 - creating a local queue 20
 - displaying events 25
 - inquire channel objects 35
 - inquiring queues 42
 - printing information 42
- MQAI (WebSphere MQ Administration Interface) 20
- MQAI (WebSphere MQ administrative interface)
 - description of 18
- MQAI, description of 185
- mqBagToBuffer 56
- mqBufferToBag 56
- MQCD structure
 - specifying that an MQI channel uses SSL 398
- MQCFH structure
 - activity report 934
 - trace-route message 959
 - trace-route reply message 966
- mqClearBag 55
- MQCLOSE options, security 540
- MQCMDMS security class 521
- MQCNO structure
 - specifying that an MQI channel uses SSL 398
- MQCONN security class 521, 530
- MQCONN call
 - specifying that an MQI channel uses SSL 398
- mqCreateBag 50
- mqCreateBag options 50
- MQCSP 426
- mqDeleteBag 50
- mqDeleteBag options 50
- mqDeleteItem 57
- MQEPH structure
 - activity report 933
 - trace-route message 958
- mqExecute 58
- MQGET in administration programs 273
- MQGET security 535
- MQI (message-queuing interface)
 - authorization specification tables 471, 502
 - authorizations 472, 502
- MQI accounting message data 985
- MQI authorizations 472, 502
- MQI channel
 - comparing link level security and application level security 435
- MQI client security 573
- MQI statistics message data 1006
- MQJMS_TRACE_DIR 1260
- MQJMS_TRACE_LEVEL 1260
- mqm group 418, 753
- MQNLIST security class 521
- MQOPEN
 - security options 540
 - user IDs used 569
- MQOPEN authorizations 472, 502
- MQOPEN/MQPUT1 options, security 533
- MQPMO structure 1303
- MQPROC security class 521
- MQPUT authorizations 472, 502
- MQPUT in administration programs 273
- MQPUT security 535
- MQPUT1
 - security options 540
 - user IDs used 569
- mqPutBag 9
- MQQUEUE security class 521
- MQS_TRACE_OPTIONS, environment variable 1236
- mq.ini configuration file
 - path to 84
- MQSC commands
 - authorization 505
 - command security 420
 - encapsulated within Escape PCF commands 418
 - escape PCFs 185
 - maximum line length 185
 - overview 184
 - resource security profiles 560
 - runmqsc command 418
 - security 551
 - STRMQMQSC command 419, 494
 - system command input queue 421
- MQSCO structure 386
- MQSERVER
 - specifying that an MQI channel uses SSL 399
- mqSetByteString 55
- mqSetByteStringFilter 55
- mqSetInteger 55
- mqSetInteger64 55
- mqSetIntegerFilter 55
- mqSetString 55
- mqSetStringFilter 55
- MQSSLKEYR
 - environment variable 386
- mqTruncateBag 55
- MQXQH structure 435
- MQZ_AUTHENTICATE_USER 426
- MQZAO, constants and authority 472, 503
- MSCS (Microsoft Cluster Server)
 - introduction 1
- MSG keyword 1338
- MSGHIST parameter
 - DEFINE COMMINFO 164
- MsgId 1300

- MsgId field, administration
 - programs 277
- MsgId parameter 1185
- Multicast 160
- Multicast Message Format 160
- multiple transactions 1301
- MUSR_MQADMIN 70
 - changing the password 70
- mutual authentication
 - comparing link level security and application level security 435
 - definition 367
 - SSPI channel exit program 487
- MXADMIN security class 524

N

- namelist security 429
 - implementing 462
 - profile 546
- namelists
 - commands 241
 - disposition 265
 - granting administrative access 719, 726
 - granting authority to access 750
 - working with 272
- names, of event queues 832
- naming convention 1272
- national language support
 - data conversion 121
- negotiations on startup 1308
- nested groups 490
- network considerations for ARM 342
- network problems 1178
- new application 1179
- NEWLOG, utility function (CSQJU003) 295, 296
- NID (network ID) 348, 350
- non-repudiation
 - digital signature 381
- nonpersistent messages 275, 1303
- NPMSPEED attribute 1303
- NSUBHIST parameter
 - DEFINE COMMINFO 164
- NTLM 487
- null event messages 858

O

- OAM 704
- OAM (object authority manager)
 - guidelines for using 518
 - sensitive operations 518
- OAM (Object Authority Manager)
 - description of 493
 - resources protected by 493
- OAM Authenticate User 426
- OAM generic profiles 705
- object authority manager (OAM) 426, 704
- object definitions
 - recording 247
- objects
 - access to 464, 492
 - administration of 4, 184

- objects (*continued*)
 - altering 271
 - automation of administration tasks 185
 - backing up definitions 309
 - creating 170
 - default configuration, Windows systems 1
 - default object attributes, displaying 89, 97
 - defining 269
 - defining and manipulating at startup 245
 - deleting 271
 - description of 112
 - displaying 271
 - disposition 238, 265
 - generic profiles 515
 - group, managing 320
 - managing objects for triggering 108, 176
 - operations and control panels 269
 - remote administration 110
 - remote queue objects 120
 - security 430
 - using MQSC commands to administer 184
- offloading, errors during 293
- OK
 - events rules 849
- OK response 11
- opening the system-command input queue 273
- operating, basic operations 248
- operations and control panels 238
 - accessing WebSphere MQ objects 422
 - changing the subsystem ID 268
 - example of 268
 - function keys 267
 - invoking 261
 - MQSC commands 421
 - queue manager default 266
 - queue manager level 266
 - RESLEVEL 563
 - rules for using 261
 - security 568
 - security checking 563
 - SYSTEM.* queue security 539
 - using 260
 - using the command line 268
 - working with object definitions 271
- operator
 - commands, no response from 1157
 - operator command, no response from 1184
- operator commands
 - IMS adapter 353
 - issuing 248
 - operations and control panels 260
- OPERCMD5 class 596
- OPPROCS attribute 1176
- ordering a specific PTF 1354
- origin context 425
- OTMA
 - connection to WebSphere MQ 1303
- out of space on active log 1364

- outage
 - messages in doubt 1301
- output, standard 79

P

- page set full 1299
- page set problem
 - I/O error 1372
 - page set full 1373
- page sets 422
 - adding 301
 - backing up 307
 - commands 241
 - copy 247
 - creating a point of recovery 307
 - defining 245
 - display usage 301
 - expanding 304
 - extract information 247
 - format 247
 - full 301, 1373
 - load balancing 302
 - managing 300
 - problems 1372
 - recovery 309
 - reset the log after copying 247
 - restart on losing 322
- panels
 - blank fields in 262
 - operations and control 260, 268
 - RESLEVEL 563
 - rules for using 261
 - security 568
 - security checking 563
 - SYSTEM.* queue security 539
- panels, dump display 1279
- password 417
- PASSWORD parameter
 - SNA LU 6.2 conversation level authentication
 - IBM i, UNIX, Windows 458
 - z/OS 458
- password stashing
 - certificate store on IBM i 622
- passwords
 - archive log data set 297
- path validation policy 660
- PCF (programmable command format)
 - administration tasks 185
 - authorization specification tables 471, 502
 - automating administrative tasks using PCF 185
 - escape PCFs 185
 - MQAI, using to simplify use of 18
- PCF (Programmable Command Format)
 - responses 10
- PCF commands 418
 - security 556
- PCF messages
 - converting from bag 9
 - sending 9
- peer recovery
 - with queue-sharing 1308
- Peer recovery 1308

- peer-shared-channel retry on z/OS
 - about 1311
- PERFM keyword 1329, 1340
- performance
 - DISPLAY commands 1115
 - effect of WebSphere MQ trace 1116
 - events
 - controlling 842
 - example SMF records 1125
 - investigating individual tasks 1121
 - monitoring 1114
 - problems 1121
 - SMF trace 1119
 - snapshots 1115
 - symptoms of reduced 1121
 - trace 1236, 1238
 - tracing Windows, performance considerations 1234
- performance considerations 1293
- performance degradation 1185
- performance event
 - control attribute 846
 - event data 845
 - event statistics 845
 - queue 832
 - types of 837, 846
- performance monitor 1114
- Performance Reporter 1120
- performance statistics 1122
- persistent messages 1300
 - shared queues 315
- personal certificate
 - adding to key repository
 - IBM i 626
 - UNIX 644
 - z/OS 679
 - creating RACF signed 678
 - creating self-signed
 - UNIX 639
 - z/OS 677
 - deleting
 - IBM i 628
 - exporting
 - z/OS 679
 - exporting, UNIX 649
 - importing to key repository
 - z/OS 679
 - importing, UNIX 650
 - introduction 371
 - renewing
 - UNIX 643
 - requesting
 - IBM i 624, 625
 - UNIX 641
 - z/OS 678
- PFSHOW, ISPF command 267
- ping
 - problem determination 1307
- PING CHANNEL command 1178
- PKCS #11
 - cryptographic hardware cards on UNIX 390
- PKCS #11 hardware
 - managing certificates on 655
 - personal certificate
 - importing 659
 - requesting 657
- plaintext 369
- point of recovery, creating 307
- policy
 - certificate 660
 - CRL 660
 - path validation 660
- prefix keyword 1330
- primary keywords 1350
- principals 756
- print log map utility 246
- print log map utility (CSQJU004)
 - BSDS time stamps 293
 - listing BSDS contents using 293
- printing dumps 1288
- printing information, sample programs 42
- privacy
 - SSL 378
- private definitions 238
- private key
 - digital certificate 371
 - introduction 369
- PRIVATE object disposition 266
- problem
 - management record 1350
 - reporting sheet 1348
 - tracking 1348
- problem determination 1178
 - applications or systems running slowly 1160
 - channel refuses to run 1308
 - channel startup negotiation errors 1308
 - channel switching 1313
 - clients 1313
 - connection switching 1313
 - conversion failure 1310
 - data structures 1312
 - dead-letter queue 1307
 - error messages 1306
 - FFST (First Failure Support Technology) 1320
 - incorrect output, definition of 1155
 - incorrect output, distributed queuing 1155
 - intermittent problems 1160
 - no response from operator commands 1157
 - performance 1121
 - problem characteristics 1163
 - problems affecting parts of a network 1160
 - problems that occur at specific times in the day 1160
 - programming errors 1161
 - questions to ask 1152
 - queue failures, problems caused by 1157
 - queue manager, problems creating or starting 1154
 - remote queues, problems affecting 1154
 - reproducing the problem 1154
 - retrying the link 1311
 - return codes 1159
 - trace 1236, 1238
 - transmission queue overflow 1307
- problem determination (*continued*)
 - triggered channels 1309
 - undelivered-message queue 1307
 - user-exit programs 1312
 - using the PING command 1307
 - validation checks 1308
- problems, solving
 - WebSphere MQ resource adapter 1204
- process definitions
 - creating 176
 - displaying 108, 176
- process security 429
 - implementing 462
 - profile 545
- processes
 - commands 241
 - granting administrative access 718, 726
 - granting authority to access 750
- processes, disposition 265
- processing a dump
 - using IPCS in batch 1288
 - using line mode IPCS 1283
 - using the CSQWDMP statement 1283
 - using the panels 1279
- processor activity
 - WebSphere MQ 1297
- profile, RACF 522
 - for alternate user security 547
 - for command resources 560
 - for command security 551
 - for connection security
 - batch connections 531
 - channel initiator 532
 - CICS 531
 - IMS 532
 - for context security 549
 - for namelists 546
 - for process security 545
 - for queue security 533
 - for topic security 542
 - switch 523
 - used to protect WebSphere MQ resources 530
- profiles, OAM generic 705
- program abend 1270
- program check 1268
- program error 1268
 - queue manager detected 1268
 - user-detected 1268
- program, administration 272
- Programmable Command Format (PCF)
 - authority checking
 - HP Integrity NonStop Server 14
 - HP OpenVMS 14
 - IBM i 14
 - UNIX systems 14
 - Windows NT 14
 - overview 6
 - responses 10
- Programmable Command Format (PCF)
 - commands
 - accessing WebSphere MQ objects 422
 - issued by WebSphere MQ Explorer 418

- programming
 - tracing 1260
- programming errors, examples 1181
- programming errors, examples of 1161
 - further checks 1160
 - secondary checks 1160
- proof of origin
 - digital signature 381
- protected resources 493
- protocol
 - SSL
 - in WebSphere MQ 386
- PSW, in dump title 1289
- PTF 1172, 1328
 - definition 1351
 - ordering 1354
- public key
 - cryptology 369
 - digital certificate 371
 - digital signature 381
 - infrastructure 376
 - introduction 369
- publications problem 1341
- Publish/Subscribe
 - security 460
- PUTAUT attribute 802
- PUTAUT channel attribute 571
- putting data bags 51

Q

- Q5ST 1132
- QMGR object disposition 266
- QMGR objects 239
- QMGM work management object 190
- QMGMADM group 419
- QMGMJOB work management object 190
- QMGMMSG work management object 190, 191
- QMGMRUN20 work management object 190
- QMGMRUN35 work management object 190
- QMGMRUN50 work management object 190
- QSGDISP, user messages from commands with 284
- querying data items 53
- queue 1114
 - cluster 1176
 - command 7
 - depth events 855
 - examples 859
 - depth limits 861
 - display 1175
 - distribution of 1267
 - full 1299
 - remote 1176
 - SYSTEM.ADMIN.COMMAND.QUEUE 7
- queue accounting message data 996
- queue browser, sample 91
- queue depth, current 89
- queue full 1300
- queue information, displaying 1273

- queue manager
 - controlling activity recording 891
 - controlling trace-route messaging 897
 - event queue 832
 - events
 - controlling 841
 - failure, handling 1198
 - monitoring 830, 889
 - processor activity 1298
 - restoring complete 203
 - restricting access to 801
- queue manager clusters
 - security 459
- queue manager detected errors 1268
- queue manager level security 420
- queue managers
 - adding a page set 301
 - adding to a queue-sharing group 313
 - attributes, changing 85
 - attributes, displaying 85
 - authorizations 518
 - CCSID, changing 121
 - changing the CCSID 121
 - cold start 325, 338
 - command server 116, 187
 - commands 241
 - default configuration, Windows systems 1
 - directories 518
 - disabling remote access 739
 - expanding a page set 304
 - granting administrative access 717, 725
 - granting authority to inquire 749
 - granting full administrative access 729
 - granting read-only access 728
 - how to increase the size of a page set 304
 - increasing a page set 304
 - increasing the size of a page set 304
 - object authority manager, description 493
 - preparing for remote administration 112
 - queue manager aliases 120
 - remote administration 110
 - removing connectivity 730
 - removing from a queue-sharing group 314
 - restarting 322
 - securing remote connectivity 732
 - showing and hiding, using the WebSphere MQ Explorer 66
 - starting 250, 327
 - startup messages 254
 - stopping 252, 329
 - stopping a queue manager manually 74
 - z/OS queue manager 118
- queue security 429
 - alter attributes 533, 551
 - implementing 462
 - profiles 533
- queue service interval events
 - algorithm for 848
 - enabling 849

- queue service interval events (*continued*)
 - examples 850
 - high 846
 - OK 846
- queue statistics message data 1017
- queue-level accounting 1146
- queue-sharing environment
 - error messages 1173
- queue-sharing group 586, 858
- queue-sharing group level security 420
- queue-sharing group security
 - alternate user 547
 - command 551
 - command resource 560
 - connection 530
 - context 549
 - incorrect access 589
 - intra-group queuing 575
 - namelist 546
 - process 545
 - queue 533
 - violation messages 588
- queue-sharing groups
 - adding a queue manager 313
 - adding to Db2 tables 313
 - and ARM 342
 - backup 334
 - cold start 326, 338
 - command scope 240
 - commands 241
 - controlling security 525
 - disaster recovery 334
 - example security scenario 608
 - load balancing 316
 - managing 313
 - moving a queue 316
 - overriding security settings 524
 - perform Db2 tasks 247
 - re-initializing queue managers 326, 338
 - recovering units of recovery 352
 - removing a queue manager 314
 - removing from Db2 tables 314
 - sharing object definitions 238
 - user messages from commands 282
- queues
 - alias 93, 174
 - alter attributes, security 533, 551
 - application queues 176
 - browsing 91
 - changing queue attributes 90, 97, 100, 172
 - clearing local queues 90, 172
 - commands 241
 - copy 247
 - current queue depth, determining 89
 - dead-letter, defining 88, 172
 - defining local 270
 - deleting a local queue 90, 172
 - deleting a subscription 101
 - disposition 265
 - distributed, incorrect output from 1155
 - empty 247
 - granting administrative access 713, 722

queues (*continued*)
 granting authority to get messages 741
 granting authority to put messages 744, 745
 implementing security 462
 initiation queues 176
 load messages 247
 local definition of a remote queue 118
 local, working with 87, 96, 172
 migrating non-shared to shared 319
 model queues 95, 175
 moving a non-shared queue 302
 moving a shared queue 316
 moving non-shared to shared 318
 preparing transmission queues for remote administration 113
 queue manager aliases 120
 remote queue objects 120
 reply-to model 273
 reply-to queues 120
 restricting access to 802
 security 533
 system-command input 273
 system-command reply-to model 273
 working with 98
 QUIESCE MODE of ARCHIVE LOG 286
 QUIESCE, stop mode 252, 329
 quiescing
 WebSphere MQ for IBM i 234
 QWHCCV 1145

R

RACF 420
 authority, dead-letter queue 538
 authorization
 ICHRIN03 582
 STARTED class 582
 started-task procedure table 582
 to WebSphere MQ data sets 583
 profiles 522
 security classes 521
 RACF profiles, IMS bridge 597
 raising an APAR 1354
 RBA (relative byte address), range specified in active log 295
 read-only access
 to queue manager 728
 real-time monitoring
 controlling 1103
 reason code
 associated with subsystem
 abend 1270
 in dump title 1289
 reason codes
 alphabetic list 1380
 receive exit
 introduction 442
 providing your own link level security 438
 receiver channel, automatic definition of 115
 receiving data bags 51
 record layouts
 CSQDQ5ST 1132

Recording 169
 Recording WebSphere MQ application definitions 169
 recoverable objects, defining and manipulating at startup 245
 recovering shared queues 315
 recovery
 active log problems 1359
 alternative site 323, 331, 335
 automatic media recovery failure, scenario 1358
 basic operations 248
 BSDS
 errors 1368
 log inventory 290
 CICS, manually recovering units of recovery 346
 COPY 311
 creating a point of 307
 disk drive failure, scenario 1356
 example scenarios 1358
 IMS
 manually recovering units of recovery 349
 resolving in-doubt units of recovery 356
 resynchronizing the bridge 364
 IMS units of recovery 350
 introduction 195
 logs 289
 long-running UOW 1377
 media images 199
 of coupling facility structures 315
 performance 233
 point of 307
 recovering a damaged queue manager object, scenario 1357
 recovering a damaged single object, scenario 1358
 RRS, manually recovering units of recovery 351
 single BSDS 298
 starting 250, 252, 327, 329
 tokens 348
 using journals 197
 WebSphere MQ-related problems
 active log problems 1359
 archive log problems 1365
 BSDS 300, 1368
 page set problems 1372
 recovery actions 1268
 recovery routine keyword 1337
 recovery routine modifier keyword 1328
 recovery with queue-sharing
 peer 1308
 redirecting input and output, MQSC commands 81, 83
 reduced performance, symptoms of 1121
 REFRESH SECURITY command 529, 578
 refreshing
 security 578
 region error options (REO) 358
 registering with ARM 341
 Registration Authority 376
 Registration time for DDNS 1311
 reinitializing a queue manager 325, 338

relative byte address (RBA), range specified in active log 295
 release in dump title 1289
 release-level keyword 1328, 1333
 remote access
 disabling
 queue manager 739
 remote administration
 administering a remote queue manager from a local one 112
 command server 116, 187
 defining channels, listeners, and transmission queues 114
 definition of remote
 administration 4, 184
 initial problems 117
 of objects 110
 preparing channels for 113
 preparing queue managers for 112
 preparing transmission queues for 113
 security, connecting remote queue managers, the WebSphere MQ Explorer 63
 using the WebSphere MQ Explorer 59
 writing Eclipse plug-ins 71
 remote connectivity
 securing
 queue manager 732
 remote issuing of MQSC commands 117
 remote queue
 initial checks 1176
 message not on queue 1301
 remote queue managers
 blocking channel access 737
 remote queue objects 120
 remote queues
 as reply-to queue aliases 120
 command resource checking 561
 defining remote queues 118
 security 537
 security considerations 518
 suggestions for remote queuing 117
 remote queuing 110
 remote system
 reply message 1302
 RemoteUserIdentifier field 759
 removing connectivity to queue managers 730
 REO (region error options) 358
 replies, examples 279
 reply messages 276
 reply-to queue
 attributes 273
 defining 273
 opening 274
 reply-to queue aliases 120
 reply-to queues
 reply-to queue aliases 120
 reporting events 832
 repository
 failure, handling 1198
 REPRO command of access method services 299
 request message 275
 requesting activity reports 891

- RESET CHANNEL command 1309
- reset queue statistics 846
- reset service timer 848
- reset the log after copying a page set 247
- resident trace table, display 1281
- RESLEVEL
 - auditing 561, 587
 - checking CICS user IDs 564
 - distributed queuing 566
 - implementing 462
 - IMS connection 565
 - intra-group queuing 566
 - introduction 561
 - usage notes 561
 - using 561
- RESLEVEL profile
 - introduction 430
- RESOLVE CHANNEL command 1309
- RESOLVE INDOUBT command, free locked resources 347
- resolving
 - in-doubt units of recovery 356
 - units of recovery 346, 351
- resolving a problem 1354
- Resource Access Control Facility (RACF)
 - authority checks on z/OS 420
- resource manager formatting
 - keywords 1284
- resource manager identifier, list of 1344
- Resource Measurement Facility (RMF) 1120
- Resource Recovery Services (RRS), units of recovery 351
- resource security
 - alias queues 561
 - API 540
 - commands 560
 - remote queues 561
- resource-level security checks
 - CICS adapter 592
 - switch profiles 528
- resources
 - security 583
- response
 - data 11
 - extended 12
 - OK 11
 - standard 11
- restart
 - after abnormal termination 322
 - after losing data sets 325, 338
 - after losing logs 322
 - after losing page sets 322
 - CICS adapter 345
 - cold start 325, 338
 - IMS adapter 349
 - introduction 195
 - long-running UOW 1377
 - media images 199
 - performance 233
 - user messages 345
 - using journals 197
 - with ARM 339
 - z/OS Automatic Restart Manager 339
- restarting queue managers 322
- restoring
 - complete queue manager 203
 - journal receivers 204
- restricting access to MQM objects 492
- restricting access to queue managers 801
- restricting access to queues 802
- restricting access using alias queues 535, 538
- restrictions
 - access to MQM objects 464
- RESUME QMGR 320, 362
- RETAIN 1350, 1351
 - searching 1327
 - symptoms 1330
- retention period, archive logs (ARCRETN) 289
- retry considerations 1311
- retrying the link, problem determination 1311
- return codes 1172
 - changed application 1179
 - problem determination 1159
- RMF 1174
- RMF (Resource Measurement Facility) 1120
- route tracing 1315
- RRE (residual recovery entry) 357
- RRS
 - abend 1270
 - active 1297
- RRS (Resource Recovery Services), units of recovery 351
- RRS, active 1176
- RSA 380
- rules
 - High events 848
 - OK events 849
 - service timer 848
- rules for using commands 77
- rules for using the operations and control panels 261
- RUNMQBRK task 189
- RUNMQCHI task 189
- RUNMQCHL task 189
- RUNMQDLQ task 189
- RUNMQLSR task 189
- runmqsc (run WebSphere MQ commands) command
 - ending 79
 - feedback 79
 - indirect mode 117
 - problems, resolving 84
 - redirecting input and output 81, 83
 - using 81, 83
 - using interactively 79
 - verifying 81
- runmqsc command
 - introduction 418
 - sending MQSC commands to a system command input queue 421
- RUNMQTRM task 189
- RVERIFY SECURITY command 577
- RVKMQMAUT command 428, 516
- sample
 - queue-sharing group security scenario 608
 - SMF accounting record 1146
 - SMF type 116 record
 - subtype 1 1147
 - subtype 2 1147
- sample ARM policy 340
- sample problem reporting sheet 1348
- sample programs
 - creating a local queue 20
 - displaying events 25
 - inquire channel objects 35
 - inquiring queues 42
 - printing information 42
- save area trace report, display 1282
- SAVLIB, using to save WebSphere MQ libraries 233
- scenario
 - security
 - two queue managers 600
- SDB (structured database) keyword
 - format 1330
- SDB format keywords 1343
- search argument
 - process 1327
 - varying 1328
- second user ID, IMS connection 565, 570
- secret key 369
- secret keys 401, 778
- secure sockets layer (SSL)
 - channel parameters 449
 - MQSC commands 449, 520
 - protecting channels 449
 - queue manager parameters 449
- Secure Sockets Layer (SSL), configuration 584
- securing remote connectivity to queue managers 732
- security 801
 - access control 432, 704, 755
 - access settings 705, 709, 711
 - activating 523
 - administration authority 751
 - alternate-user authority 757
 - API quick-reference table 540
 - auditing considerations 586
 - authentication 613
 - authority, alternate-user 757
 - authority, context 758
 - authorizations to use the WebSphere MQ Explorer 62
 - automating starting of CKTI 593
 - batch connections 530
 - blank user IDs 576
 - channel 584, 590
 - channel exits 449, 452
 - channels 447
 - checks 754
 - checks, preventing 712
 - CICS adapter 592
 - transactions 593
 - user IDs 593
 - CICS bridge 594
 - CKTI 593
 - CKTI user IDs 593
 - clustering 592

S

SAF 420

- security (*continued*)
 - command summary table 551
 - commands 241
 - connecting to remote queue managers,
 - the WebSphere MQ Explorer 63
 - connection 530, 568
 - considerations 492
 - context authority 518, 758
 - customizing 587
 - data sets 582
 - default user ID 568
 - displaying status 580
 - dmpmqaut command 705, 709
 - domain controller 489
 - dspmqaout command 711
 - example queue-sharing group
 - scenario 608
 - exit 416
 - groups 756
 - identifiers 757
 - implementation checklist 462
 - IMS 596
 - incorrect access 589
 - installation tasks 581
 - INTERVAL attribute 577
 - intra-group queuing 575
 - Linux 470
 - management 577
 - MQCLOSE/MQOPEN/MQPUT1
 - options 540
 - MQI authorizations 472, 502
 - mqm group 753
 - nested groups 490
 - OAM 704
 - object authority manager
 - (OAM) 493, 704
 - objects
 - UNIX systems 430
 - Windows systems 430
 - on a WebSphere MQ MQI client 613
 - operations and control panels 568
 - OTMA 596
 - password 417
 - principals 756
 - problem determination 462
 - profile, RESLEVEL 561
 - queue-sharing group 524
 - queues
 - alias 535, 561
 - dead-letter 538
 - dynamic 537
 - model 536
 - profiles 533, 551
 - remote 537, 561
 - transmission 540, 550, 561
 - refreshing 578
 - remote queues 518
 - resources 583
 - resources protected by the OAM 493
 - scenarios
 - two queue managers 600
 - security for the WebSphere MQ Explorer 62, 68
 - sensitive operations, OAM 518
 - setmqaut command 704
 - setting checks off 524
 - Solaris 470

- security (*continued*)
 - tasks 712
 - template files 490
 - terminal user IDs 593
 - TIMEOUT attribute 577
 - topics
 - profiles 542
 - transmission queues 449
 - undelivered messages 538
 - universal access (UACC) levels 576
 - user ID 417, 755
 - user ID timeouts 577
 - user IDs 568, 590
 - user IDs for MQI clients 573
 - using RACF classes 521
 - utilities 568
 - WebSphere MQ authorities 494
 - WebSphere MQ objects
 - UNIX systems 753
 - Windows 753
 - WebSphere MQ Services 490
 - Windows 487
 - Windows 2003 465
 - Windows systems 757
 - Windows XP 465
- security exit
 - introduction 441
 - providing your own link level
 - security 438
 - SSPI channel exit program 487
- security exits
 - WebSphere MQ Explorer 64
- security mechanisms 367
- security messages 441
- security services
 - access control
 - API exit 760
 - authority to administer WebSphere MQ 418
 - authority to work with WebSphere MQ objects 422
 - introduction 368
 - user written message exit 760
 - user written security exit 758
- application level
 - introduction 439
 - providing your own 440
- authentication
 - API exit 694
 - introduction 367
 - SNA LU 6.2 conversation level
 - authentication 456
 - SNA LU 6.2 session level
 - authentication 454
 - SSPI channel exit program 487
 - user written message exit 693
 - user written security exit 692
- confidentiality
 - introduction 368
 - SNA LU 6.2 session level
 - cryptology 453
 - user written security exit 780
- identification
 - API exit 694
 - introduction 367
 - SSPI channel exit program 487
 - user written message exit 693

- security services (*continued*)
 - identification (*continued*)
 - user written security exit 692
 - introduction 367
 - link level
 - introduction 437
 - providing your own 437
 - SNA LU 6.2 453
 - Security Support Provider Interface (SSPI)
 - channel exit program 487
 - security switches
 - example 529
 - valid combinations 527
 - selecting a dump 1280, 1283
 - self-defining section
 - SMF type 115 records 1124
 - SMF type 116 records 1139, 1141
 - self-signed certificate
 - creating
 - UNIX 639
 - z/OS 677
 - extracting, UNIX 647
 - send exit
 - introduction 442
 - providing your own link level
 - security 438
 - sending documentation 1353
 - sending PCF messages 9
 - sensitive operations, OAM 518
 - server certificate
 - adding to key repository
 - IBM i 626
 - requesting
 - IBM i 624, 625
 - server-connection channel, automatic
 - definition of 115
 - service definitions
 - common tokens 105
 - service timer
 - resetting 848
 - rules for 848
 - services 102
 - granting administrative access 720, 727
 - setmqaut command
 - examples 427
 - introduction 418
 - setmqscp command
 - specifying that an MQI channel uses SSL 398
 - severity level 1350
 - shared channels after Db2 failure 1375
 - shared definition 239
 - SHARED object disposition 266
 - SHARED objects 239
 - shared queue
 - message not on queue 1300
 - shared queues 1177
 - coordinating queue manager 858
 - example security scenario 608
 - initial checks 1177
 - load balancing 316
 - managing 315
 - moving a queue 316
 - moving to non-shared 318
 - persistent messages 315
 - profiles for security 533

- shared queues (*continued*)
 - queue depth events 858
 - recovering 315
 - recovering units of recovery 352
 - security 533
 - sharing object definitions 238
 - user messages from commands 282
- short of space on active log 1364
- SIDs (security identifiers) 757
- signer certificate
 - introduction 371
- SMF (System Management Facility)
 - accounting record sample 1146
 - buffers 1120
 - common header 1144
 - introduction 1119
 - processing type 115 records 1127
 - processing type 116 records 1143
 - recording trace data for 1119
 - reporting data in (IFASMFDP) 1120
 - self-defining section
 - type 115 1124
 - type 116 records 1139, 1141
 - statistics records example 1125
 - type 115
 - header 1123
 - record layout 1123
 - record subtypes 1122
 - self-defining section 1124
 - type 116
 - header 1138, 1141
 - record layout 1138
- SMF global accounting interval 1120
- SNA LU 6.2
 - conversation level authentication
 - introduction 456
 - PASSWORD parameter, IBM i, UNIX, Windows 458
 - PASSWORD parameter, z/OS 458
 - security type, IBM i, UNIX, Windows 457
 - security type, z/OS 458
 - USERID parameter, IBM i, UNIX, Windows 458
 - USERID parameter, z/OS 458
 - end user verification 456
 - LU-LU verification 454
 - security services 453
 - session level authentication 454
 - session level cryptography 453
- snap dump 1277, 1291
- snapshots, performance 1115
- software support database,
 - searching 1327
- Solaris
 - MQAI support for 18
 - security 470
 - trace 1236
- Solaris client
 - trace 1236
- solving problems
 - WebSphere MQ resource
 - adapter 1204
- specifying
 - CipherSpec
 - Java client and JMS 778, 801
 - key repository
 - WebSphere MQ MQI client on UNIX 638
 - specifying coded character sets 121
 - specifying the installed authorization
 - service 516
 - SPT (started-task procedure table) 582
 - SRVCONN 452
 - SSL 390
 - certificate store
 - IBM i 621
 - configuration options 386
 - handshake 377, 386
 - platforms 386
 - protocol 386
 - setting up
 - introduction 761, 768, 782, 791
 - UNIX systems 631
 - WebSphere MQ MQI client 397
 - Windows systems 631
 - z/OS 674
 - z/OS, user requirements 674
 - SSL (Secure Sockets Layer),
 - configuration 584
 - SSL Distinguished Name
 - blocking channel access 738
 - SSLCIPH parameter 398
 - specifying CipherSpecs 388, 774, 797
 - SSLCipherSpec field 398
 - SSLKEYR parameter
 - z/OS 676
 - SSLKeyRepository field 386
 - SSLTASKS parameter
 - setting on z/OS 675
 - SSM (subsystem member)
 - contains control information 354
 - error options 358
 - specified on EXEC parameter 358
 - SSPI 487
 - stand-alone dump 1277
 - START CMDSERV 1184
 - START CMDSERV command 274
 - START QMGR command
 - from z/OS console 250, 327
 - START REGION, command of IMS 360
 - START SUBSYS, command of IMS 353
 - START TMEMBER, command of IMS 364
 - START TRACE command 1116
 - STARTED RACF class
 - authorization to data sets 583
 - authorizing procedures 582
 - started task procedure
 - authorization 582
 - starting
 - a channel 115
 - a command server 116
 - a listener 115
 - after an abend 252, 329
 - command server 274
 - IMS-WebSphere MQ connection 354
 - queue manager 72, 250, 252, 327, 329
 - WebSphere MQ 250, 252, 327, 329
 - with ARM 339
- starting (*continued*)
 - z/OS Automatic Restart
 - Manager 339
 - starting the GTF 1245
 - starting the trace 1246
 - starting WebSphere MQ trace 1116
 - startup messages (queue manager) 254
 - statistics
 - buffer manager 1129
 - buffer pool 1130
 - CICS adapter 1116
 - coupling facility manager 1136
 - Coupling Facility manager 1136
 - data manager 1129
 - DB2 manager 1132
 - example SMF records 1125
 - eye-catcher strings 1124
 - lock manager 1132
 - log manager 1128
 - message
 - format 981
 - message manager 1129
 - security 586
 - SMDS 1137
 - storage manager 1127
 - statistics monitoring
 - channel messages 1024
 - MQI messages 1006
 - queue messages 1017
 - statistics, events 845
 - stdin, on runmqsc 81
 - stdout, on runmqsc 81
 - STOP CMDSERV command 275
 - STOP QMGR 1297
 - STOP QMGR command
 - MODE(FORCE) 253, 330
 - MODE(QUIESCE) 253, 330
 - MODE(RESTART) 252, 329
 - STOP REGION, command of IMS 360
 - STOP SUBSYS, command of IMS 353, 360
 - STOP TMEMBER, command of IMS 363
 - STOP TRACE command 1116
 - stopping
 - a queue manager manually 74
 - command server 116
 - stopping the GTF 1245
 - storage classes
 - commands 241
 - storage classes, disposition 265
 - storage management subsystem (SMS) 289
 - storage manager statistics 1127
 - storage medium full
 - display usage 301
 - recovery scenario 1374
 - stream cipher algorithm 369
 - strength of encryption 369
 - STRMQMMQSC command 419, 494
 - strmqtrc trace command 1234, 1236
 - structure of accounting and statistics
 - messages 982
 - structured database (SDB) keyword
 - format 1330
 - structures
 - MQCFH
 - activity report 934

- structures (*continued*)
 - MQCFH (*continued*)
 - trace-route message 959
 - trace-route reply message 966
 - MQEPH
 - activity report 933
 - trace-route message 958
 - MQMD
 - activity report 929
 - trace-route message 954
 - trace-route reply message 965
- subscriptions 98
 - attributes of subscriptions, displaying 99
 - working with subscriptions 98
- SUBSYS= parameter 1281, 1284, 1286
- subsystem ID, changing 268
- subsystem name
 - finding 1282
 - in dump formatting 1281
 - in dump title 1289
- subsystem security 420, 525
- subsystem termination 1269, 1270
- Suite B 382, 399, 403
- SUMDUMP= parameter 1281, 1286
- summary dump in dump
 - formatting 1281
- summary portion of a dump 1277
- suppressing dumps 1292
- SUSPEND QMGR 320, 362
- SVC dump 1272, 1277
 - printing 1288
 - processing
 - using IPCS in batch 1288
 - using line mode IPCS 1283
 - using the CSQWDMP statement 1283
 - using the dump display panels 1279
 - suppression 1292
 - title 1288
 - title, variation with PSW and ASID 1289
- switch profiles
 - authority checks associated with MQI calls 429
 - introduction 420
- switches, security 523
- switching channels 1313
- symmetric cryptography algorithm 369
- symptom keywords 1331
- symptom string
 - display 1280
 - introduction 1273
- symptom-to-keyword
 - cross-reference 1343
- syncpoint 1184, 1300
- SYS1.DUMPxx data set 1280
- SYS1.LOGREC
 - analyzing 1292
 - data 1280
 - finding the applicable information 1292
 - introduction 1272
- system
 - changes 1180
 - displaying status 1173
 - system (*continued*)
 - running slowly 1185
 - stopped 1174
 - system abend completion code 1269
 - system administration
 - using application programs 272
 - WebSphere MQ commands 248, 250, 327
 - System Authorization Facility (SAF) 420
 - system bag 48
 - system command input queue 421
 - system control commands for starting WebSphere MQ 250, 327
 - system diagnostic work area, display 1280
 - system loading
 - time of day 1178
 - system monitoring, DISPLAY commands 1115
 - system security 582
 - System SSLtrace
 - trace 1252
 - system-command input queue
 - defining 273
 - introduction 249
 - opening 273
 - putting messages on 275
 - system-command reply-to model queue 273
 - SYSTEM.* queues
 - security 539
 - SYSTEM.ADMIN.COMMAND.QUEUE 7
 - SYSTEM.CLUSTER.COMMAND.QUEUE 1198
 - SYSUDUMP
 - introduction 1277
 - sample 1290
- T**
 - tampering 371
 - target queue manager 266
 - tasks, WebSphere MQ 189
 - TCP/IP and ARM 342
 - TCP/IP channels, user IDs used 571
 - TCP/IP connection 1301
 - template files, security 490
 - terminal user IDs, CICS adapter 593
 - terminating
 - queue manager 252, 329
 - WebSphere MQ-IMS connection 360
 - termination, abnormal 1269
 - testing WebSphere MQ classes for Java programs 1260
 - thread cross reference data 1145
 - thread-level accounting 1146
 - threads
 - active 344
 - attachment in IMS 355
 - commands 241
 - displaying 344
 - displaying, IMS adapter 356
 - IMS termination 360
 - in-doubt 344
 - stopping WebSphere MQ 252, 329
 - thresholds for queue depth 855
 - time since reset 845
 - time stamps
 - from BSDS 293
 - unequal in BSDS 1370
 - timed out responses from MQSC commands 117
 - TIMEOUT security attribute 577
 - timer
 - service 848
 - Tivoli Decision Support 1120
 - TLS
 - platforms 386
 - topic
 - security 542
 - topic security
 - implementing 462
 - topics
 - deleting an administrative topic 98
 - granting administrative access 715, 723
 - granting authority to publish messages 747
 - granting authority to subscribe 748
 - tpipe, deleting 366
 - trace
 - AIX 1236
 - channel initiator 1250
 - CICS adapter 1252
 - commands 241
 - controlling WebSphere MQ 1116
 - effect on performance 1116
 - entries for CICS adapter 1252
 - format identifier 1247
 - formatting 1247
 - HP-UX 1236
 - identifying WebSphere MQ control blocks 1247
 - interpreting 1248
 - performance considerations 1236, 1238, 1267
 - Solaris 1236
 - specifying destinations 1116
 - specifying the CLASS 1246
 - starting 1246
 - System SSL 1252
 - user parameters 1244
 - using
 - on HP Integrity NonStop Server 1239
 - on IBM i 1240
 - Windows 1234
 - Windows, performance considerations 1234
 - z/OS 1252
 - trace command
 - dspmqttrc 1236
 - endmqttrc 1234, 1236
 - strmqtrc 1234, 1236
 - trace table, display 1281
 - trace types 1245
 - trace-route message
 - format 953, 954
 - message descriptor 954
 - structure 954
 - trace-route message data 961
 - trace-route messages
 - generating 899

- trace-route messages (*continued*)
 - queue manager control 897
 - TraceRoute PCF group 900
 - trace-route messaging 894
 - trace-route reply message
 - format 964, 965
 - message descriptor 965
 - structure 965
 - trace-route reply message data 966
 - TRACE, command of IMS 353
 - TraceRoute PCF group 900
 - tracing
 - Java 1262
 - programs 1260
 - WebSphere MQ classes for JMS 1256
 - MQJMS_TRACE_LEVEL 1260
 - transmission queue 1302
 - overflow 1307
 - transmission queue header structure (MQXQH)
 - comparing link level security and application level security 435
 - message exit 442
 - transmission queues
 - creating 120
 - default transmission queues 120
 - defining transmission queues remote administration 114
 - preparing transmission queues for remote administration 113
 - security 449
 - transmission segment
 - introduction 442
 - user written send and receive exits 438
 - trigger
 - definition 1299
 - monitor 1301
 - options 1301
 - process 1301
 - trigger messages, from event queues 843
 - trigger monitor
 - job submission attributes 176
 - triggered channels, problem determination 1309
 - triggered event queues 843
 - triggering
 - defining an application queue for triggering 176
 - managing objects for triggering 108, 176
 - truncating a bag 55
 - TSO
 - abend 1182, 1270
 - application loop 1298
 - wait 1296
 - tuning WebSphere MQ 1114
 - type-of-failure keyword
 - ABEND 1334
 - determining 1334
 - DOC 1341
 - INCORROUT 1342
 - LOOP 1338
 - MSG 1338
 - PERFM 1340
 - symptom-to-keyword cross-reference 1343
 - type-of-failure keyword (*continued*)
 - WAIT 1338
 - types of data bag 48
 - types of data items 51
- ## U
- U3042 abend (IMS) 361
 - unavailable event queues 843
 - undelivered messages, security 538
 - unit of recovery
 - CICS, recovering manually 346
 - displaying in-doubt 347
 - IMS
 - in-doubt resolution 356
 - recovering manually 349
 - in-doubt
 - displaying in IMS 356
 - recovering in IMS 356
 - recovering on another queue manager 352
 - RRS, recovering manually 351
 - unit of work
 - CICS adapter 345
 - long running 1377
 - RESOLVE INDOUBT command 356
 - unit of work, and events 837
 - unit of work, long-running 1176
 - units of work
 - active 344
 - unresolved 344
 - universal access (UACC) levels 576
 - UNIX operating system
 - levels supported by the WebSphere MQ Explorer 61
 - UNIX systems
 - using distributed queuing object security 431
 - unresolved units of work 344
 - UOW
 - CICS adapter 345
 - long running 1377
 - RESOLVE INDOUBT command 356
 - updating coded character sets 121
 - use of the MQAI 19
 - user bag 48
 - user certificate
 - introduction 371
 - User context 425
 - user data related errors 1268
 - user ID 417, 432, 755
 - implementing timeouts 462
 - intra-group queuing 575
 - MQI clients 573
 - user ID security
 - batch connections
 - checks 568
 - introduction 531
 - RESLEVEL 563
 - blank 576
 - channel initiator, connection 532
 - checking 569
 - CICS
 - connection 531, 568
 - CICS adapter
 - transactions 592
 - user ID security (*continued*)
 - CICS connection
 - address space 569
 - connection 564
 - RESLEVEL 564
 - task 569
 - transactions 569
 - CKTI 593
 - CSQUTIL 563
 - distributed queuing
 - connection 566
 - RESLEVEL 566
 - IMS
 - connection 532, 568
 - IMS connection
 - address space 570
 - connection 565
 - RESLEVEL 565
 - second user ID 565, 570
 - intra-group queuing, RESLEVEL 566, 567
 - number checked 561
 - operations and control panels 563
 - RESLEVEL profile 561
 - reverification 577
 - timeouts 577
 - user ID, client asserted
 - blocking channel access 738
 - user identifier, in dump title 1289
 - user IDs
 - blocking 735
 - user messages
 - at startup 254
 - COUNT field 278
 - from WebSphere MQ commands, replies 279
 - user parameter trace 1244
 - user-defined message formats 121
 - user-exit
 - programs 1312
 - user-exit programs
 - problem determination 1312
 - USERID parameter
 - SNA LU 6.2 conversation level authentication
 - IBM i, UNIX, Windows 458
 - z/OS 458
 - userid, in dump title 1289
 - UserIdentifier field
 - authentication in a user written message exit 693
 - authentication in an API exit 691, 694
 - message containing an MQSC command 421
 - message context 425
 - use by a server application 760
 - using activity reports 889
 - using CL commands 169
 - using commands, rules for 77
 - using events 830
 - using SAVLIB to save WebSphere MQ libraries 233
 - using trace-route messaging 894
 - utilities, time stamp 293
 - utility programs
 - active log 247

- utility programs (*continued*)
 - change log inventory 246
 - CSQ1LOGP 246
 - CSQ5PQSG 247
 - CSQJU003 246
 - CSQJU004 246
 - CSQJUFMT 247
 - CSQUCVX 246
 - CSQUDLQH 247
 - CSQUTIL 247
 - data conversion 246
 - dead-letter queue handler 247
 - log print 246
 - print log map 246
 - queue-sharing group 247
- Utoken 568

V

- validation
 - checks 1308
- validation policy
 - basic 667
 - standard 670
- variable recording area 1272
- variable recording area, display 1280
- VERBEXITs 1286
- verifying MQSC commands 81
- violation messages, security 588

W

- wait
 - batch 1296
 - causes 1295
 - CICS transaction 1296
 - Db2 1297
 - distinguishing from a loop 1294
 - TSO 1296
 - WebSphere MQ 1297
- WAIT keyword 1329, 1338
- waiting for replies to messages 276
- WaitInterval parameter 1184
- WebSphere MQ
 - commands 76
 - Commands (MQSC) 1
 - creating objects 170
 - issuing MQSC commands using an ASCII file 76
 - loop 1298
 - runmqsc command, to issue MQSC commands 76
 - running slowly 1185
 - wait 1297
- WebSphere MQ Administration Interface
 - concepts and terminology 19
 - creating a local queue 20
 - displaying events 25
 - examples 19
 - inquiring queues 42
 - introduction 19
 - printing information 42
 - sample programs 19
 - use 19
- WebSphere MQ Administration Interface (MQAI) 20
- WebSphere MQ Advanced Message Security 439
- WebSphere MQ alert monitor, using 72
- WebSphere MQ channel protocol flows
 - comparing link level security and application level security 435
 - send and receive exits 442
- WebSphere MQ classes for Java 422
- WebSphere MQ classes for Java Message Service (JMS) 422
- WebSphere MQ command files
 - input 81
 - output reports 81
 - running 81
- WebSphere MQ commands
 - ARCHIVE LOG 286
 - authorization 475
 - BACKUP CFSTRUCT 315
 - command files, input 81
 - command files, output reports 81
 - command files, running 81
 - DEFINE QLOCAL 316
 - DELETE QLOCAL 316
 - DISPLAY CONN 253, 330, 344, 347, 356, 358
 - DISPLAY SECURITY 529
 - ending interactive input 79
 - issuing interactively 79
 - issuing MQSC commands remotely 117
 - maximum line length 81
 - MOVE QLOCAL 316
 - overview 76
 - problems using MQSC commands remotely 117
 - problems, list 84
 - problems, resolving 84
 - RECOVER BSDS 298
 - RECOVER CFSTRUCT 315
 - redirecting input and output 81, 83
 - REFRESH SECURITY 529
 - remote queue manager 274
 - RESOLVE INDOUBT 347, 350, 356
 - RESUME QMGR 320, 333, 362
 - runmqsc control command, modes 2, 76
 - security profiles 551
 - SET ARCHIVE 288
 - SET LOG 288
 - SET SYSTEM 288
 - START TRACE 1116
 - STOP TRACE 1116
 - SUSPEND QMGR 320, 333, 362
 - syntax errors 79
 - timed out command responses 117
 - using 81, 83
 - verifying 81
- WebSphere MQ Db2 tables
 - CSQ45STB job 1353
- WebSphere MQ display route application 908
- WebSphere MQ ESE 439
- WebSphere MQ Explorer
 - alert monitor application 72
 - AMQ7604 70
 - authority to use 418
 - authorizations to use 62

- WebSphere MQ Explorer (*continued*)
 - cluster membership 67
 - connecting to remote queue managers, security 63
 - description of 1, 186
 - introduction 1
 - MQ Services
 - changing the password 70
 - MUSR_MQADMIN
 - changing the password 70
 - performance considerations 61
 - Prepare WebSphere MQ Wizard 68
 - required resource definitions 62
 - security exits, using 64
 - security implications 62, 68
 - showing and hiding queue managers and clusters 66
 - SSL security, the WebSphere MQ Explorer 64
 - SSL security, using 64
 - user rights for WebSphere MQ Windows service 69
 - using 71
- WebSphere MQ for IBM i
 - backups of data 200
 - journal management 201
 - journal usage 197
 - journals 196
 - media images 199
 - quiescing 234
 - restoring a complete queue manager 203
 - restoring journal receivers 204
- WebSphere MQ internet pass-thru security 461
- WebSphere MQ MQI client
 - SSL 397
- WebSphere MQ objects 422
- WebSphere MQ objects, creating 170
- WebSphere MQ resource adapter
 - problem determination
 - creating connections for outbound communication 1205
 - deploying message driven beans (MDBs) 1205
 - deploying the resource adapter 1204
 - introduction 1204
- WebSphere MQ Script commands 418
- WebSphere MQ Services
 - security 490
- WebSphere MQ Services snap-in
 - MQ Services
 - changing the password 70
 - MUSR_MQADMIN
 - changing the password 70
 - security implications 68
- WebSphere MQ Taskbar application using 71
- WebSphere MQ tasks 189
- WebSphere MQ utility program (CSQUTIL)
 - accessing WebSphere MQ objects 422
 - creating certificate revocation information in a client channel definition file 703
 - issuing commands from 249

- WebSphere MQ utility program (CSQUTIL) *(continued)*
 - MQSC commands 421
- WebSphere MQ utility program (CSQUTIL), security 568
- WebSphere MQ-IMS bridge, message not arriving 1303
- Windows 1114
 - performance monitor 1114
- Windows 2003
 - security 465
- Windows client
 - trace 1234
- Windows NT LAN Manager (NTLM) 487
- Windows operating system
 - default configuration 1
 - Event Viewer application, problem determination 1230
 - FFST, examining 1315
 - levels supported by the WebSphere MQ Explorer 61
 - MQAI support for 18
 - security 757
 - tracing, considerations 1234
 - using the WebSphere MQ Explorer 59
 - writing Eclipse plug-ins 71
- Windows systems
 - using distributed queuing object security 432
- Windows XP
 - security 465
- Windows, security 487
- work management
 - objects 190
 - tasks 189
 - using 191
- work, units of 346
- working with 102
- working without authority profiles 517
- writing programs to administer WebSphere MQ 272
- WRKMQMAUT command 428
- WRKMQMAUTD command 428
- WTOR, WebSphere MQ-related 252, 329

- z/OSTrace
 - trace 1252

X

- X.509 standard
 - digital certificates comply with 372
 - DN identifies entity 373
 - public key infrastructure (PKI) 376
- XES
 - abend 1270

Z

- z/OS
 - abend 1271
 - levels supported by the WebSphere MQ Explorer 61
 - using the dump command 1277
- z/OS queue manager 118
- z/OS under stress 1293

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

Programming interface information, if provided, is intended to help you create application software for use with this program.

This book contains information on intended programming interfaces that allow the customer to write programs to obtain the services of WebSphere MQ.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Important: Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks

IBM, the IBM logo, [ibm.com](http://www.ibm.com)[®], are trademarks of IBM Corporation, registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the Web at “Copyright and trademark information”www.ibm.com/legal/copytrade.shtml. Other product and service names might be trademarks of IBM or other companies.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

This product includes software developed by the Eclipse Project (<http://www.eclipse.org/>).

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Sending your comments to IBM

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Use one of the following methods to send us your comments:

- Send an email to ibmkc@us.ibm.com
- Use the form on the web here: www.ibm.com/software/data/rcf/

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate, without incurring any obligation to you.

Include the following information:

- Your name and address
- Your email address
- Your telephone or fax number
- The publication title and order number
- The topic and page number related to your comment
- The text of your comment

IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you submit.

Thank you for your participation.

