Troubleshooting and Support for IBM MQ



Note
Before using this information and the product it supports, read the information in <u>"Notices" on page 269</u> .
edition applies to version 9 release 0 of IBM® MQ and to all subsequent releases and modifications until other ated in new editions.
you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any vers appropriate without incurring any obligation to you.

IBM Corp.

Contents

Troubleshooting and support	7
Troubleshooting overview	
Making initial checks	
Making initial checks on UNIX, Linux, and Windows	
Making initial checks on IBM i	
Making initial checks on z/OS	
Collecting troubleshooting information	
Using error logs	
Error logs on UNIX, Linux, and Windows	
Error logs on IBM i	
Error logs in IBM MQ classes for JMS	
Suppressing channel error messages from error logs on Multiplatforms	
First Failure Support Technology (FFST)	
FFST: IBM MQ classes for JMS	
FFST: IBM MQ for Windows	
FFST: IBM MQ for UNIX and Linux systems	
FFST: IBM MQ for IBM i	
Using trace	
Using trace on Windows	
Using trace on UNIX and Linux systems	
Using trace with IBM MQ server on IBM i	
Using trace with IBM MQ client on IBM i	
Using trace for problem determination on z/OS	
Tracing the Advanced Message Queuing Protocol (AMQP) Service	
Tracing TLS: runmqakm, strmqikm, and runmqckm functions	
Tracing IBM MQ classes for JMS applications	
Tracing IBM MQ classes for Java applications	91
Tracing the IBM MQ resource adapter	95
Tracing additional IBM MQ Java components	97
Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ cla	
for JMS	100
Tracing IBM MQ .NET applications	
Tracing the IBM MQ Console and REST API	
Tracing the IBM MQ Bridge to Salesforce	
Tracing the IBM MQ Bridge to blockchain	
Enabling dynamic tracing of LDAP client library code	
Problem determination on z/OS	106
IBM MQ for z/OS performance constraints	
IBM MQ for z/OS recovery actions	
IBM MQ for z/OS abends	
Diagnostic information produced on IBM MQ for z/OS	
Other sources of problem determination information for IBM MQ for z/OS	
Diagnostic aids for CICS	
Diagnostic aids for IMS	
Diagnostic aids for Db2	
Troubleshooting the connection to Product Insights from IBM MQ for z/OS	
IBM MQ for z/OS dumps	117
Dealing with performance problems on z/OS	
Dealing with incorrect output on z/OS	
Dealing with issues when capturing SMF data for the channel initiator (CHINIT)	
Problem determination in DQM	
Error message from channel control	151

	D' .	4 - 4
	Ping	
	Dead-letter queue considerations	
	Validation checks	
	In-doubt relationship	152
	Channel startup negotiation errors	152
	Shared channel recovery	152
	When a channel refuses to run	153
	Retrying the link	155
	Data structures	
	User exit problems	
	Disaster recovery	
	Channel switching	
	Connection switching.	
	Client problems	
	Error logs	
	Message monitoring	
Cha	annel authentication records troubleshooting	
	mmands troubleshooting	
	tributed publish/subscribe troubleshooting	
	Routing for publish/subscribe clusters: Notes on behavior	
	Checking proxy subscription locations	
	Resynchronization of proxy subscriptions	
	Loop detection in a distributed publish/subscribe network	
Jav	va and JMS troubleshooting	
	Troubleshooting IBM MQ classes for JMS problems	
	JMS provider version troubleshooting	
	PCF processing in JMS	
	JMS connection pool error handling	
	Troubleshooting JMSCC0108 messages	
	CWSJY0003W warning messages in WebSphere Application Server SystemOut.log file	
	J2CA0027E messages containing the error The method 'xa_end' has failed with	
	errorCode '100'	175
	2035 MQRC_NOT_AUTHORIZED when connecting to IBM MQ from WebSphere Application	
	Server	176
	Problem determination for the IBM MQ resource adapter	
	Using IBM MQ connection property override	
Tro	ubleshooting IBM MQ.NET problems	
	solving problems with IBM MQ MQI clients	
IVC.	ECONNRESET error over an IBM MQ MQI client channel connection	
Мп	lticast troubleshooting	
ı ıu	Testing multicast applications on a non-multicast network	
	Setting the appropriate network for multicast traffic	190
	Setting the appropriate network for multicast traffic	190 191
O	Setting the appropriate network for multicast traffic	190 191 191
Qu	Setting the appropriate network for multicast traffic	190 191 191 193
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195 198
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195 198 199
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195 198 199
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200 201 201
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200 201 201 201
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200 201 201 202 202
Qu	Setting the appropriate network for multicast traffic	190 191 193 195 198 199 200 201 201 202 202 203
Qu	Setting the appropriate network for multicast traffic	190 191 191 193 195 198 199 200 201 202 203 203 203

A queue manager does not rejoin the cluster	204
Workload balancing set on a cluster-sender channel is not working	
Out of date information in a restored cluster	205
Cluster queue manager force removed from a full repository by mistake	206
Possible repository messages deleted	
Two full repositories moved at the same time	
Unknown state of a cluster	
What happens when a cluster queue manager fails	208
What happens when a repository fails	
What happens if a cluster queue is disabled for MQPUT	
Potential issues when switching transmission queues	
Queue managers troubleshooting	211
Undelivered messages troubleshooting	212
TLS troubleshooting information	212
Troubleshooting RDQM configurations	220
RDQM HA architecture	221
Example RDQM HA configurations and errors	222
MQ Telemetry troubleshooting	230
Location of telemetry logs, error logs, and configuration files	230
Tracing the telemetry (MQXR) service	232
Resolving problem: MQTT client does not connect	234
Resolving problem: MQTT client connection dropped	236
Resolving problem: Lost messages in an MQTT application	237
Resolving problem: Telemetry (MQXR) service does not start	239
Resolving problem: JAAS login module not called by the telemetry service	240
Recovering after failure	242
Disk drive failures	243
Damaged queue manager object	244
Damaged single object	245
Automatic media recovery failure	245
Example recovery procedures on z/OS	245
Notices	269
Programming interface information	
Trademarks	

IBM MQ Troubleshooting and support

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

For an introduction to troubleshooting and support, see "Troubleshooting overview" on page 7.

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 UNIX, Linux, and Windows" on page 9
- IBM i "Making initial checks on IBM i" on page 18
- Making initial checks on z/OS" on page 27

For information about solving problems, see "Making initial checks" on page 8.

For information about solving problems for MQ Telemetry, see "MQ Telemetry troubleshooting" on page 230.

For information about solving problems when you are using channel authentication records, see <u>"Channel authentication records troubleshooting"</u> on page 159.

Information that is produced by IBM MQ can help you to find and resolve problems. For more information, see the following topics:

- "Using error logs" on page 42
- "Using trace" on page 63
- Problem determination on z/OS" on page 106
- "First Failure Support Technology (FFST)" on page 51

For information about recovering after a problem, see "Recovering after failure" on page 242.

See also "Collecting troubleshooting information" on page 41.

If an IBM 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 <u>Messages and</u> reason codes.

Related tasks

Troubleshooting and support 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" on page 7
- 2. "Re-creating the problem" on page 8
- 3. "Eliminating possible causes" on page 8

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 UNIX, Linux, and Windows" on page 9.

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.

Making initial checks

There are some initial checks that you can make that may provide answers to common problems that you may have.

About this task

Use the information and general advice given in the subtopics to help you to carry out the initial checks for your platform and rectify the problem.

Procedure

- Carry out the initial checks for your platform:
 - "Making initial checks on UNIX, Linux, and Windows" on page 9

- "Making initial checks on z/OS" on page 27
- IBM i "Making initial checks on IBM i" on page 18

Tips for system administrators

- Check the error logs for messages for your operating system:
 - "Error logs on UNIX, Linux, and Windows" on page 44
 - "Error logs on IBM i" on page 47
 - "Diagnostic information produced on IBM MQ for z/OS" on page 112
- Check the contents of qm. ini for any configuration changes or errors.

For more information on changing configuration information, see:

- Changing configuration information on UNIX, Linux, and Windows
- IBM i Changing configuration information on IBM i
- Customizing your queue managers on z/OS
- 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 63.

Tips for application developers

- Check the return codes from the MQI calls in your applications.
 - For a list of reason codes, see API completion and reason codes. 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 63. For more information about handling errors in MQI applications, see Handling program errors.

Related concepts

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

Troubleshooting and support reference

ULW Making initial checks on UNIX, Linux, and Windows

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

About this task

The cause of your problem could be in:

- IBM MO
- · The network
- The application

• Other applications that you have configured to work with IBM MQ

Procedure

• Consider the following list of questions.

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 MQ run successfully before?" on page 10
- "Have any changes been made since the last successful run?" on page 11
- "Are there any error messages or return codes to explain the problem?" on page 11
- "Can you reproduce the problem?" on page 12
- "Are you receiving an error code when creating or starting a queue manager on Windows?" on page 12
- "Does the problem affect only remote queues?" on page 12
- "Have you obtained incorrect output?" on page 12
- "Are some of your queues failing?" on page 15
- "Have you failed to receive a response from a PCF command?" on page 15
- "Has the application run successfully before?" on page 16
- "Is your application or system running slowly?" on page 17
- "Does the problem affect specific parts of the network?" on page 17
- "Does the problem occur at specific times of the day?" on page 18
- "Is the problem intermittent?" on page 18

Related tasks

"Making initial checks on z/OS" on page 27

Before you start problem determination in detail on z/OS, 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.

"Making initial checks on IBM i" on page 18

Before you start problem determination in detail on IBM i, 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.

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

Troubleshooting and support reference

Related reference

Messages and reason codes

PCF reason codes

Has IBM MQ run successfully before?

If IBM MQ has not run successfully before, it is likely that you have not yet set it up correctly. See Installing IBM MQ 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 your IBM MQ installation* for the platform, or platforms, that your enterprise use.

Also look at Configuring for information about post-installation configuration of IBM MQ.

Have any changes been made since the last successful run?

Changes that have been made to your IBM MQ configuration, maintenance updates, or changes to other programs that interact with IBM MQ could be the cause of your problem.

When you are considering changes that might recently have been made, think about the IBM 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 IBM 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 IBM MQ? For example, have you modified the Windows Registry.

Have you applied any maintenance updates?

If you have applied a maintenance update to IBM 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 IBM 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 IBM MO 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 MO 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.

IBM MQ also logs errors in the Windows Application Event Log. On Windows, check if the Windows Application Event Log shows any IBM 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 UNIX, Linux, and Windows" on page 44

For each IBM MQ Message Queue Interface (MQI) and IBM 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 (MOI) 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.

Detailed information on return codes is contained within the description of each MQI call.

Related tasks

Troubleshooting and support reference

Related reference

Diagnostic messages: AMQ4000-9999

PCF reason codes

Transport Layer Security (TLS) return codes

WCF custom channel exceptions

z/0S

IBM MQ for z/OS messages, completion, and reason codes

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 IBM 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 on Windows?

If the IBM 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 MQ Windows service is running has insufficient rights.

Ensure that the user with which the IBM MQ Windows service is configured has the rights described in User rights required for an IBM MQ Windows Service. By default this service is configured to run as the MUSR_MQADMIN user. For subsequent installations, the Prepare IBM 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 (MsqId 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 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 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 MsqId of that message. However, in the meantime, another transaction issues a successful MOGET 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 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 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 <u>Using the dead-letter (undelivered message) queue</u> and <u>MQDLH - Dead-letter header for information about the dead-letter queue header structure.</u>

• 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 Distributed queuing and clusters 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 <u>Data conversion</u> 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 MOOPEN 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 dspmgcsv 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 strmgcsv 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 and Using the dead-letter (undelivered message) queue 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 UNIX, Linux, and Windows" on page 46 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 <u>WaitInterval</u> (MQLONG) for information about the <code>WaitInterval</code> 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 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 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 IBM 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 IBM MQ systems?

Could it be that there is something different about the way that this IBM 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 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 IBM MQ applications.

Common programming errors

The errors in the following list illustrate the most common causes of problems encountered while running IBM MQ programs. Consider the possibility that the problem with your IBM 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 IBM 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 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.

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 IBM 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 IBM 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 IBM MQ network extends across more than one time zone, peak system loading might seem to occur at some other time of day.)

ULW 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.

IBM i Making initial checks on IBM i

Before you start problem determination in detail on IBM i, 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.

About this task

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 MQ product
- · Your IBM MQ application
- · Other applications
- · Site operating procedures

Some preliminary questions for you to consider are listed in the following procedure. If you are able to find the cause of the problem by working through these preliminary checks, you can then, if needed, use the information in other sections of the IBM MQ product documentation, and in the libraries of other licensed programs, to help you resolve the problem.

If you are not able to identify the cause of the probllem by carrying out the preliminary checks, and so need to carry out a more detailed investigation there are further questions for you to consider in the subtopics. As you work through the lists of questions, 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.

Procedure

· Consider the following questions.

The following steps are intended to help you isolate the problem and are taken from the viewpoint of an IBM MQ application. Check all the suggestions at each stage.

1. Has IBM MQ for IBM i run successfully before?

Proceed to Step "2" on page 19.

No

It is likely that you have not installed or set up IBM MQ correctly.

2. Has the IBM MQ application run successfully before?

Yes

Proceed to Step "3" on page 19.

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, for information about 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 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 MQI Applications reference for your programming language for a description of each return code.

d. If you have applied any PTF to IBM 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 <u>"4" on page</u> 20.

4. Has the server system remained unchanged since the last successful run?

Yes

Proceed to "Identifying characteristics of the problem on IBM i" on page 21.

No

Consider all aspects of the system and review the appropriate documentation on how the change might have affected the IBM MQ application. For example :

- Interfaces with other applications
- Installation of new operating system or hardware
- Application of PTFs
- Changes in operating procedures

What to do next

Related tasks

"Manually applying required authority for commands and programs" on page 23

Some IBM 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 MQ program code, for example a queue manager, relies on using IBM i system programs.

"Making initial checks on UNIX, Linux, and Windows" on page 9

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

"Making initial checks on z/OS" on page 27

Before you start problem determination in detail on z/OS, 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.

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

Related reference

"Determining problems with applications, commands and messages" on page 24

If you encounter problems with IBM MQ applications, commands, and messages, there are a number of questions that you can consider to help you determine the cause of the problem.

Messages and reason codes

PCF reason codes

Troubleshooting and support reference

Identifying characteristics of the problem on IBM i

If you have not been able to identify the cause of the problem by using the preliminary checks, you should now start to look at the characteristics of the problem in greater detail.

Use the following questions as pointers to help you to identify the cause of the problem:

- "Can you reproduce the problem?" on page 21
- "Is the problem intermittent?" on page 21
- "Problems with commands" on page 21
- "Does the problem affect all users of the IBM MQ for IBM i application?" on page 22
- "Does the problem affect specific parts of the network?" on page 22
- "Does the problem occur only on IBM MQ?" on page 22
- "Does the problem occur at specific times of the day?" on page 22
- "Have you failed to receive a response from a command?" on page 22

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 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 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 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 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 IBM MQ?

If the problem occurs only on this version of IBM MQ, check the appropriate database on RETAIN, or the https://www.ibm.com/support/entry/portal/Overview/Software/WebSphere // WebSphere_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 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 DSPMOMCSVR 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 STRMOMCSVR command to start it.
 - If the response to the command indicates that the SYSTEM.ADMIN.COMMAND.QUEUE is not enabled for MOGET requests, enable the queue for MOGET 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 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 47 for further information.

- Are the gueues 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 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 MsqId fields correctly?

Set the values of MsqId and CorrelId in your application to ensure that you receive all messages from the queue.

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MO applications, use the techniques described to help you diagnose and solve the problems.

Related tasks

"Manually applying required authority for commands and programs" on page 23 Some IBM 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 MQ program code, for example a queue manager, relies on using IBM i system programs.

Related reference

"Determining problems with applications, commands and messages" on page 24 If you encounter problems with IBM MQ applications, commands, and messages, there are a number of questions that you can consider to help you determine the cause of the problem.

IBM I Manually applying required authority for commands and programs

Some IBM 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 MQ program code, for example a queue manager, relies on using IBM i system programs.

About this task

To enable this reliance, the commands and programs must either have *PUBLIC *USE authority, or explicit *USE authority to the IBM MQ user profiles QMQM and QMQMADM.

Such authority is applied automatically as part of the installation process, and you do not need to apply it yourself. However, if you encounter problems, you can set the authorities manually as described in the following steps.

Procedure

1. Set the authorities for commands using GRTOBJAUT with an OBJTYPE(*CMD) parameter, for example:

GRTOBJAUT OBJ(QSYS/ADDLIBLE) OBJTYPE(*CMD) USER(QMQMADM) AUT(*USE)

You can set authorities for the following commands:

- QSYS/ADDLIBLE
- OSYS/ADDPFM
- · QSYS/CALL
- · QSYS/CHGCURLIB
- · QSYS/CHGJOB
- QSYS/CRTJRN
- QSYS/CRTJRNRCV
- QSYS/CRTJOBQ
- QSYS/CRTJOBD
- · QSYS/CRTLIB
- QSYS/CRTMSGQ
- QSYS/CRTPF
- QSYS/CRTPGM
- QSYS/CRTSRCPF
- QSYS/DLTJRN
- QSYS/DLTJRNRCV
- QSYS/DLTLIB
- QSYS/DLTMSGQ
- QSYS/OVRPRTF
- QSYS/RCLACTGRP
- QSYS/RTVJRNE
- QSYS/RCVJRNE
- OSYS/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)

You can set authorities for the following programs:

- QSYS/QWTSETP(*PGM)
- QSYS/QSYRLSPH(*PGM)
- QSYS/QSYGETPH(*PGM)

Determining problems with applications, commands and messages

If you encounter problems with IBM MQ applications, commands, and messages, there are a number of questions that you can consider to help you determine the cause of the problem.

Use the following questions as pointers to help you to identify 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 WRKMQMQSTS 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 Distributed queuing and clusters.
- 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 (MsqId or CorrelId)?

Check that you are waiting for a message with the correct <code>MsgId</code> or <code>CorrelId</code>. 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 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 26.

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 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. For information about the

dead-letter header structure, see MQDLH - Dead-letter header. For IBM i, see also IBM i Application Programming Reference (ILE/RPG).

• 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 Distributed queuing and clusters.

z/0S Making initial checks on z/OS

Before you start problem determination in detail on z/OS, 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.

About this task

The cause of your problem could be in:

- IBM MQ
- The network
- The application
- Other applications that you have configured to work with IBM MQ

Procedure

- Consider the following list of questions. 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 MQ for z/OS run successfully before?" on page 28
 - "Have you applied any APARs or PTFs?" on page 29
 - "Are there any error messages, return codes or other error conditions?" on page 29
 - "Has your application or IBM MQ for z/OS stopped processing work?" on page 31
 - "Is there a problem with the IBM MQ queues?" on page 31
 - "Are some of your queues working?" on page 32
 - "Are the correct queues defined?" on page 33
 - "Does the problem affect only remote or cluster queues?" on page 33
 - "Does the problem affect only shared queues?" on page 33

- "Does the problem affect specific parts of the network?" on page 34
- "Problems that occur at specific times of the day or affect specific users" on page 34
- "Is the problem intermittent or does the problem occur with all z/OS, CICS, or IMS systems?" on page 35
- "Has the application run successfully before?" on page 35
- "Have any changes been made since the last successful run?" on page 36
- "Do you have a program error?" on page 37
- "Has there been an abend?" on page 38
- "Have you obtained incorrect output?" on page 39
- "Can you reproduce the problem?" on page 39
- "Have you failed to receive a response from an MQSC command?" on page 39
- "Is your application or IBM MQ for z/OS running slowly?" on page 41

Related tasks

"Making initial checks on UNIX, Linux, and Windows" on page 9

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

"Making initial checks on IBM i" on page 18

Before you start problem determination in detail on IBM i, 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.

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

Related reference

Messages and reason codes

PCF reason codes

Troubleshooting and support reference

Has IBM MQ for z/OS run successfully before?

Knowing whether IBM 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 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 <u>Installing the IBM MQ for</u> z/OS product 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 MQ as an installed subsystem. To determine if IBM 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 IBM MQ. These fixes can have been applied to IBM MQ or to other z/OS systems.

If an APAR or PTF has been applied to IBM 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 IBM 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 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 MQ for z/OS messages, completion, and reason codes 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.

Unusual messages

Be aware of unusual messages associated with the startup of IBM 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 <u>Return codes</u> 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 MQ or from another program.

IBM MQ error messages

IBM MQ for z/OS error messages are prefixed with the letters CSQ.

If you get an unexpected IBM MQ error message (for example, in the console log, or the CICS log), see IBM MQ for z/OS messages, completion, and reason codes for an explanation.

IBM MQ for z/OS messages, completion, and reason codes 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 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)

Unexpected return codes

If your application has received an unexpected return code from IBM MQ, see <u>Return codes</u> for information about how your application can handle IBM 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 <u>DISPLAY THREAD</u>. 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, IBM 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 IBM 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 IBM MQ, find out whether a system dump was produced, see IBM MQ dumps.

IBM MQ for z/OS might still be running

Consider also that IBM 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 <u>Is your application or IBM MQ for z/OS</u> running slowly. Refer to Dealing with performance problems for advice about what to do next.

Has your application or IBM 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.

IBM MQ problems

Your queues might have become disabled for MQPUT or MQGET calls, the dead-letter queue might be full, or IBM 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 IBM MQ queues?

Use this topic for investigating potential problems with IBM 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 <u>"Are</u> some of your queues working?" on page 32.

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 IBM 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

Refer to "Collecting troubleshooting information" on page 41 for further guidance.

Related concepts

"Are the correct queues defined?" on page 33

IBM MQ requires certain predefined queues. Problems can occur if these queues are not defined correctly.

"Does the problem affect only remote or cluster queues?" on page 33

Use this topic for further investigation if the problem only occurs on remote or cluster queues.

"Does the problem affect only shared queues?" on page 33

Use this topic to investigate possible queue sharing group issues which can cause problems for shared queues.

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 gueue 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) OPPROCS IPPROCS to see how many tasks are putting messages on to, and getting messages from the queue. In a queue-sharing environment, check OPPROCS 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 MOOPEN 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 IBM 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

Refer to "Collecting troubleshooting information" on page 41 for further guidance.

Are the correct queues defined?

IBM 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 MOOPEN 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 <u>"Dealing</u> with incorrect output on z/OS" on page 143).

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 Configuring distributed queuing for more information.

Look for error messages on the startup of the channel initiator and listener. See <u>IBM MQ for z/OS</u> messages, completion, and reason codes and Configuring distributed queuing 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 <u>Configuring distributed queuing</u> 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
CSOU503I
            QSG=SQ02, DB2 DSG=DSN710P5, DB2 ssid=DFP5
            XCF group CSQGSQ02 already defined
Summary information for XCF group CSQGSQ02
Member=MQ04, state=QUIESCED, system=MV4A
CSÕU517I
CSQU520I
CSQU522I
CSQU523I
            User data=D4E5F4C15AD4D8F0F4404040C4C5.
CSOU522I
            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
            QMGR MQ04 in table entry 1 not set in usage map
CSQU573E
            OMGR 27 in usage map has no entry in table
Usage map T01_ARRAY_STRUC and DB2 table CSQ.ADMIN_B_STRUCTURE inconsistent
Structure APPL2 in table entry 4 not set in usage map
CSQU574E
CSQU572E
CSQU575E
CSQU576E
            Structure 55 in usage map has no entry in table
            Usage map T03_LH_ARRAY and DB2 table CSQ.OBJ_B_QUEUE inconsistent Queue MYSQ in table entry 13 not set in usage map for structure APPL1
CSOU572E
CSQU577E
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 BYTSSENT 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 IBM 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 IBM 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 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 MQ for z/OS system has a performance problem, see "Dealing with performance problems on z/OS" on page 136.

If the problem only affects some users, is it because some users do not have the correct security authorization? See <u>User IDs for security checking</u> for information about user IDs checked by IBM 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 IBM 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 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 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 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 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 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 37 for some examples of common errors that cause problems with IBM MQ applications.

Do applications report errors from IBM 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 IBM 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 not 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 IBM 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 an IBM MQ problem.

The examples that follow illustrate the most common causes of problems encountered while running IBM 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 IBM 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 IBM 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 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 <u>"SYSUDUMP information on z/OS" on page 134</u>. 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 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 "IBM MQ for z/OS abends" on page 109 for information about the abend codes, and see Return codes for an explanation of the reason code.

Abnormal program termination

If your program has terminated abnormally, see "Dealing with abends on IBM MQ for z/OS" on page 111.

If your system has terminated abnormally, and you want to analyze the dump produced, see <u>"IBM MQ for z/OS dumps"</u> on page 117. This section tells you how to format the dump, and how to interpret the data contained in it.

Z/0S 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

IBM 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 IBM MQ that you were not expecting, refer to "Are there any error messages, return codes or other error conditions?" on page 29.

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 z/OS" on page 143.

Z/0S Can you reproduce the problem?

Reproducing the problem can be used to assist problem determination for IBM 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 IBM 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.

It is a second to the secon

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 <u>Reason (MQLONG)</u> 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 <u>WaitInterval</u> (MQLONG) and <u>MQGET - Get message</u> 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 <u>CSQ0016E</u> 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 <u>"IBM MQ for z/OS dumps" on</u> page 117.

Were any error messages issued?

Check to see if any error messages were issued that might indicate the nature of the error.

See <u>Issuing commands</u> for information about the different methods you can use to enter MQSC commands.

Is your application or IBM MQ for z/OS running slowly?

Slow applications can be caused by the application itself or underlying software including IBM 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 MQ for z/OS running slowly?

The following symptoms might indicate that IBM 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 <u>"Dealing with applications that are running slowly or have stopped on z/OS" on page 137</u>, and on dealing with performance problems in <u>"Dealing with performance problems on z/OS"</u> on page 136.

Collecting troubleshooting information

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

About this task

The IBM MQ Support pages within the IBM Support Site are:

- Multi IBM MQ for Multiplatforms Support web page
- IBM MQ for z/OS Support web page

To receive notifications about IBM MQ fixes, troubleshooting and other news, you can <u>subscribe to</u> notifications.

For more information, including how to register for support, see the IBM Support Guide.

Note: Running the **runmqras** command will help you in collecting troubleshooting information. For more information, see runmqras (collect IBM MQ troubleshooting information).

Procedure

1. Determine the business severity level for the problem.

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.

2. Describe the problem and gather background information.

You might find the information you need in your own in-house tracking system for problems.

Be as specific as possible. Include all relevant background information. Know the answers to these questions:

- What was the source of the problem within your system software; that is, the program that seems to be the cause of the problem.
- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms?
- 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 changes
 - Operating system upgrades
 - Networking software updates
 - Changes in the level of licensed programs
 - PTFs applied
 - Additional features used
 - Application programs changed
 - Unusual operator action
 - **Z/OS** Regenerations
- Are you currently using a workaround for this problem?

Using error logs

There are a variety of error logs that you can use to help with problem determination and troubleshooting. On Multiplatforms, use the following links to find out about the error logs available for your platform and how to use them:

- "Error logs on UNIX, Linux, and Windows" on page 44
- Error logs on IBM i" on page 47
- On z/OS error messages are written to:
- The z/OS system console
- The channel-initiator job log

For information about error messages, console logs, and dumps on IBM MQ for z/OS, see <u>Problem</u> determination on z/OS.

Suppressing or excluding messages from error logs

It is possible to suppress or exclude some messages on both Multiplatforms and z/OS systems.:

- Multi For details of suppressing some messages on Multiplatforms, see "Suppressing channel error messages from error logs on Multiplatforms" on page 50.
- Z/05 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 IBM MQ for z/OS concepts.

AMQ_DIAGNOSTIC_MSG_SEVERITY



From IBM MQ 9.0.3, if you set the environment variable **AMQ_DIAGNOSTIC_MSG_SEVERITY**, for an IBM MQ process, when that IBM MQ process writes a message to an error log or to the console, the message severity is appended to the message number as a single uppercase alphabetic character as follows:

Table 1.		
Type of message	Character	
Informational (0)	I	
Warning (10)	W	
Error (20 or 30)	Е	
Severe (40)	S	
Termination (50)	Т	

For example:

```
AMQ5051I: The queue manager task 'LOGGER-IO' has started.
AMQ7075W: Unknown attribute foo at /var/mqm/qmgrs/QM1/qm.ini in
the configuration data.
AMQ9510E: Messages cannot be retrieved from a queue.
AMQ8506S: Command server MQGET failed with reason code 2009.
AMQ8301T: IBM MQ storage monitor job could not be started.
```

Notes:

- 1. Because the queue manager writes messages, the environment variable has to be set in the environment where the queue manager is started. This is especially important on Windows, where it might be the Windows service that starts the queue manager.
- 2. AMO_DIAGNOSTIC_MSG_SEVERITY also affects messages printed by a program.

V 9.0.4

From IBM MQ 9.0.4, the behavior that **AMQ_DIAGNOSTIC_MSG_SEVERITY** enables, is set by default. You can turn off this behavior by setting the environment variable to 0.

Note that the new services always add the severity character.



ISO 8601 Time

When IBM MQ processes write a message to an error log, the message time in ISO 8601 format, in Coordinated Universal Time (UTC), is included as a Time () attribute.

For example, where the Z time zone indicates UTC:

Rename on Rollover



Prior to IBM MQ 9.0.4, when AMQERR01.LOG reaches the maximum configured size, AMQERR02.LOG is renamed to be AMQERR03.LOG.

Then, the contents of AMQERR01.LOG are copied into AMQERR02.LOG, and AMQERR01.LOG is truncated to empty. This meant that it was possible for certain tools to miss messages that the tool has not processed, before those messages were copied into AMQERR02.LOG.

From IBM MQ 9.0.4, the logic has been changed, so that AMQERR01.LOG is renamed to AMQERR02.LOG.

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Troubleshooting overview" on page 7

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?"

"First Failure Support Technology (FFST)" on page 51

First Failure Support Technology (FFST) for IBM MQ provides information about events that, in the case of an error, can help IBM support personnel to diagnose the problem.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

Error logs on UNIX, Linux, and Windows

The errors subdirectory, which is created when you install IBM MQ, can contain up to three error log files.

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\MQ\ file path under Windows systems. The errors subdirectory can contain up to three error log files named:

- AMQERR01.LOG
- AMQERR02.LOG
- AMOERR03.LOG

For more information about directories where log files are stored, see <u>"Error log directories on UNIX, Linux, and Windows"</u> on page 46.

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 AMQERRO3, and each has a default capacity of V 9.0.4 32 MB (33554432 bytes). The capacity can be altered in the Extended queue manager properties page from the IBM 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 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\MQ\qmgrs\ qmname \errors file path under Windows systems.

V 9.0.4 As error messages are generated, they are placed in AMQERR01. When AMQERR01 gets bigger than 32 MB it is renamed to AMQERR02.

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 1 on page 45 shows an extract from an IBM 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 mgm 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.

Ignoring error codes under Windows systems

On Windows systems, the error message is written to both the IBM MQ error log and the Windows Application Event Log. The error messages written to the Application Event Log includes messages of error severity, warning severity and information severity. 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:

where MQ_INSTALLATION_NAME is the installation name associated with a particular installation of IBM 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 IBM 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

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

"Problem determination on z/OS" on page 106

IBM MQ for z/OS, CICS, Db2, and IMS produce diagnostic information which can be used for problem determination.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

Related reference

"Error logs on IBM i" on page 47

Use this information to understand the IBM MQ for IBM i error logs.

Error log directories on UNIX, Linux, and Windows

IBM MQ uses a number of error logs to capture messages concerning its own operation of IBM 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 IBM MQ is installed.

• If the queue manager name is known, the location of the error log is shown in Table 2 on page 46.

Table 2. Queue manager error log directory		
Platform	Directory	
Linux UNIX UNIX and Linux systems	/var/mqm/qmgrs/ qmname /errors	
Windows Windows systems	<pre>MQ_INSTALLATION_PATH\QMGRS\ qmname \ERRORS\AMQERR01.LOG</pre>	

• If the queue manager name is not known, the location of the error log is shown in Table 3 on page 47.

Table 3. System error log directory		
Platform	Directory	
Linux UNIX UNIX and Linux systems	/var/mqm/errors	
Windows Windows systems	<pre>MQ_INSTALLATION_PATH\QMGRS\@SYSTEM\ERRORS\AMQERR01 .LOG</pre>	

• If an error has occurred with a client application, the location of the error log on the client is shown in Table 4 on page 47.

Table 4. Client error log directory		
Platform	Directory	
Linux UNIX UNIX and Linux systems	/var/mqm/errors	
Windows Windows systems	MQ_DATA_PATH\ERRORS\AMQERR01.LOG	

In IBM 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. IBM 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\MQ).

If IBM 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\MQ, errors are logged in C:\Program Files\IBM\MQ\errors.

For further information about configuration files, see <u>Changing IBM MQ and queue manager configuration</u> information.

Note: Errors in the Windows Registry are notified by messages when a queue manager is started.

Error logs on IBM i

Use this information to understand the IBM MQ for IBM i error logs.

By default, only members of the QMQMADM group can access error logs. To give users access to error logs, who are not members of this group, set **ValidateAuth** to *No* and grant those users *PUBLIC authority. See Filesystem for more information.

IBM MQ uses a number of error logs to capture messages concerning the operation of IBM 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/mgm/errors'

Alternatively, you can use option 23 against the queue manager from the WRKMOM 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 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 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 MQ error log

Figure 2 on page 49 shows a typical extract from an IBM MQ error log.

```
EXPLANATION:
Cause . . . . : The repository manager ended normally. Recovery . . : None.
Technical Description . . . . . . : None.
07/19/02 11:15:57 AMQ9542: Queue manager is ending.
EXPLANATION:
                   The program will end because the queue manager is quiescing.
Cause . . . . :
           . . : Ine p
Recovery
Technical Description . . . . . : None. ---- amqrimna.c : 773 -----
07/19/02 11:16:00 AMQ8004: IBM MQ queue manager 'mick' ended.
EXPLANATION:
Cause . . . . : IBM MQ queue manager 'mick' ended. Recovery . . : None.
                    None.
Technical Description . . . . . . . . . None.
07/19/02 11:16:48 AMQ7163: IBM MQ job number 18429 started.
EXPLANATION:
                   This job has started to perform work for Queue Manager
Cause . . . . :
                    mick, The job's PID is 18429 the CCSID is 37. The job name is
                    582775/MQUSER/AMQZXMAO.
Recovery . . . : None
07/19/02 11:16:49 AMQ7163: IBM 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: IBM 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: IBM MQ job number 18432 started.
EXPLANATION:
                   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
Cause . . . . :
                    582778/MQUŠER/AMQALMPX.
Recovery . . . : None
```

Figure 2. Extract from an IBM MQ error log

Related concepts

"Error logs on UNIX, Linux, and Windows" on page 44

The errors subdirectory, which is created when you install IBM MQ, can contain up to three error log files.

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

"Problem determination on z/OS" on page 106

IBM MQ for z/OS, CICS, Db2, and IMS produce diagnostic information which can be used for problem determination.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

Error logs in IBM MQ classes for JMS

Information about runtime problems that might require corrective action by the user is written to the IBM MQ classes for JMS log.

For example, if an application attempts to set a property of a connection factory, but the name of the property is not recognized, IBM MQ classes for JMS writes information about the problem to its log.

By default, the file containing the log is called mqjms.log and is in the current working directory. However, you can change the name and location of the log file by setting the com.ibm.msg.client.commonservices.log.outputName property in the IBM MQ classes for JMS configuration file. For information about the IBM MQ classes for JMS configuration file, see The IBM MQ classes for JMS configuration file, and for more detail about valid values for the com.ibm.msg.client.commonservices.log.outputName property, see "Logging errors for IBM MQ classes for JMS" on page 165.

Suppressing channel error messages from error logs on Multiplatforms

You can prevent selected messages from being sent to the error logs for a specified time interval, for example if your IBM MQ system produces a large number of information messages that fill the error logs.

About this task

There are two ways of suppressing messages for a given time interval:

- By using SuppressMessage and SuppressInterval in the QMErrorLog stanza in the qm.ini file.
- By using the environment variables MQ_CHANNEL_SUPPRESS_MSGS and MQ_CHANNEL_SUPPRESS_INTERVAL.

Procedure

• To suppress messages for a given time interval by using the QMErrorLog stanza in the qm.ini file, specify the messages that are to be written to the queue manager error log once only during a given time interval with SuppressMessage, and specify the time interval for which the messages are to be suppressed with SuppressInterval.

For example, to suppress the messages AMQ9999, AMQ9002, AMQ9209 for 30 seconds, include the following information in the QMErrorLog stanza of the qm.ini file:

SuppressMessage=9001,9002,9202 SuppressInterval=30 Alternatively, instead of editing the qm.ini file directly, you can use the Extended Queue Manager properties page in IBM MQ Explorer to exclude and suppress messages.

- To suppress messages for a given time interval by using the environment variables MQ_CHANNEL_SUPPRESS_MSGS and MQ_CHANNEL_SUPPRESS_MSGS, complete the following steps:
 - a) Specify the messages that are to be suppressed with MQ_CHANNEL_SUPPRESS_MSGS.

You can include up to 20 channel error message codes in a comma-separated list. There is no comprehensive list of message ids that can be included in the **MQ_CHANNEL_SUPPRESS_MSGS** environment variable. However, the message ids must be channel messages (that is <u>AMQ9xxx</u>: messages).

The following examples are for messages AMQ9999, AMQ9002, AMQ9209.



b) Specify the time interval for which the messages are to be suppressed with **MQ_CHANNEL_SUPPRESS_INTERVAL**.

The default value is 60,5 which means that after the first five occurrences of a given message in a 60 second interval, any further occurrences of that message are suppressed until the end of that 60 second interval. A value of 0,0 means always suppress. A value of 0,n where n > 0 means never suppress.

Related concepts

QMErrorLog stanza on UNIX, Linux, and Windows QMErrorLog stanza on IBM i

Related reference

Environment variables descriptions

Queue manager properties

First Failure Support Technology (FFST)

First Failure Support Technology (FFST) for IBM MQ provides information about events that, in the case of an error, can help IBM support personnel to diagnose the problem.

First Failure Data Capture (FFDC) provides an automated snapshot of the system environment when an internal event occurs. In the case of an error, this snapshot is used by IBM support personnel to provide a better understanding of the state of the system and IBM MQ when the problem occurred.

The information about an event is contained in an FFST file. In IBM MQ, FFST files have a file type of FDC. FFST files do not always indicate an error. An FFST might be informational.

Monitoring and housekeeping

Here are some tips to help you with managing FFST events:

- Monitor FFST events for your system, and ensure that appropriate and timely remedial action is taken
 when an event occurs. In some cases, the FDC files might be expected and can therefore be ignored,
 for example FFST events that arise when IBM MQ processes are ended by the user. By appropriate
 monitoring, you can determine which events are expected, and which events are not.
- FFST events are also produced for events outside IBM MQ. For example, if there is a problem with the IO subsystem or network, this problem can be reported in an FDC type file. These types of event are outside the control of IBM MQ and you might need to engage third parties to investigate the root cause.

• Ensure that good housekeeping of FFST files is carried out. The files must be archived and the directory or folder must be cleared to ensure that only the most recent and relevant FDC files are available, should the support team need them.

Use the information in the following links to find out the names, locations, and contents of FFST files in different platforms.

- "FFST: IBM MQ classes for JMS" on page 52
- "FFST: IBM MQ for Windows" on page 57
- "FFST: IBM MQ for UNIX and Linux systems" on page 59
- IBM i "FFST: IBM MQ for IBM i" on page 61

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Troubleshooting overview" on page 7

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?"

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

"Problem determination on z/OS" on page 106

IBM MQ for z/OS, CICS, Db2, and IMS produce diagnostic information which can be used for problem determination.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

FFST: IBM MQ classes for JMS

Describes the name, location, and contents of the First Failure Support Technology (FFST) files that are generated by the IBM MQ classes for JMS.

When using the IBM MQ classes for JMS, FFST information is recorded in a file in a directory that is called FFDC, which by default is a subdirectory of the current working directory for the IBM MQ classes for JMS application that was running when the FFST was generated. If the property com.ibm.msg.client.commonservices.trace.outputName has been set in the IBM MQ classes for JMS configuration file, the FFDC directory is a subdirectory of the directory that the property points to. For information about the IBM MQ classes for JMS , see The IBM MQ classes for JMS configuration file.

An FFST file contains one FFST record. 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 an internal error within the IBM MQ classes for JMS.

FFST files are named JMSC *nnnn*. FDC, where *nnnn* starts at 1. If the full file name already exists, this value is incremented by one until a unique FFST file name is found.

An instance of an IBM MQ classes for JMS application writes FFST information to multiple FFST files. If multiple errors occur during a single execution of the application, each FFST record is written to a different FFST file.

Sections of an FFST record

An FFST record that is generated by the IBM MQ classes for JMS contains the following sections:

The header

A header, indicating the time when the FFST record was created, the platform that the IBM MQ classes for JMS application is running on, and the internal method that was being called. The header also contains a probe identifier, which uniquely identifies the place within the IBM MQ classes for JMS that generated the FFST record.

Data

Some internal data that is associated with the FFST record.

Version information

Information about the version of the IBM MQ classes for JMS being used by the application that generated the FFST record.

Stack Trace

The Java stack trace for the thread that generated the FFST record.

Property Store Contents

A list of all of the Java system properties that have been set on the Java Runtime Environment that the IBM MQ classes for JMS application is running in.

WorkQueueMananger Contents

Information about the internal thread pool that is used by the IBM MQ classes for JMS.

Runtime properties

Details about the amount of memory and the number of processors available on the system where the IBM MQ classes for JMS application is running.

Component Manager Contents

Some information about the internal components that are loaded by the IBM MQ classes for JMS.

Provider Specific information

Information about all of the active JMS Connections, JMS Sessions, MessageProducer, and MessageConsumer objects currently being used by the IBM MQ classes for JMS application that was running when the FFST was generated. This information includes the name of the queue manager that JMS Connections and JMS Sessions are connected to, and the name of the IBM MQ queue or topic objects that are being used by MessageProducers and MessageConsumers.

All Thread information

Details about the state of all of the active threads in the Java Runtime Environment that the IBM MQ classes for JMS application was running in when the FFST record was generated. The name of each thread is shown, together with a Java stack trace for every thread.

Example FFST log file

```
c:\JBoss-6.0.0\bin\FFDC\JMSCC0007.FDC PID:4472

JMS Common Client First Failure Symptom Report

Product :- IBM MQ classes for JMS
Date/Time :- Mon Feb 03 14:14:46 GMT 2014
System time :- 1391436886081
Operating System :- Windows Server 2008
UserID :- pault
Java Vendor :- IBM Corporation
Java Version :- 2.6

Source Class :- com.ibm.msg.client.commonservices.j2se.wmqsupport.PropertyStoreImpl
Source Method :- getBooleanProperty(String)
ProbeID :- XS002005
Thread :- name=pool-1-thread-3 priority=5 group=workmanager-threads
ccl=BaseClassLoader@ef1c3794{vfs://C:/JBoss-6.0.0/server/default/deploy/basicMDB.ear}

Data
----
```

```
name :- com.ibm.mq.connector.performJavaEEContainerChecks
Version information
Java Message Service Client
7.5.0.2
p750-002-130627
Production
IBM MQ classes for Java Message Service
7.5.0.2
p750-002-130627
Production
IBM MQ JMS Provider
7.5.0.2
p750-002-130627
Production
Common Services for Java Platform, Standard Edition
7.5.0.2
p750-002-130627
Production
Stack trace
Stack trace to show the location of the FFST call
   FFST Location :- java.lang.Exception
       at com.ibm.msg.client.commonservices.trace.Trace.getCurrentPosition(Trace.java:1972)
       at com.ibm.msg.client.commonservices.trace.Trace.createFFSTString(Trace.java:1911)
       at com.ibm.msg.client.commonservices.trace.Trace.ffstInternal(Trace.java:1800)
       at com.ibm.msg.client.commonservices.trace.Trace.ffst(Trace.java:1624)
       at
com.ibm.msg.client.commonservices.j2se.propertystore.PropertyStoreImpl.getBooleanProperty(
PropertyStoreImpl.java:322)
       at
com.ibm.msg.client.commonservices.propertystore.PropertyStore.getBooleanPropertyObject(Pr
opertyStore.java:302)
       at
com.ibm.mq.connector.outbound.ConnectionWrapper.jcaMethodAllowed(ConnectionWrapper.java:510)
       at
\verb|com.ibm.mq.connector.outbound.ConnectionWrapper.setExceptionListener(ConnectionWrapper.java:244)| \\
       at com.ibm.basicMDB.MDB.onMessage(MDB.java:45)
Property Store Contents
All currently set properties
   awt.toolkit
                                          :- sun.awt.windows.WToolkit
   catalina.ext.dirs
                                             :- C:\JBoss-6.0.0\server\default\lib
                                           :- C:\JBoss-6.0.0\server\default
   catalina.home
   com.ibm.cpu.endian
                                              :- little
   com.ibm.jcl.checkClassPath
   com.ibm.mq.connector.performJavaEEContainerChecks
com.ibm.oti.configuration :- s
                                                             :- false
                                                 :- scar
                                               :- 20131013_170512
   com.ibm.oti.jcl.build
   com.ibm.oti.shared.enabled
                                                  :- false
   com.ibm.oti.vm.bootstrap.library.path
                                                       :- C:\Program
Files\IBM\Java70\jre\bin\compressedrefs;C:\Program Files\IBM\Java70\jre\bin
   com.ibm.oti.vm.library.version
                                                    :- 26
                                                 :- C:\Program
   com.ibm.system.agent.path
Files\IBM\Java70\jre\bin
   com.ibm.util.extralibs.properties
   com.ibm.vm.bitmode
                                              :- 64
   com.ibm.zero.version
                                               :- 2
                                             :- Cp850
   console.encoding
   file.encoding
                                           :- Cp1252
                                             :- sun.io
  file.encoding.pkg
WorkQueueMananger Contents
   Current ThreadPool size :- 2
   Maintain ThreadPool size :- false
   Maximum ThreadPool size :- -1
```

```
| ThreadPool inactive timeout :- 0
Runtime properties
   Available processors
   Free memory in bytes (now) :- 54674936
Max memory in bytes :- 536870912
  Total memory in bytes (now) :- 235012096
Component Manager Contents
   Class Name :- class
Common Services Components:
   Class Name :- class com.ibm.msg.client.commonservices.j2se.J2SEComponent
Component Name :- com.ibm.msg.client.commonservices.j2se
Component Title :- Common Services for Java Platform, Standard Edition
Factory Class :- class com.ibm.msg.client.commonservices.j2se.CommonServicesImplementation
Version :- 7.5.0.2
   inPreferenceTo[0] :- com.ibm.msg.client.commonservices.j2me
Messaging Provider Components:
| CMVC :- p750-002-1306
   CMVC :- p750-002-130627
Class Name :- class com.ibm.msg.client.wmq.factories.WMQComponent
   Component Name :- com.ibm.msg.client.wmq
Component Title :- IBM MQ JMS Provider
Factory Class :- class com.ibm.msg.client.wmq.factories.WMQFactoryFactory
Version :- 7.5.0.2
Provider Specific Information
Overview of JMS System
Num. Connections: 3
Num. Sessions : 3
Num. Consumers : 0
Num. Producers : 0
Detailed JMS System Information
Connections
seqNumberWrap:99999999 maxMsgLength:104857600 putAuthority:1 dataConversion:0 userIdentifier:'<null>' mcaUserIdentifier:'<null>' mcaType:1 remoteUserIdentifier:'' msgRetryExit:'<null>' msgRetryUserData:'<null>' msgRetryCount:10
heartbeatInterval:1 batchInterval:0 nonPersistentMsgSpeed:2 clustersDefined:0
networkPriority:0
00000000 sslCipherSpec:'<null>' sslPeerName:'<null>' sslClientAuth:0 keepAliveInterval:-1
channelMonitoring:0 channelStatistics:0 exitNameLength:128 exitDataLength:32 sendExitsDefined:0 sendExit:'<null>'
sendExitPtr:<null> sendUserDataPtr:<null> receiveExitsDefined:0 receiveExit:'<null>'
propertyControl:0 maxInstances:999999999 maxInstancesPerClient:9999999999
clientChannelWeight:0 connectionAffinity:1 batchDataLimit:5000 useDLQ:2 defReconnect:0 ]
sslConfig:[version:1 keyRepository:'<null>' cryptoHardware:'<null>' authInfoRecCount:0
keyResetCount:0 fipsRequired:0 encryptionPolicySuiteB:(1,0,0,0)certificateValPolicy:0 ]
connectionId:414D514374657374202020202020208CA3E2522028FD02 securityParms:[<null>]
   exceptionListener
\verb|com.ibm.msg.client.jms.internal.JmsProviderExceptionListener@f17b3583|
com.ibm.msg.client.wmq.internal.WMQConsumerOwnerShadow@adabbe93
| queueManagerName
                       :- test
Sessions : 3
```

```
:- com.ibm.msg.client.wmq.internal.WMQXASession@f5c63f0a
       Parent Connection :- com.ibm.msg.client.wmq.internal.WMQXAConnection@228b45cb
      asfConsumer
                                       :- <null>
       asfDestination :- <null>
       asfSubName
                                         :- <null>
      asyncPutCounter :-
com.ibm.msg.client.wmq.internal.WMQSession$AsyncPutCounter@88db6ec0
      didRecovAsyncPut :- false
     helper
com.ibm.msg.client.wmq.internal.WMQConsumerOwnerShadow@28192ad1
     inSyncpoint
                                         :- false
     queueManagerName :- test
Consumers
Producers
All Thread Information
Name : DispatchThread:
Priority : 5
ThreadGroup : java.lang.ThreadGroup[name=JMSCCThreadPool,maxpri=10]
ID: 86
State : TIMED_WAITING
Stack : java.lang.Object.wait(Object.java:-2)
: java.lang.Object.wait(Object.java:196)
com.ibm.mq.jmqi.remote.impl.RemoteDispatchThread.waitOnSleepingEvent(RemoteDispatchThread
.java:151)
com.ibm.mq.jmqi.remote.impl.RemoteDispatchThread.sleepPhase(RemoteDispatchThread.java:636)
com.ibm.mq.jmqi.remote.impl.RemoteDispatchThread.run(RemoteDispatchThread.java:385)
com.ibm.msg.client.commonservices.workqueue.WorkQueueItem.runTask(WorkQueueItem.java:214)
com.ibm.msg.client.commonservices.workqueue.SimpleWorkQueueItem.runItem(SimpleWorkQueueIt
em.java:105)
com.ibm.msg.client.commonservices.workqueue.WorkQueueItem.run(WorkQueueItem.java:229)
com.ibm.msg.client.commonservices.workqueue.WorkQueueManager.runWorkQueueItem(WorkQueueMa
nager.java:303)
\verb|com.ibm.msg.client.commonservices.j2se.workqueue.WorkQueueManagerImplementation\\| \$ThreadPooling to the common services of the common
lWorker.run(WorkQueueManagerImplementation.java:1219)
Name : RcvThread:
com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection@269522111[qmid=test_2014-01-
24_15.55.24, fap=10, channel=MY.SVRCONN, ccsid=850, sharecnv=10, hbint=300, peer=/9.20.124.119(
1414),localport=65243,ssl=no,hConns=0,LastDataSend=1391436871409 (Oms ago
),LastDataRecv=1391436871409 (Oms ago),]
Priority: 5
ThreadGroup: java.lang.ThreadGroup[name=JMSCCThreadPool,maxpri=10]
ID: 84
State : RUNNABLE
Stack:
java.net.SocketInputStream.socketRead0(SocketInputStream.java:-2)
java.net.SocketInputStream.read(SocketInputStream.java:163)
java.net.SocketInputStream.read(SocketInputStream.java:133)
com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.receive(RemoteTCPConnection.java:1545)
com.ibm.mq.jmqi.remote.impl.RemoteRcvThread.receiveBuffer(RemoteRcvThread.java:794)
com.ibm.mq.jmqi.remote.impl.RemoteRcvThread.receiveOneTSH(RemoteRcvThread.java:757)
com.ibm.mq.jmqi.remote.impl.RemoteRcvThread.run(RemoteRcvThread.java:150)
com.ibm.msg.client.commonservices.workqueue.WorkQueueItem.runTask(WorkQueueItem.java:214)
\verb|com.ibm.msg.client.commonservices.workqueue.Simple Work Queue Item.run Item (Simple Work Queue Item.run Item)| | The commonservices and the common services are considered as a common services and the common services are considered as a common service are considered as a common 
m.java:105)
com.ibm.msg.client.commonservices.workqueue.WorkQueueItem.run(WorkQueueItem.java:229)
com.ibm.msg.client.commonservices.workqueue.WorkQueueManager.runWorkQueueItem(WorkQueueManager.j
ava:303)
:
```

The information in the header, Data, and Stack Trace sections of the FFST record 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.

Suppressing FFST records

An FFST file that is generated by the IBM MQ classes for JMS contain one FFST record. If a problem occurs multiple times during the execution of an IBM MQ classes for JMS application, multiple FFST files with the same probe identifier are generated. This might not be desirable. The property com.ibm.msg.client.commonservices.ffst.suppress can be used to suppress the production of FFST files. This property must be set in the IBM MQ classes for JMS configuration file used by the application, and can take the following values:

- 0: Output all FFDC files (default).
- -1: Output only the first FFST file for a probe identifier.

integer: Suppress all FFST files for a probe identifier except those files that are a multiple of this number.

Windows FFST: IBM MQ for Windows

Describes the name, location, and contents of the First Failure Support Technology (FFST) files for Windows systems.

In IBM MQ for Windows, FFST information is recorded in a file in the C:\Program Files\IBM\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 an IBM MQ internal error.

FFST files are named AMQ nnnnn.mm.FDC, where:

nnnnn

Is the ID of the process reporting the error

mm

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 3 on page 58.

```
WebSphere MQ First Failure Symptom Report
   _____
                          :- Mon January 28 2008 21:59:06 GMT
   Date/Time
   UTC Time/Zone :- 1201539869.892015 0 GMT
Host Name :- 99VXY09 (Windows 7 Build 2600: Service Pack 1)
   Host Name
PIDS
                        :- 5724H7200
   LVLS
                          :- 7.0.0.0
   Product Long Name :- IBM MQ for Windows
   Vendor :- IBM
Probe Id :- HL010004
   Application Name :- MQM
Component :- hlgReserveLogSpace
SCCS Info :- lib/logger/amqhlge0.c, 1.26
   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\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
                    :- MSGAMQ6709
   Probe Type
   Probe Severity
   Probe Description :- AMQ6709: The log for the Queue manager is full.
 | FDCSequenceNumber :- 0
 MQM Function Stack
 zľaMainThread
 zlaProcessMessage
 zlaProcessMQIRequest
 zlaMQPUT
 zsqMQPUT
 kpiMOPUT
 kqiPutIt
 kqiPutMsgSegments
 apiPutMessage
 agmPutMessage
 aqhPutMessage
 adqWriteMsg
 aqqWriteMsgData
 aqlReservePutSpace
 almReserveSpace
 hlgReserveLogSpace
 xcsFFST
 MQM Trace History
 -----} xllLongLockRequest rc=OK
Figure 3. Sample IBM 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\MQ\errors directory. A dump file will have the same name as the FFST file, in the form AMOnnnnn.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 IBM MQ MQI client installation directory.

These are normally severe, unrecoverable errors and indicate either a configuration problem with the system or an IBM 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: IBM 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 IBM 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 an IBM MQ internal error.

FFST files are named AMQ nnnnn.mm.FDC, where:

nnnnn

Is the ID of the process reporting the error

mm

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 4 on page 60.

```
WebSphere MQ First Failure Symptom Report
   _____
   Date/Time :- Mon January 28 2008 21:59:06 GMT UTC Time/Zone :- 1201539869.892015 0 GMT
                       :- mqperfh2 (HP-UX B.11.23)
   Host Name
   PIDS
                         :- 5724H7202
   LVLS
                         :- 7.0.0.0
   Product Long Name :- IBM MQ for HP-UX Vendor :- IBM Probe Id :- XC034255
   Application Name :- MQM :- xcsWaitEventSem :- 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
                        :- 15497
   Process
   Thread
                        :- CSIM
   QueueManager
    ConnId(2) QM
                         :- 4
                        :- OK
   Major Errorcode
                       :- OK
   Minor Errorcode
                         :- INCORROUT
   Probe Type
   Probe Severity
                         :- 4
   Probe Description :- AMQ6109: An internal IBM MQ error has occurred.
   FDCSequenceNumber :- 0
 MQM Function Stack
 amqzmuc0
 xcsWaitEventSem
 xcsFFST
 MQM Trace History
 Data: 0x00003c87
 --} xcsCheckProcess rc=0K
     xcsRequestMutexSem
 --} xcsRequestMutexSem rc=OK
Figure 4. FFST report for IBM 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 IBM 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 an IBM 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: IBM MQ for IBM i

Describes the name, location, and contents of the First Failure Support Technology (FFST) files for IBM i systems.

For IBM i, FFST information is recorded in a stream file in the /QIBM/UserData/mgm/errors directory.

These errors are normally severe, unrecoverable errors, and indicate either a configuration problem with the system or an IBM MQ internal error.

The stream files are named AMQ nnnnn.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.

```
IBM MQ First Failure Symptom Report
  _____
  Date/Time
                          :- Mon January 28 2008 21:59:06 GMT
  UTC Time/Zone :- Mon January 28 2008 21:5

UTC Time/Zone :- 1201539869.892015 0 GMT

Host Name :- WINAS12B.HURSLEY.IBM.COM
PIDS :- 5733A38

LVIS :- 520
                         :- 520
  LVLS
  Product Long Name :- IBM MQ for IBMi
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 (QMQM) (QMQM/STRMQM_R)
  Process
                          :- 00001689
  Thread
                          :- 00000001
  QueueManager :- TEST.AS400.0E.P
Major Errorcode :- STOP
  Minor Errorcode :- OK
  Probe Type
                          :- HALT6109
  Probe Severity
                          :- 1
  Probe Description :- 0
                  :- 1 1
:- 00d
  Arith1
                           :- 00d0
| Comment1
MQM Function Stack
lpiSPIMQConnect
zstMQConnect
ziiMOCONN
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=0K
xihGetConnSPDetailsFromList
<-- xihGetConnSPDetailsFromList rc=OK
-->
xstCreateConnExtentList
<-- xstCreateConnExtentList rc=OK
-->
xstConnSetInSPbyHandle
-->
xstSerialiseSPList
xllSpinLockRequest
<-- xllSpinLockRequest rc=0K
<-- xstSerialiseSPList rc=0K
xstGetSetDetailsFromSPByHandle
<-- xstGetSetDetailsFromSPByHandle rc=OK
xstConnSharedMemSet
xstConnectExtent
-->
xstAddConnExtentToList
<-- xstAddConnExtentToList rc=0K
<-- xstConnectExtent rc=0K
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 000000000 XIHP...\..r...
SPP:0000 :1aefSTRMQM_R MAYFCT
                                                 020100 :8bbb0:1:6d
SPP:0000 :1aefSTRMOM R MAYFCT
                                                 80000000 000000000 EC161F7C FC002DB0
                                                                                     020100 :8bbc0:2:6d
SPP:0000 :1aefSTRMQM_R MAYFCT
                                                 80000000 000000000 EC161F7C FC002DB0
                              020100 :8bbd0:3:6d
                                                                                     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 000000000
                                                                                    XIHT.....
SPP:0000 :1aefSTRMOM_R MAYFCT
                              020100 :1dc0:21:6d
                                                 00000001 00000000 00000000 00000000
                                                                                     SPP:0000 :1aefSTRMQM_R MAYFCT
                              020100 :1dd0:22:6d
                                                 80000000 00000000 DD13C17B 81001000
SPP:0000 :1aefSTRMQM_R MAYFCT
                              020100 :1de0:23:6d
                                                 00000000 00000046 00000002 00000001
                                                                                     . . . . . . . . . . . . . . . .
                                                 SPP:0000 :1aefSTRMQM_R MAYFCT
                              020100 :1df0:24:6d
                                                                                     . . . . . . . . . . . . . . . .
SPP:0000 :1aefSTRMQM_R MAYFCT 020100 :2064:128:6d 000000000
```

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.

Using trace

You can use different types of trace to help you with problem determination and troubleshooting.

About this task

Use this information to find out about the different types of trace, and how to run trace for your platform.

- Windows "Using trace on Windows" on page 63
- Linux "Using trace on UNIX and Linux systems" on page 65
- "Using trace with IBM MQ server on IBM i" on page 68
- Using trace with IBM MQ client on IBM i" on page 71
- Using trace for problem determination on z/OS" on page 73
- "Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86
- "Tracing IBM MQ classes for JMS applications" on page 87
- "Tracing IBM MQ classes for Java applications" on page 91
- "Tracing the IBM MQ resource adapter" on page 95
- "Tracing additional IBM MQ Java components" on page 97
- "Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS" on page 100

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Troubleshooting overview" on page 7

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?"

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

"First Failure Support Technology (FFST)" on page 51

First Failure Support Technology (FFST) for IBM MQ provides information about events that, in the case of an error, can help IBM support personnel to diagnose the problem.

Related tasks

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

Windows Using trace on Windows

Use the **strmqtrc** and **endmqtrc** commands or the IBM 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 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 <u>strmqtrc</u>). To stop tracing, use the **endmqtrc** control command (see endmqtrc).

In IBM MQ for Windows systems, you can also start and stop tracing using the IBM MQ Explorer, as follows:

- 1. Start the IBM MQ Explorer from the Start menu.
- 2. In the Navigator View, right-click the **IBM 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.

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 strmgtrc.

Related concepts

"Using trace on UNIX and Linux systems" on page 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace with IBM MQ server on IBM i" on page 68

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

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 IBM MQ MQI client trace facility:

strmatrc

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 67.

Trace files on IBM 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 MQ for UNIX and Linux systems, you enable or modify tracing using the **strmqtrc** control command (see <u>strmqtrc</u>). To stop tracing, you use the **endmqtrc** control command (see <u>endmqtrc</u>). On IBM MQ for Linux (x86 and x86-64 platforms) systems, you can alternatively use the IBM 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.

Selective component tracing on IBM 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 Transport Layer Security (TLS) channel security, use:

```
strmqtrc -m QM1 -t ssl
```

For detailed information about the trace command, see strmqtrc.

Selective component tracing on IBM 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 enabled.

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 IBM 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 strmgtrc.

Related concepts

"Using trace with IBM MQ server on IBM i" on page 68

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

Tracing with the AIX system trace

In addition to the IBM MQ trace, IBM MQ for AIX users can use the standard AIX system trace.

Note: You should use the *aix* option, only when directed to do so by IBM service personnel.

AIX system tracing is a three-step process:

- 1. Set the **-o** parameter on the strmqtrc command to aix.
- 2. Gather the data, and run the endmqtrc command once you have done so.
- 3. Format the results.

IBM MQ uses two trace hook identifiers:

X,30D,

This event is recorded by IBM MQ on entry to or exit from a subroutine.

X'30E'

This event is recorded by IBM 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 IBM MQ is installed.

report.file is the name of the file where you want to put the formatted trace output.

Note: All IBM MQ activity on the machine is traced while the trace is active.

Using trace with IBM 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 /OIBM/UserData/mgm/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 MQ trace output in the directory specified by the TRCDIR parameter. This value is the default value.

*MQMFMT

Collect formatted IBM MQ trace output in the directory specified by the TRCDIR parameter.

*PEX

Collect Performance Explorer (PEX) trace output

*ALL

Collect both IBM 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 MO 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.

*LOMFLOW

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 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/mgm/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 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

Using trace with IBM MQ client on IBM i

On IBM i, there is no Control Language (CL) command to capture the trace when using a stand-alone IBM 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 <u>Table 1</u>. 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 5.	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.
comme ntary	Output data for trace points associated with comments in the IBM 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 s	Output data for trace points associated with processing flow in common services.
detail	Activate tracing at high-detail level for flow processing trace points.
lqmdat a	Output data for trace points associated with internal data buffers in the local queue manager.
lqmflo ws	Output data for trace points associated with processing flow in the local queue manager.
otherd ata	Output data for trace points associated with internal data buffers in other components.
otherfl ows	Output data for trace points associated with processing flow in other components.

Table 5. TraceType values (continued)		
Value	Description	
parms	Activate tracing at default-detail level for flow processing trace points.	
remote data	Output data for trace points associated with internal data buffers in the communications component.	
remote flows	Output data for trace points associated with processing flow in the communications component.	
service data	Output data for trace points associated with internal data buffers in the service component.	
service flows	Output data for trace points associated with processing flow in the service component.	
version data	Output data for trace points associated with the version of IBM 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 (AMQpppppp.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 AMQpppppp.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

• 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 endmqtrc command

• To display a formatted trace file:

```
/QSYS.LIB/QMQM.LIB/DSPMQTRC.pgm
```

To examine First Failure Support Technology (FFST) files, see "FFST: IBM MQ for IBM i" on page 61.

Related concepts

"Using trace on UNIX and Linux systems" on page 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

Using trace for problem determination on z/OS

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

The trace facilities available with IBM 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 IBM MQ.

- Controlling the GTF for your z/OS system
- · Controlling the IBM MQ trace for each queue manager subsystem for which you want to collect data
- "Formatting and identifying the control block information on z/OS" on page 76
- "Interpreting the trace information on z/OS" on page 77

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 z/OS" on page 79.

Related concepts

"Using trace on UNIX and Linux systems" on page 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace with IBM MQ server on IBM i" on page 68

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

"Tracing TLS: runmgakm, strmgikm, and runmgckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

The MQI call and user parameter, and GTF on z/OS

Use this topic to understand how to control the z/OS generalized trace facility (GTF) and IBM MQ trace.

You can obtain information about MQI calls and user parameters passed by some IBM MQ calls on entry to, and exit from, IBM MQ. To do this, use the global trace in conjunction with the z/OS generalized trace facility (GTF).

Starting and stopping the GTF

On z/OS, you can use the generalized trace facility (GTF) to record and diagnose system and program problems.

About this task

You can obtain information about MQI calls and user parameters passed by some IBM MQ calls on entry to, and exit from, IBM MQ. To do this, use the global trace in conjunction with the z/OS generalized trace facility (GTF).

Procedure

• Start the GTF at the console by entering a **START GTF** command.

When you start the GTF, specify the USRP option. You are prompted to enter a list of event identifiers (EIDs). The EIDs used by IBM MQ are:

5**E**9

To collect information about control blocks on entry to IBM MQ

5EA

To collect information about control blocks on exit from IBM MO

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 IBM 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. The following example shows a sample startup for the GTF, specifying the four EIDs, and a jobname. The lines shown in **bold** 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 Starting GTF.

```
START GTFxx.yy
 #HASP100 GTFxx.yy ON STCINRDR
#HASP373 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
 IEE600I REPLY TO 04 IS;U
 AHLO31I GTF INITIALIZATION COMPLETE
```

where

- xx is the name of the GTF procedure to use (optional)
- yy is an identifier for this occurrence of GTF trace
- xxxx is the name of the queue manager
- zzzzzzzz is a batch job or CICS region name

Up to 5 job names can be listed.

When using GTF, specify the primary job name (CHINIT, CICS, or batch) in addition to the queue manager name (xxxxMSTR).

Stop the GTF at the console.

When you enter the stop command for the GTF, include the additional identifier (yy) that you used at startup, as shown in the following example:

```
STOP yy
```

Related information

Generating IBM MQ GTF trace on IBM z/OS

Controlling the trace within IBM MQ for z/OS

IBM 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 IBM 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 MQSC commands.

Formatting and identifying the control block information on z/OS

After capturing a trace, the output must be formatted and the IBM MQ control blocks identified.

- Formatting the information
- Identifying the control blocks associated with IBM MQ
- Identifying the event identifier associated with the control block

Formatting the information

To format the user parameter data that is collected by the global trace, use either the batch job that is shown in Figure 5 on page 76 or the IPCS GTFTRACE USR (xxx) command, where xxx is:

5E9

To format information about control blocks on entry to IBM MQ MQI calls.

5EA

To format information about control blocks on exit from IBM MQ MQI calls.

5EE

To format information about IBM 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 5. Formatting the GTF output in batch
```

Identifying the control blocks associated with IBM MQ

The format identifier for the IBM 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 start of each control block. These messages indicate whether the data was captured on entry to or exit from IBM MQ:

- CSQW072I ENTRY: MQ user parameter trace
- CSQW073I EXIT: MQ user parameter trace

Related tasks

"Starting and stopping the GTF" on page 74

On z/OS, you can use the generalized trace facility (GTF) to record and diagnose system and program problems.

Interpreting the trace information on z/OS

The GTFTRACE produced by IBM MQ can be examined to determine possible errors with invalid addresses, invalid control blocks, and invalid data.

Start the GTFTRACE subcommand to format generalized trace facility (GTF) records contained in a dump or in a trace data set. For more information on GTF, see "Starting and stopping the GTF" on page 74.

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 occurred on exit from IBM MQ, it is possible that IBM 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 6 on page 77 illustrates which control blocks are traced for different MQI calls.

Table 6. Control blocks traced for IBM 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 SelectorCount is greater than 0)	Selectors (if SelectorCount is greater than 0)	
		Integer attributes (if IntAttrCount is greater than 0)	
		Character attributes (if CharAttrLength is greater than 0)	

Table 6. Control blocks traced for IBM MQ MQI calls (continued)			
MQI call	Entry	Exit	
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 SelectorCount is greater than 0)	Selectors (if SelectorCount is greater than 0)	
	Integer attributes (if IntAttrCount is greater than 0)	Integer attributes (if IntAttrCount is greater than 0)	
	Character attributes (if CharAttrLength is greater than 0)	Character attributes (if CharAttrLength 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.SubUser 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 appears 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	
Action	Requested action	
BufferL	Buffer length	
CBD	Address of callback descriptor	
CompCode	Completion code	
CharAttL	Character attributes length	
DataL	Data length	
Hobj	Object handle	
Hsub	Subscription handle	
IntAttC	Count of integer attributes	
pObjDesc	Object descriptor	
Oper	Operation	
Options	Options	
pBuffer	Address of buffer	

Identifier	Description	
pCharAtt	Address of character attributes	
pCTLO	Address of control callback options	
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	
pSD	Address of subscription descriptor	
pSelect	Address of selectors	
pSRQOpt	Address of subscription request options	
pSTS	Address of status structure	
Reason	Reason code	
RSVn	Reserved for IBM	
SelectC	Selector count	
Thread	Thread	
Туре	Requested type	
UOWInfo	Information about the unit of work	
Userid	CICS or IMS user ID, for batch or TSO this is zero	

Other types of trace on z/OS

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 IBM MQ.

- · The channel initiator trace
- The line trace
- The CICS adapter trace
- System SSL trace
- z/OS traces

The channel initiator trace

See Figure 10 on page 119 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.

You can run the trace using the <u>START TRACE</u> command. You can also set this trace to start automatically using the TRAXSTR queue manager attribute. For more information about how to do this, see <u>ALTER</u> QMGR.

You can display this trace information by entering the IPCS command:

```
LIST 1000. DSPNAME(CSQXTRDS)
```

You can format it using the command:

```
CTRACE COMP(CSQXssnm)
```

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 <u>"Snap dumps on z/OS" on page 135</u> 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.

i) 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
```

ii) To stop the trace, issue the following command:

```
+ssid STOP TRACE(CHINIT) TNO(01)
```

- iii) You can then enter another START TRACE command with a DISPLAY CHSTATUS SAVED command to gather more line trace data to CSQSNAP.
- 4. 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 APO 000, 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" on page 81.

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 IBM MQ. For information about using and interpreting this trace facility, see the <u>z/OS MVS</u> 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 APO xxx (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 7. CICS adapter trace entries				
Name	Description	Trace sequence	Trace data	
CSQCABNT	Abnormal termination	Before issuing END_THREAD ABNORMAL to IBM 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.	
CSQCBACK	Syncpoint backout	Before issuing BACKOUT to IBM MQ. This is due to an explicit backout request from the application.	Unit of work information.	
CSQCCONX	MQCONNX	Before issuing MQCONNX to IBM MQ.	Connection tag.	

Name Description		Trace sequence	Trace data	
CSQCCCRC	Completion code and reason code	After unsuccessful return from API call.	Completion code and reason code.	
CSQCCOMM	Syncpoint commit	Before issuing COMMIT to IBM 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 a 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 IBM 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 IBM 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 gives the number of in-doubt of work.		
CSQCINDW	Syncpoint in doubt	During syncpoint processing, CICS is in doubt as to the disposition of the unit of work.		
CSQCINQH	INQ handle	Before issuing MQINQ to IBM MQ.	Object handle.	
CSQCLOSH	CLOSE handle	Before issuing MQCLOSE to IBM MQ.	Object handle.	
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.	

Name	Description	Trace sequence	Trace data	
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 IBM 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 IBM MQ.	Object name.	
CSQCPMGD	PUT message data	Before issuing MQPUT to IBM MQ.	Up to 40 bytes of the message data.	
CSQCPMGH	PUT message handle	Before issuing MQPUT to IBM MQ.	Object handle.	
CSQCPMGI	PUT message ID	After successful MQPUT from IBM MQ.	Message ID and correlation ID of the message.	
CSQCPREP	Syncpoint prepare	Before issuing PREPARE to IBM 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 IBM MQ. Up to 40 bytes of data of message.		
CSQCP1MI	PUTONE message ID	After successful return from Message ID and correlative MQPUT1. Message ID and correlative message.		
CSQCP1ON	PUTONE object name	Before issuing MQPUT1 to IBM MQ. Object name.		
CSQCRBAK	Resolved backout	Before issuing Unit of work information RESOLVE_ROLLBACK to IBM MQ.		
CSQCRCMT	Resolved commit	Before issuing RESOLVE_COMMIT Unit of work information. to IBM MQ.		
CSQCRMIR	RMI response	Before returning to the CICS RMI (resource manager interface) from a specific invocation. Architected RMI responsementing depends of the invocation. To determine of invocation, look at preentries produced by the 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.	

Table 7. CICS adapter trace entries (continued)			
Name	Description Trace sequence Trace data		Trace data
CSQCSETH	SET handle	Before issuing MQSET to IBM 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.

Enabling internal trace for the AMSM system

Trace for the AMSM address space can be enabled using the _AMS_MSG_LEVEL variable, which is passed into the AMSM address space through the ENVARS DD card.

A sample data set for the ENVARS DD card is in thlqual.SCSQPROC(CSQ40ENV).

Trace is written to the SYSOUT of the AMSM address space.

The _AMS_MSG_LEVEL variable specifies the subcomponent and message level that is to be logged. An asterisk indicates all subcomponents to be logged; currently there is only one subcomponent.

The severity levels are:

- S severe messages only
- E error and severe messages only
- W warning, error, and severe messages only
- I informational, warning, error, and severe messages. This is the default value
- D debug mode, all messages with additional debug diagnostics
- V verbose mode, all of the preceding, plus buffer dumps



Attention: You should only enable debug or verbose mode on the advice of an IBM service representative.

For example, to enable the default for _AMS_MSG_LEVEL, issue the following:

```
_AMS_MSG_LEVEL=*.i
```

To enable verbose mode, issue the following:

```
_AMS_MSG_LEVEL=*.v
```

Tracing the Advanced Message Queuing Protocol (AMQP) Service

The trace facility provided by the Advanced Message Queuing Protocol (AMQP) Service is provided to help IBM Support to diagnose customer issues related to the service.

About this task

There are two ways to control trace for the IBM MQ AMQP service:

- By using the strmqtrc and endmqtrc commands, to start and stop trace. Enabling trace, using the strmqtrc command, generates trace information for the entire queue manager where the IBM MQ AMQP service is running. This includes the IBM MQ AMQP service itself, and the underlying Java Message Queuing Interface (JMQI) that the service uses to communicate with other queue manager components.
- By running the **controlAMQPChannel** command. Note, that turning trace on using the **controlAMQPChannel** command traces only the IBM MQ AMQP service.

If you are unsure which option to use, contact your IBM Support representative and they will be able to advise you on the best way to collect trace for the issue that you are seeing.

Procedure

- 1. Method one
 - a) Bring up a command prompt and navigate to the directory: MQ_INSTALLATION_PATH\bin
 - b) Run the **strmqtrc** command to enable trace:

```
strmqtrc -m qmgr_name
```

where qmgr_name is the name of the queue manager where the IBM MQ AMQP service is running.

- c) Reproduce the issue.
- d) Stop trace, by running the command:

```
endmqtrc -m qmgr_name
```

- 2. Method two.
 - a) Bring up a command prompt and navigate to the directory:

```
MQ_INSTALLATION_PATH\bin
```

- b) Run the following command to enable trace:
 - . Windows

```
controlAMQPChannel -qmgr=qmgr_name -mode=starttrace
```

Linux UNIX

```
./controlAMQPChannel.sh -qmgr=qmgr_name -mode=starttrace
```

where qmgr_name is the name of the queue manager where the AMQP Service is running.

c) Reproduce the issue.

Windows

- d) When the issue occurs, stop trace by running the following command:
 - controlAMOPChannel -amou

```
\verb|controlAMQPC| + qmgr=qmgr\_name| - mode=stoptrace|
```

Linux VINIX

```
./{\tt controlAMQPChannel.sh\ -qmgr=\it qmgr\_\it name\ -mode=stoptrace\ [clientid=ClientIdentifier]}
```

where qmgr_name is the name of the queue manager where the AMQP Service is running.

Results

To view the trace output, go to the following directory:

- Windows MQ_DATA_PATH\trace.
- Linux UNIX /var/mqm/trace.

The trace files containing the information from the AMQP Service are called amqp_N.trc, where N is a number.

Trace information generated by the JMQI is written to a trace file called amqp_PPPPP.trc, where PPPPP is the process identifier for the AMQP Service.

Windows Linux AIX Additional diagnostics using the

controlAMQPChannel command

Using the **controlAMQPChannel** command to provide additional diagnostic information about the AMQP service.

Procedure

Run the following command to provide useful diagnostic information from the MQXR service:

```
< MQ\_INSTALLATION\_PATH>\\ amqp\\bin\\controlAMQPChannel -qmgr=< QMGR\_NAME> -mode= diagnostics -diagnostics type=< number>
```

The diagnostic information generated depends on the value of the **-diagnosticstype**=<*number*> parameter:

-diagnosticstype= 0

Thread dump written to the console

-diagnosticstype= 1

FDC with some internal service statistics

-diagnosticstype= 2

FDC with internal statistics, plus information about the clients that are currently connected

-diagnosticstype= 3

Heap dump

-diagnosticstype= 4

Javacore

-diagnosticstype= 5

Full system dump

-diagnosticstype= 6

Detailed information about a specific client. Note that you must also supply the **-clientid** parameter for that client as well.

Tracing TLS: runmqakm, strmqikm, and runmqckm functions

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

strmqikm and runmqckm trace

To request **strmqikm** tracing, run the **strmqikm** command for your platform with the following -D flags.

On UNIX, Linux, and Windows:

```
strmqikm -Dkeyman.debug=true -Dkeyman.jnitracing=ON
```

To request **runmqckm** tracing, run the **runmqckm** command for your platform with the following -D flags.

On UNIX, Linux, and Windows:

```
runmqckm -Dkeyman.debug=true -Dkeyman.jnitracing=ON
```

strmqikm and **runmqckm** write three trace files to the directory from which you start them, so consider starting iKeyman or **runmqckm** from the trace directory to which the runtime TLS 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 IBM MQ is installed.

The trace file generated by **strmqikm** and **runmqckm** has the following format:

```
debugTrace. n
```

where n is an incrementing number starting at 0.

runmqakm trace

To request **runmqakm** tracing, run the **runmqakm** command with the following flags:

```
runmqakm -trace filename
```

where *filename* is the name of the trace file to create. You cannot format the **runmqakm** trace file. Send it unchanged to IBM support. The **runmqakm** trace file is a binary file and, if it is transferred to IBM support via FTP, it must be transferred in binary transfer mode.

Runtime TLS trace

On UNIX, Linux, and Windows systems, you can independently request trace information for **strmqikm**, **runmqckm**, the runtime TLS functions, or a combination of these.

The runtime TLS trace files have the names AMQ.TLS.TRC and AMQ.TLS.TRC.1 and the TLS trace files have the names AMQ.SSL.TRC and AMQ.SSL.TRC.1. You cannot format any of the TLS trace files; send them unchanged to IBM support. The TLS 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 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace with IBM MQ server on IBM i" on page 68

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

Tracing IBM MQ classes for JMS applications

The trace facility in IBM MQ classes for JMS is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

If you are asked to provide trace output to investigate an issue, use one of the options mentioned below:

- If the issue is easy to recreate, then collect an IBM MQ classes for JMS trace by using a Java System Property. For more information, see "Collecting an IBM MQ classes for JMS trace by using a Java system property" on page 88.
- If an application needs to run for a period of time before the issue occurs, collect an IBM MQ classes for JMS trace by using the IBM MQ classes for JMS configuration file. For more information, see "Collecting an IBM MQ classes for JMS trace by using the IBM MQ classes for JMS configuration file" on page 89.
- To generate a trace from an application that is currently running, collect the IBM MQ classes for JMS trace dynamically by using the traceControl utility. For more information, see "Collecting an IBM MQ classes for JMS trace dynamically by using the traceControl utility" on page 90.

If you are unsure which option to use, contact your IBM Support representative and they will be able to advise you on the best way to collect trace for the issue that you are seeing.

If a severe or unrecoverable error occurs, First Failure Support Technology (FFST) information is recorded in a file with a name of the format JMSCC xxxx. FDC where xxxx is a four-digit number. This number 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 IBM MQ classes for JMS, see <u>"FFST: IBM MQ classes for JMS" on page 52.</u>

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 that closes all handlers and sets the log level to null, which in effect turns 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 shown 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();
    7
        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();
    3
```

The shutdown hook is necessary to ensure that trace is properly shut down 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 ...
```

Collecting an IBM MQ classes for JMS trace by using a Java system property

For issues that can be reproduced in a short amount of time, IBM MQ classes for JMS trace should be collected by setting a Java system property when starting the application.

About this task

To collect a trace by using a Java system property, complete the following steps.

Procedure

• Run the application that is going to be traced by using the following command:

```
java -Dcom.ibm.msg.client.commonservices.trace.status=ON application_name
```

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in:

- For IBM MQ classes for JMS for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mgjavaclient %PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

The application stops writing information to the trace file when it is stopped.

If the application has to run for a long period of time before the issue that the trace is being collected for occurs, then the trace file could potentially be very large. In this situation, consider collecting trace by using the IBM MQ classes for JMS configuration file (see "Collecting an IBM MQ classes for JMS trace by using the IBM MQ classes for JMS configuration file" on page 89). When enabling trace in this way, it is possible to control the amount of trace data that the IBM MQ classes for JMS generate.

Collecting an IBM MQ classes for JMS trace by using the IBM MQ classes for JMS configuration file

If an application must run for a long period of time before an issue occurs, IBM MQ classes for JMS trace should be collected by using the IBM MQ classes for JMS configuration file. The configuration file allows you to specify various options to control the amount of trace data that is collected.

About this task

To collect a trace by using the IBM MQ classes for JMS configuration file, complete the following steps.

Procedure

- Create an IBM MQ classes for JMS configuration file.
 For more information about this file, see The IBM MQ classes for JMS configuration file.
- 2. Edit the IBM MQ classes for JMS configuration file so that the property **com.ibm.msg.client.commonservices.trace.status** is set to the value ON.
- 3. Optional: Edit the other properties that are listed in the IBM MQ classes for JMS configuration file Java Standard Edition Trace Settings.
- 4. Run the IBM MO classes for JMS application by using the following command:

```
java -Dcom.ibm.msg.client.config.location=config_file_url
application_name
```

where *config_file_url* is a uniform resource locator (URL) that specifies the name and location of the IBM MQ classes for JMS configuration file. URLs of the following types are supported: http, file, ftp, and jar.

Here is an example of a Java command:

```
java -Dcom.ibm.msg.client.config.location=file:/D:/mydir/myjms.config
MyAppClass
```

This command identifies the IBM MQ classes for JMS configuration file as the file D:\mydir\myjms.config on the local Windows system.

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in:

- For IBM MQ classes for JMS for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 5.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

To change the name of the trace file, and the location where it is written, ensure that the IBM MQ classes for JMS configuration file that the application uses contains an entry for the property com.ibm.msg.client.commonservices.trace.outputName. The value for the property can be either of the following:

- The name of the trace file that is created in the application's working directory.
- The fully qualified name of the trace file, including the directory in which the file is created.

For example, to configure the IBM MQ classes for JMS to write trace information for an application to a file called C:\Trace\trace.trc, the IBM MQ classes for JMS configuration file that the application uses needs to contain the following entry:

com.ibm.msg.client.commonservices.trace.outputName=C:\Trace\trace.trc

Collecting an IBM MQ classes for JMS trace dynamically by using the traceControl utility

The traceControl utility that is shipped with the IBM MQ classes for JMS allows trace to be collected from a running application. This can be very useful if IBM Support need to see a trace from an application once an issue has occurred, or if trace needs to be collected from a critical application that cannot be stopped.

About this task

Important: This function is supported for IBM Java runtime environments (JREs) only.

For more information about the traceControl utility, see "Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS" on page 100.

To collect a trace by using the traceControl utility, complete the following steps.

Procedure

- 1. Bring up a command prompt, and navigate to the directory MQ_INSTALLATION_PATH\java\lib.
- 2. Run the command:

```
java -jar com.ibm.mq.traceControl.jar -list
```

This command brings up a list of all of the Java processes on the system.

3. Identify the process identifier for the IBM MQ classes for JMS application that needs to be traced, and run the command:

```
java -jar com.ibm.mq.traceControl.jar -i processidentifier -enable
```

Trace is now turned on for the application.

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in:

- For IBM MQ classes for JMS for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 5.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the JAR file com.ibm.mqjms.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for JMS from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

4. To turn trace off, run the command:

```
java -jar com.ibm.mq.traceControl.jar -i processidentifier -disable
```

Tracing IBM MQ classes for Java applications

The trace facility in IBM MQ classes for Java is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

About this task

If you are asked to provide trace output to investigate an issue, use one of the options mentioned below:

- If the issue is easy to recreate, then collect an IBM MQ classes for Java trace by using a Java System Property. For more information, see "Collecting an IBM MQ classes for Java trace by using a Java system property" on page 92.
- If an application needs to run for a period of time before the issue occurs, collect an IBM MQ classes for Java trace by using the IBM MQ classes for Java configuration file. For more information, see "Collecting an IBM MQ classes for Java trace by using the IBM MQ classes for Java configuration file" on page 93.
- To generate a trace from an application that is currently running, collect the IBM MQ classes for Java trace dynamically by using the traceControl utility. For more information, see "Collecting an IBM MQ classes for Java trace dynamically by using the traceControl utility" on page 94.

If you are unsure which option to use, contact your IBM Support representative and they will be able to advise you on the best way to collect trace for the issue you are seeing.

If a severe or unrecoverable error occurs, First Failure Support Technology (FFST) information is recorded in a file with a name of the format JAVACC xxxx. FDC where 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.

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, which in effect turns 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();
                 3
         3
                 public JmsLogManager(){
                  // add shutdown hook to ensure final cleanup
                 Runtime.getRuntime().addShutdownHook(new ShutdownHook());
         3
                 public void reset() throws SecurityException {
                  // does nothing
         public void doReset(){
                 super.reset();
```

The shutdown hook is necessary to ensure that trace is properly shut down 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 ...
```

Collecting an IBM MQ classes for Java trace by using a Java system property

For issues that can be reproduced in a short amount of time, IBM MQ classes for Java trace should be collected by setting a Java system property when starting the application.

About this task

To collect a trace by using a Java system property, complete the following steps.

Procedure

Run the application that is going to be traced by using the following command:

```
java -Dcom.ibm.msg.client.commonservices.trace.status=ON application_name
```

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in: :

- For IBM MQ classes for Java for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mg.jar, trace is written to a file called mgjava %PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mq.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mg.allclient.jar, trace is written to a file called mqjavaclient_%PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

The application stops writing information to the trace file when it is stopped.

If the application has to run for a long period of time before the issue that the trace is being collected for occurs, then the trace file could potentially be very large. In this situation, consider collecting trace by using the IBM MQ classes for Java configuration file (see "Collecting an IBM MQ classes for Java trace by using the IBM MO classes for Java configuration file" on page 93). When enabling trace in this way, it is possible to control the amount of trace data that the IBM MQ classes for Java generate.

Collecting an IBM MQ classes for Java trace by using the IBM MQ classes for Java configuration file

If an application must run for a long period of time before an issue occurs, IBM MQ classes for Java trace should be collected by using the IBM MQ classes for Java configuration file. The configuration file allows you to specify various options to control the amount of trace data that is collected.

About this task

To collect a trace by using the IBM MO classes for Java configuration file, complete the following steps.

Procedure

- 1. Create an IBM MQ classes for Java configuration file. For more information about this file, see The IBM MQ classes for Java configuration file.
- 2. Edit the IBM MQ classes for Java configuration file so that the property com.ibm.msg.client.commonservices.trace.status is set to the value ON.
- 3. Optional: Edit the other properties that are listed in the IBM MQ classes for Java configuration file Java Standard Edition Trace Settings.
- 4. Run the IBM MQ classes for Java application by using the following command:

```
java -Dcom.ibm.msg.client.config.location=config_file_url
application_name
```

where config file url is a uniform resource locator (URL) that specifies the name and location of the IBM MQ classes for Java configuration file. URLs of the following types are supported: http, file, ftp, and jar.

Here is an example of a Java command:

```
java -Dcom.ibm.msg.client.config.location=file:/D:/mydir/myJava.config
MyAppClass
```

This command identifies the IBM MQ classes for Java configuration file as the file D:\mydir\myJava.config on the local Windows system.

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in:

- For IBM MQ classes for Java for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 5.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mq.jar, trace is written to a file called mqjava_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mq.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

To change the name of the trace file, and the location where it is written, ensure that the IBM MQ classes for Java configuration file that the application uses contains an entry for the property com.ibm.msg.client.commonservices.trace.outputName. The value for the property can be either of the following:

- The name of the trace file that is created in the application's working directory.
- The fully qualified name of the trace file, including the directory in which the file is created.

For example, to configure the IBM MQ classes for Java to write trace information for an application to a file called C:\Trace\trace.trc, the IBM MQ classes for Java configuration file that the application uses needs to contain the following entry:

com.ibm.msg.client.commonservices.trace.outputName=C:\Trace\trace.trc

Collecting an IBM MQ classes for Java trace dynamically by using the traceControl utility

The traceControl utility that is shipped with the IBM MQ classes for Java allows trace to be collected from a running application. This can be very useful if IBM Support need to see a trace from an application once an issue has occurred, or if trace needs to be collected from a critical application that cannot be stopped.

About this task

For more information about the traceControl utility, see "Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS" on page 100.

To collect a trace by using the traceControl utility, complete the following steps.

Procedure

- 1. Bring up a command prompt, and navigate to the directory MQ_INSTALLATION_PATH\java\lib.
- 2. Run the command:

```
java -jar com.ibm.mq.traceControl.jar ...
```

This command brings up a list of all of the Java processes on the system.

3. Identify the process identifier for the IBM MQ classes for Java application that needs to be traced, and run the command:

```
java -jar com.ibm.mq.traceControl -i process identifier -enable
```

Trace is now turned on for the application.

By default, trace information is written to a trace file in the current working directory of the application. The name of the trace file depends upon the environment that the application is running in:

- For IBM MQ classes for Java for IBM MQ 9.0.0 Fix Pack 1 or earlier, trace is written to a file called mqjms_%PID%.trc.
- V 5.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mq.jar, trace is written to a file called mqjava_%PID%.trc.
- V 9.0.0.2 From IBM MQ 9.0.0 Fix Pack 2, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the JAR file com.ibm.mq.jar, trace is written to a file called mqjava_%PID%.cl%u.trc.
- V 9.0.0.10 From IBM MQ 9.0.0 Fix Pack 10, if the application has loaded the IBM MQ classes for Java from the relocatable JAR file com.ibm.mq.allclient.jar, trace is written to a file called mqjavaclient_%PID%.cl%u.trc.

where %PID% is the process identifier of the application that is being traced, and %u is a unique number to differentiate files between threads running trace under different Java classloaders.

4. To turn trace off, run the command:

```
java -jar com.ibm.mq.traceControl -i process identifier -disable
```

Tracing the IBM MQ resource adapter

The ResourceAdapter object encapsulates the global properties of the IBM MQ resource adapter. To enable trace of the IBM MQ resource adapter, properties need to be defined in the ResourceAdapter object.

The ResourceAdapter object has two sets of properties:

- · Properties associated with diagnostic tracing
- · Properties associated with the connection pool managed by the resource adapter

The way you define these properties depends on the administration interfaces provided by your application server.

Table 8 on page 96 lists the properties of the ResourceAdapter object that are associated with diagnostic tracing.

Table 8. Properties of the ResourceAdapter object that are associated with diagnostic tracing			
Name of property	Туре	Default value	Description
traceEnabled	String	false	A flag to enable or disable diagnostic tracing. If the value is false, tracing is turned off.
traceLevel	String	3	The level of detail in a diagnostic trace. The value can be in the range 0, which produces no trace, to 10, which provides the most detail. See Table 9 on page 96 for a description of each level. If trace is enabled, traceLevel should be set to the value 10, unless otherwise specified by IBM Support.
logWriterEnabled	String	true	A flag to enable or disable the sending of a diagnostic trace to a LogWriter object provided by the application server. If the value is true, the trace is sent to a LogWriter object. If the value is false, any LogWriter object provided by the application server is not used.

Table 9 on page 96 describes the levels of detail for diagnostic tracing.

Table 9. T	Table 9. The levels of detail for diagnostic tracing		
Level number	Level of detail		
0	No trace.		
1	The trace contains error messages.		
3	The trace contains error and warning messages.		
6	The trace contains error, warning, and information messages.		
8	The trace contains error, warning, and information messages, and entry and exit information for methods.		
9	The trace contains error, warning, and information messages, entry and exit information for methods, and diagnostic data.		
10	The trace contains all trace information.		

Note: Any level that is not included in this table is equivalent to the next lowest level. For example, specifying a trace level of 4 is equivalent to specifying a trace level of 3. However, the levels that are not included might be used in future releases of the IBM MQ resource adapter, so it is better to avoid using these levels.

If diagnostic tracing is turned off, error and warning messages are written to the system error stream. If diagnostic tracing is turned on, error messages are written to the system error stream and to the trace destination, but warning messages are written only to the trace destination. However, the trace contains warning messages only if the trace level is 3 or higher. By default, the trace destination is the current working directory, but if the logWriterEnabled property is set, the trace is sent to the application server.

In general, the ResourceAdapter object requires no administration. However, to enable diagnostic tracing on UNIX and Linux systems for example, you can set the following properties:

traceEnabled: true
traceLevel: 10

These properties have no effect if the resource adapter has not been started, which is the case, for example, when applications using IBM MQ resources are running only in the client container. In this situation, you can set the properties for diagnostic tracing as Java virtual machine (JVM) system

properties. You can set the properties by using the -D flag on the **java** command, as in the following example:

```
java ... -DtraceEnabled=true -DtraceLevel=10
```

Hints and tips

You do not need to define all the properties of the ResourceAdapter object. Any properties that remain unspecified take their default values.

In a managed environment, it is better not to mix the two ways of specifying properties. If you do mix them, the JVM system properties take precedence over the properties of the ResourceAdapter object.

When using WebSphere Application Server traditional 9.0 with the IBM MQ 9.0 resource adapter, as the Java EE Dependency Injection is now a common Java EE paradigm, the standard trace string should be updated to include com.ibm.ws.cdi.jms*=all. This means that the full string is:

```
*=info:jmsApi=all:Messaging=all:com.ibm.mq.*=all:JMSApi=all:com.ibm.ws.cdi.jms*=all
```

For more information about using trace with WebSphere Application Server traditional, see the technote Enabling Java Message Service (JMS) trace for WebSphere Application Server.

Tracing additional IBM MQ Java components

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM 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 IBM MQ facilities or the facilities of IBM MQ classes for Java or IBM MQ classes for JMS, as appropriate. Generally use the IBM 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 IBM MQ trace.

To request and configure diagnostic output, two system properties are used when starting an IBM 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 98.

For instructions on locating trace information and FFDC files, see <u>"Java trace and FFDC files" on page</u> 99.

Related concepts

"Using trace on UNIX and Linux systems" on page 65

Use the **strmqtrc** and **endmqtrc** commands to start and end tracing, and **dspmqtrc** to display a trace file

"Using trace with IBM MQ server on IBM i" on page 68

Use the TRCMQM command to start and stop tracing and specify the type of trace that you require.

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Tracing TLS: runmqakm, strmqikm, and runmqckm functions" on page 86

How to trace Transport Layer Security (TLS), and request **runmqakm** tracing and **strmqikm** (iKeyman) and **runmqckm** (iKeycmd) tracing.

Related reference

"Using trace on Windows" on page 63

Use the **strmqtrc** and **endmqtrc** commands or the IBM MQ Explorer interface to start and end tracing.

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 IBM MQ.

Note that case is significant in all these entries:

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 IBM MQ Explorer, "soap" refers to the diagnostics from the running process within IBM MQ Transport for SOAP, and "wmqjavaclasses" refers to the diagnostics from the underlying IBM 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 IBM 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= dirnαme

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 <u>Figure 6 on page 99</u>. Lines beginning with the number sign (#) are treated as comments.

```
# Java diagnostics for IBM MQ Transport for SOAP
# and the IBM 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 6. 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" installation option.

Java trace and FFDC files

File name conventions for Java trace and FFDC files.

When Java trace is generated for IBM MQ Transport for SOAP, it is written to a file with a name of the format AMQ. <code>diagid</code>. <code>counter.TRC</code>. Here, <code>diagid</code> is the value of the system property com.ibm.mq.commonservices.diagid associated with this Java process, as described earlier in this section, and <code>counter</code> 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 MO trace.

If com.ibm.mq.commonservices.diagid is not specified, the value of *diagid* is the current time, in the format YYYYMMDDhhmmssmmm.

When Java trace is generated for the IBM MQ Explorer, it is written to file with a name of the format AMQYYYYMMDDHHmmssmmm.TRC.n. Each time IBM MQ Explorer trace is run, the trace facility renames all previous trace files by incrementing the file suffix.n by one. The trace facility then creates a new file with the suffix.0 that is always the latest.

The IBM MQ Java classes trace file has a name based on the equivalent IBM MQ Transport for 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 MQ Explorer or for IBM 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 MQ Explorer and for IBM 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 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.

Controlling trace in a running process by using IBM MQ classes for Java and IBM MQ classes for JMS

The IBM MQ classes for Java and IBM MQ classes for JMS register a Standard MBean that allows suitable Java Management Extensions (JMX) tools to control certain aspects of trace behavior for a client process.

Principles

As an alternative to the well-known general-purpose tools like jconsole you can use a command-line tool in the form of an executable JAR file to access these facilities.

The JAR file is called com.ibm.mq.traceControl.jar and is stored in the java/lib subdirectory of the IBM MQ installation (see What is installed for IBM MQ classes for JMS and Installation directories for IBM MO classes for Java.

Note: Depending on configuration, JMX tools can be used either locally (on the same system as the process) or remotely. The local case is discussed initially.

Finding the process

To control a process, you must establish a JMX connection it. To control a process locally, you must specify its identifier.

To display a summary of running Java processes with their identifiers, run the executable JAR file with the option -list. This option produces a list of identifiers and descriptions for the processes that are found.

Examining trace status

When you have found the identifier for the relevant process, run the executable JAR file with the options <code>-i identifier</code> -status, where <code>identifier</code> is the identifier of the process you want to change. These options display the status, either enabled or <code>disabled</code> for the process, and the information about where the process is running, the name of the trace file, and a tree that represents the inclusion and exclusion of packages in trace.

Enabling and disabling trace

To enable trace for a process, run the executable JAR file with the options -i identifier -enable.

To disable trace for a process, run the executable JAR file with the options -i identifier -disable.

Note: You can choose only one option from the set -status, -enable, and -disable.

Including and excluding packages

To include a package in trace for a process, run the executable JAR file with the options -i identifier -ip package_name, where package_name is the name of your package.

To exclude a package from trace for a process, run the executable JAR file with the options -i identifier -ep package_name.

Note: You can use multiple -ip and -ep options. These options are not checked for consistency.

When you specify a package for exclusion or inclusion, the handling of packages that have matching prefixes is not affected. For example, excluding the package com.ibm.mq.jms from trace would not exclude com.ibm.mq, com.ibm.msq.client.jms, or com.ibm.mq.remote.api, but it would exclude com.ibm.mq.jms.internal.

```
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -list
10008 : 'MQSample'
9004 : 'MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -list'

C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -status
Tracing enabled : false
User Directory : C:\Users\IBM_ADMIN\RTCworkspace\sandpit
```

```
Trace File Name : mqjms.trc
Package Include/Exclude tree
root - Included
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -enable
Enabling trace
Tracing enabled : true
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -status
Tracing enabled : true
User Directory : C:\Users\IBM_ADMIN\RTCworkspace\sandpit
Trace File Name : mgjms 10008.cl0.trc
Package Include/Exclude tree
root - Included
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -ip
com.ibm.mq.jms
Adding 'com.ibm.mq.jms' to the list of packages included in trace
C:>java -jar MQ INSTALLATION PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -status
Tracing enabled: true
User Directory : C:\Users\IBM_ADMIN\RTCworkspace\sandpit
Trace File Name : mqjms_10008.cl0.trc
Package Include/Exclude tree
root - Include com - Included
       Included
ibm - Included
mq - Included
jms - Included
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -ip
com.acme.banana -ep com.acme.banana.split -ip com.acme.banana.shake
Adding 'com.acme.banana' to the list of packages included in trace
Adding 'com.acme.banana' to the list of packages included in trace Adding 'com.acme.banana.shake' to the list of packages included in trace
Adding 'com.acme.banana.split' to the list of packages excluded from trace
C:>java -jar MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.traceControl.jar -i 10008 -status
Tracing enabled : true User Directory : C:\Users\IBM_ADMIN\RTCworkspace\sandpit
Trace File Name : mqjms_10008.cl0.trc
Package Include/Exclude tree
root - Included
com - Included
acme - Included
banana - Included
shake - Included
split - Excluded
ibm - Included
mq - Included
jms - Included
```

The package inclusion-exclusion tree

The tracing mechanism for IBM MQ classes for Java and IBM MQ classes for JMS tracks the inclusion and exclusion of packages by means of a tree structure, starting from a root node. In the tree structure each node represents one element of a package name, identified by the package name element and containing a trace status which can be either Included or Excluded. For example the package *com.ibm.mq* would be represented by three nodes identified by the strings com, ibm, and mq.

Initially, the tree usually contains entries to include most packages, but the header and pcf packages are excluded as they generate a lot of noise. So the initial tree will look something like this

```
root - Included
com - Included
ibm - Included
mq - Included
headers - Excluded
pcf - Excluded
```

When the trace facility is determining whether to include or exclude a package, it matches leading portions of the package name to the nodes in the tree as far as possible and takes the status of the last matched node. At the initial state of the tree, the packages com.ibm.msg.client and com.ibm.mq.jms would be included, as the last nodes in the tree that matches them (com->ibm and com->ibm->mq respectively) are marked as *Included*. Conversely, the package

com.ibm.headers.internal would be excluded as the last matching node in the tree (com->ibm->mq->headers) is marked as *Excluded*.

As further changes are made to the tree by using the com.ibm.mq.TraceControl.jar, it is important to remember that inclusion or exclusion only affects a package and child packages. So, given the initial state that is shown previously, specifying -ep com.ibm.mq.jms, would update the tree to look like this:

```
root - Included
com - Included
ibm - Included
mq - Included
headers - Excluded
jms - Excluded
pcf - Excluded
```

This update would exclude packages com.ibm.mq.jms, and com.ibm.mq.jms.internal, without affecting packages outside the com.ibm.mq.jms.* hierarchy.

If -ip com.ibm.mq.jms.admin is specified next, the tree would look like this:

```
root - Included
com - Included
ibm - Included
mq - Included
headers - Excluded
jms - Excluded
admin - Included
pcf - Excluded
```

This update would still exclude packages com.ibm.mq.jms, com.ibm.mq.jms.internal, but now the packages com.ibm.mq.jms.admin, and com.ibm.mq.jms.admin.internal are included in trace.

Connecting remotely

You can connect remotely only if the process was started with a JMX agent that is enabled for remote connection, and that uses the -Dcom.sun.management.jmxremote.port=port_number system setting.

After you have started with this system setting, you can run the executable JAR file with the options -h host_name -p port_number in place of the -i identifier option, where host_name is the name of the host you want to connect to and port_number is the name of the port to be used.

Note: You must ensure that you take appropriate steps to minimize security risks by enabling TLS for the connection. See the Oracle documentation on JMX for further details https://www.oracle.com.

Limitations

The following limitations exist:

• For non-IBM JVMs, the tool must be started with tools.jar added to its class path. The command that is on these platforms is:

```
java -cp MQ_INSTALL_DIR/java/lib/com.ibm.mq.traceControl.jar; JAVA_HOME/lib/tools.jar
com.ibm.msg.client.commonservices.trace.TraceController
```

• Local attach is controlled by user ID. The tool must be run under the same ID as the process that is to be controlled.

Tracing IBM MQ .NET applications

In IBM MQ .NET, you start and control the trace facility as in IBM MQ programs using the MQI.

However, the -i and -p parameters of the strmqtrc command, which allow you to specify process and thread identifiers, and named processes, have no effect.

You normally need to use the trace facility only at the request of IBM service.

See Using trace on Windows for information on trace commands.

Windows Z/OS Linux V 9.0.1 AIX Tracing the IBM MQ Console and REST API

The trace facilities in the IBM MQ Console and REST API are provided to help IBM staff to diagnose customer problems. Various properties control the behavior of these facilities.

The IBM MQ Console and REST API consist of three functional areas, each with their own trace mechanisms:

- The IBM MQ Console JavaScript code that executes in the browser.
- The IBM MQ Console and REST API code that runs in the mqweb server.
- The IBM MQ Classes for JMS code that runs in the mqweb server.

Enabling trace for the IBM MQ Console JavaScript code that runs in the browser

This trace is output only from the browser that it is enabled in. After you log out of the IBM MQ Console, trace is automatically disabled.

To enable trace for the IBM MQ Console JavaScript code that runs in the browser:

- 1. Log on to the IBM MQ Console
- 2. Click the dashboard menu icon, and select **Diagnostics**.
- 3. Select **Enable** for the IBM MQ Console browser trace, and click **OK**.
- 4. Follow the steps that are outlined to enable trace for the IBM MQ Console and REST API code running in the mqweb server.

Actions that are performed in your browser then start to be traced. This trace is periodically sent to the IBM MQ Console code that runs in the mgweb server, and is output in the mgweb server trace logs.

Enabling trace for the IBM MQ Console and REST API code that runs in the mqweb server

- 1. Locate the mqwebuser.xml file in one of the following directories:
 - Windows Linux AIX MQ_DATA_DIRECTORY/web/installations/installationName/servers/mqweb
 - WLP_user_directory/servers/mqweb

Where WLP_user_directory is the directory that was specified when the crtmqweb.sh script ran to create the mqweb server definition.

2. Add the following XML to the mawebuser.xml file, between the server tags:

```
<variable name="traceSpec"
value="*=info:com.ibm.mq*=all:com.ibm.mq.rest*=all:js.mq*=all" />
```

If the traceSpec variable exists in the mqwebuser.xml file, add the value attributes to the variable. Separate each value attribute with a colon.

If the may be server is running, trace is immediately enabled.

Trace is output to a set of files. The active file is called trace.log. Historical trace is kept in files that are called trace_timestamp.log. The size of these trace files, and the number of historical files that are

kept can be configured by setting the maxTraceFileSize and maxTraceFiles variables. V 9.0.1 For more information, see Configuring logging.

Enabling trace for the IBM MQ Classes for JMS code that runs in the mqweb server

- 1. Create a file called jmstrace.config in one of the following directories:
 - Windows Linux AlX MQ_DATA_DIRECTORY/web/installations/installationName/servers/mqweb
 - WLP user directory/servers/mgweb

Where WLP_user_directory is the directory that was specified when the crtmqweb.sh script ran to create the mqweb server definition.

2. Add the following lines to the jmstrace.config file:

```
com.ibm.msg.client.commonservices.trace.outputName=PATH/logs/jmstrace.txt
com.ibm.msg.client.commonservices.trace.limit=104857600
com.ibm.msg.client.commonservices.trace.count=10
com.ibm.msg.client.commonservices.trace.status=0N
```

Where PATH specifies the full path to the directory where you want the jmstrace.txt file to be written.

These lines set the maximum trace file size to 100 MB, and set the maximum number of trace files to 10. Ensure that you have disk space available for these files.

- 3. In the same directory as the jmstrace.config file, open or create the jvm.options file.
- 4. Add the following lines to the jvm. options file:

```
-Dcom.ibm.msg.client.commonservices.trace.startup=TRUE
-Dcom.ibm.msg.client.config.location=CONFIG_PATH/jmstrace.config
```

Where CONFIG_PATH specifies the full path to the directory where the jmstrace.config file is located, as a URL. For example, file:c:/ProgramData/IBM/MQ/web/installations/Installation2/servers/mgweb/.

5. Restart the may be server by using the **endmayeb** and **strmqweb** commands on the command line.

Information to provide to IBM Service

Include the following files and directories when you gather diagnostic information for IBM Service:

- The mayeb.xml file.
- The contents of the directory that contains the maweb server definition:
 - Windows Linux AIX MQ_DATA_DIRECTORY/web/installations/installationName
 - _ z/0S

The directory that was specified when the crtmqweb.sh script ran to create the mqweb server definition. By default, this directory is /var/mqm/web/installation1.

Tracing the IBM MQ Bridge to Salesforce

The trace facilities for the IBM MQ Bridge to Salesforce are provided to help IBM staff to diagnose customer problems. Enable the trace for the IBM MQ Bridge to Salesforce and define the debug level when you issue the **runmqsfb** command to start the bridge.

Procedure

1. Set the environment variable MQSFB_EXTRA_JAVA_OPTIONS to specify the **-D** Java option and turn on the IBM MQ classes for JMS trace.

```
export MQSFB_EXTRA_JAVA_OPTIONS="-Dcom.ibm.msg.client.commonservices.trace.status=ON"
```

2. Set the debug level to verbose mode **-d** 2 when you issue the **runmqsfb** command at run time.

```
runmqsfb -f new_config.cfg -r logFile.log -d 2
```

Your logFile.log contains information that might be helpful in resolving your problem with the IBM MQ Bridge to Salesforce.

3. Optional: You can achieve finer control over the exact trace by creating the IBM MQ classes for JMS configuration file. For more information, see "Tracing IBM MQ classes for JMS applications" on page 87 and follow the advice that is provided by your IBM service support representative.

Related tasks

Running the IBM MQ Bridge to Salesforce Monitoring the IBM MQ Bridge to Salesforce

Related reference

runmqsfb (run IBM MQ Bridge to Salesforce)



blockchain

The trace facilities for the IBM MQ Bridge to blockchain are provided to help IBM staff to diagnose customer problems. Enable the trace for the IBM MQ Bridge to blockchain and define the debug level when you issue the **runmqbcb** command to start the bridge.

Procedure

1. Set the environment variable MQBCB_EXTRA_JAVA_OPTIONS to specify the **-D** Java option and turn on the IBM MQ classes for JMS trace.

```
export MQBCB_EXTRA_JAVA_OPTIONS="-Dcom.ibm.msg.client.commonservices.trace.status=ON"
```

2. Set the debug level to verbose mode -d 2 when you issue the runmmbcb command at run time.

```
./runmqbcb.sh -f new_config.cfg -r logFile.log -d 2
```

Your logFile.log contains information that might be helpful in resolving your problem with the IBM MQ Bridge to blockchain.

3. Optional: You can achieve finer control over the exact trace by creating the IBM MQ classes for JMS configuration file. For more information, see "Tracing IBM MQ classes for JMS applications" on page 87 and follow the advice that is provided by your IBM service support representative.

Related tasks

Running the IBM MQ Bridge to blockchain

Related reference

runmqbcb (run IBM MQ Bridge to Blockchain)

Y 9.0.0.3 Enabling dynamic tracing of LDAP client library code

From IBM MQ 9.0.0 Fix Pack 9, it is possible to switch LDAP client trace on and off without also stopping or starting the queue manager.

About this task

Before IBM MQ 9.0.0 Fix Pack 9, it was not possible to switch the LDAP client trace on and off without also stopping or starting the queue manager.

From IBM MQ 9.0.0 Fix Pack 9, you can switch LDAP client trace on with the **strmqtrc** command and off with the **endmqtrc** command without needing to stop or start the queue manager. To enable this behavior, it is also necessary to set an environment variable **AMQ_LDAP_TRACE** to a non-null value.

When the **AMQ_LDAP_TRACE** is set to a non-null value, and the LDAP functionality is used, some queue manager processes create zero length files under /var/mqm/trace. When the trace is then switched on using the **strmqtrc** command, some trace information is written to these files. Later, when trace is switched off with the **endmqtrc** command, trace information ceases to be written to the files, but handles to the files remain open until the queue manager ends.

On UNIX platforms, filesystem space cannot be released completely simply by unlinking these files with the **rm** command. This is a side-effect from the fact that the handles remain open. Therefore, a queue manager end should be performed, whenever disk space in /var/mqm/trace needs to be released.

Procedure

- Set the environment variable AMQ_LDAP_TRACE to a non-null value.
- Use the **strmqtrc** command to switch the trace on:

```
strmqtrc -m QMNAME -t servicedata
```

• Use the **endmqtrc** command to switch the trace off.

Problem determination on z/OS

IBM 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.
- IBM MQ for z/OS abends, and the information produced when an abend occurs.
- The diagnostic information produced by IBM 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 subtopics for more information about problem determination and diagnostic information on IBM MQ for z/OS.

- "IBM MQ for z/OS performance constraints" on page 107
- "IBM MQ for z/OS recovery actions" on page 109
- "IBM MQ for z/OS abends" on page 109
- "Diagnostic information produced on IBM MQ for z/OS" on page 112
- "Other sources of problem determination information for IBM MQ for z/OS" on page 115
- "Diagnostic aids for CICS" on page 116
- "Diagnostic aids for IMS" on page 116
- "Diagnostic aids for Db2" on page 116
- "IBM MQ for z/OS dumps" on page 117
- "Dealing with performance problems on z/OS" on page 136
- "Dealing with incorrect output on z/OS" on page 143

Related concepts

"Troubleshooting overview" on page 7

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?"

"Using error logs" on page 42

There are a variety of error logs that you can use to help with problem determination and troubleshooting.

"First Failure Support Technology (FFST)" on page 51

First Failure Support Technology (FFST) for IBM MQ provides information about events that, in the case of an error, can help IBM support personnel to diagnose the problem.

Related tasks

"Using trace" on page 63

You can use different types of trace to help you with problem determination and troubleshooting.

IBM 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 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 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 Using CSQ6LOGP 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.

Incorrect specification of buffer pool size can adversely affect IBM 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 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 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 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 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 <u>Using CSQ6SYSP</u>. 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 START TRACE for more information.

IBM MQ for z/OS recovery actions

Use this topic to understand some of the recovery actions for user detected and queue manager detected errors.

IBM 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 MQ for z/OS messages, completion, and reason codes.

Program errors

Program errors might be associated with user application program code or IBM 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 IBM MQ function has completed. Rerunning the application with the IBM 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 73 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

IBM 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), IBM 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 IBM 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.

IBM MQ for z/OS abends

Abends can occur in WebSphere for z/OS or other z/OS systems. Use this topic to understand the IBM MQ system abend codes and how to investigate abends which occur in CICS, IMS, and z/OS.

IBM 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 IBM MQ has detected an internal error and has terminated an internal task (TCB) or a user-connected task abnormally. Errors associated with an 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 IBM MQ has detected a severe error and has terminated the queue manager abnormally. When an X'6C6' is issued, IBM MQ has determined that continued operation could result in the loss of data integrity. Errors associated with an 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 IBM MQ.

Table 10 on page 110 summarizes the actions and diagnostic information available to IBM 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 MQ for z/OS" on page 112.

Table 10. Abend completion codes		
	X'5C6'	X'6C6'
Explanation	Error during IBM MQ normal operation	Severe error; continued operation might jeopardize data integrity
System action	Internal IBM MQ task is abended Connected user task is abended	The entire IBM MQ subsystem is abended User task with an active IBM MQ connection might be abnormally terminated with an X'6C6' code Possible MEMTERM (memory termination) of connected allied address space
Diagnostic information	SVC dump SYS1.LOGREC entry VRA data entries	SYS1.LOGREC VRA data entries
Associated reason codes	IBM MQ abend reason code Associated z/OS system codes	 Subsystem termination reason code z/OS system completion codes and X'5C6' codes that precede the X'6C6' abend
Location of accompanying codes	 SVC dump title Message CSQW050I Register 15 of SDWA section General Purpose Registers at Time of Error SYS1.LOGREC entries VRA data entries 	 SYS1.LOGREC VRA data entries Message CSQV086E, which is sent to z/OS system operator

"Dealing with abends on IBM MQ for z/OS" on page 111

Abends can occur with applications and other z/OS systems. Use this topic to investigate program abends, batch abends, CICS transaction abends, and IMS transaction abends.

"CICS, IMS, and z/OS abends" on page 112

Use this topic to investigate abends from CICS, IMS, and z/OS.

"Diagnostic information produced on IBM MQ for z/OS" on page 112

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.

"IBM MQ for z/OS dumps" on page 117

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 an IBM MQ for z/OS address space.

Dealing with abends on IBM MQ for z/OS

Abends can occur with applications and other z/OS systems. Use this topic to investigate program abends, batch abends, CICS transaction abends, and IMS transaction abends.

Types of abend

Program abends can be caused by applications failing to check, and respond to, reason codes from IBM 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 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 MQ for z/OS abends (abend codes X ' 5C6 ' and X ' 6C6 '), see IBM MQ for z/OS messages, completion, and reason codes
- For batch abends, the z/OS MVS System Codes manual
- · For CICS abends, CICS Messages
- For IMS abends, IMS Messages and Codes
- For Db2 abends, Messages
- Db2
- For RRS abends, z/OS MVS System Messages, Volume 3
- For XES abends, z/OS MVS System Messages, Volume 10

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 "IBM MQ for z/OS dumps" on page 117).

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 z/OS" on page 141 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 <u>CICS Troubleshooting</u>, formerly the <u>CICS Problem Determination Guide</u> 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 Troubleshooting for IMS.

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 MQ for z/OS messages, completion, and reason codes

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 IBM 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 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.

IBM 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 IBM 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 information is 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 IBM 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.

IBM 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

IBM MQ produces an error message when a problem is detected. IBM MQ diagnostic messages begin with the prefix CSQ. Each error message generated by IBM MQ is unique; that is, it is generated for one and only one error. Information about the error can be found in IBM MQ for z/OS messages, completion, and reason codes.

The first three characters of the names of IBM 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. 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 MQ for z/OS messages, completion, and reason codes

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 "IBM MQ for z/OS dumps" on page 117.

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. "IBM MQ for z/OS dumps" on page 117 contains guidance about using dumps to locate problems in your IBM 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 IBM 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 7 on page 114 shows an example of a symptom string.

PIDS/ 5655R3600 RIDS/CSQMAIN1 AB/S6C6 PRCS/0E30003

Figure 7. 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.

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.

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 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"Other sources of problem determination information for IBM MQ for z/OS" on page 115
Use this topic to investigate other sources of information for IBM MQ for z/OS problem determination.

"Diagnostic aids for CICS" on page 116

You can use the CICS diagnostic transactions to display information about queue manager tasks, and MQI calls. Use this topic to investigate these facilities.

"Diagnostic aids for IMS" on page 116

Use this topic to investigate IMS diagnostic facilities.

"Diagnostic aids for Db2" on page 116

Use this topic to investigate references for Db2 diagnostic tools.

Other sources of problem determination information for IBM MQ for z/OS

Use this topic to investigate other sources of information for IBM MQ for z/OS problem determination.

You might find the following items of documentation useful when solving problems with IBM 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 IBM 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 IBM 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 IBM 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 <u>Administering IBM MQ for z/OS</u> 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 IBM 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.

The CICS adapter also writes trace entries to the CICS trace. These entries are described in <u>"CICS adapter"</u> trace entries" on page 81.

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 IBM 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.

208 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

z/08 V 9.0.3 MQ Adv. VUE Troubleshooting the connection to Product

Insights from IBM MQ for z/OS

The IBM Cloud® Product Insights service is no longer available. For more information, see this blog post: Service Deprecation: IBM Cloud Product Insights.

IBM MQ for z/OS 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 an IBM MQ for z/OS address space.

How to use dumps for problem determination

When solving problems with your IBM MQ for z/OS system, you can use dumps in two ways:

- To examine the way IBM 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 IBM 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 with IBM MQ for z/OS" on page 118
- "Using the z/OS DUMP command" on page 118
- "Processing a dump using the IBM MQ for z/OS dump display panels" on page 120
- "Processing an IBM MQ for z/OS dump using line mode IPCS" on page 124
- "Processing an IBM MQ for z/OS dump using IPCS in batch" on page 131

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 z/OS" on page 132. For information about the IBM MQ for z/OS abend codes, see "IBM MQ for z/OS abends" on page 109, and abend reason codes are documented in IBM MQ for z/OS messages, completion, and reason codes.

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 z/OS" on page 134
- "Snap dumps on z/OS" on page 135
- "SYS1.LOGREC information on z/OS" on page 135
- "SVC dumps on z/OS" on page 136

Related concepts

"Using trace for problem determination on z/OS" on page 73

There are different trace options that can be used for problem determination with IBM MQ. Use this topic to understand the different options and how to control trace.

"IBM MQ for z/OS abends" on page 109

Abends can occur in WebSphere for z/OS or other z/OS systems. Use this topic to understand the IBM MQ system abend codes and how to investigate abends which occur in CICS, IMS, and z/OS.

"Diagnostic information produced on IBM MQ for z/OS" on page 112

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.

Getting a dump with IBM MQ for z/OS

Use this topic to understand the different dump types for IBM MQ for z/OS problem determination.

The following table shows information about the types of dump used with IBM MQ for z/OS and how they are initiated. It also shows how the dump is formatted:

Table 11. Types of dump used with IBM MQ for z/OS				
Dump type	Data set	Output type	Formatted by	Caused by
SVC	Defined by system	Machine readable	IPCS in conjunction with an IBM MQ for z/OS verb exit	z/OS or IBM 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 an IBM MQ for z/OS verb exit	Operator IPL of the stand- alone dump program

IBM MQ for z/OS recovery routines request SVC dumps for most X'5C6' abends. The exceptions are listed in <u>"SVC dumps on z/OS" on page 136</u>. SVC dumps issued by IBM MQ for z/OS are the primary source of diagnostic information for problems.

If the dump is initiated by the IBM 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 z/OS 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 IBM MQ address space
- · Channel initiator address space
- · Coupling facility application structure
- Coupling facility administration structure for your queue sharing group

<u>Figure 8 on page 119</u> through to <u>Figure 12 on page 120</u> 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 8. Dumping the IBM 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 9. Dumping the IBM 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 10. 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 11. Dumping the IBM 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 12. Dumping a coupling facility structure

Processing a dump using the IBM 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.

IBM MQ for z/OS provides a set of panels to help you process dumps. The following section describes how to use these panels:

- 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 required. 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
           Component trace summary
  CTRACE
  DAEDATA DAE header data
  DIVDATA Data-in-virtual storage
```

Figure 13. IPCS MVS Dump Component Data Analysis panel

3. Select **CSQMAIN IBM MQ dump formatter panel interface** by typing s next to 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 for more information about installing the IBM 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 IBM MQ dump formatter** to display the formatted contents again, using the default options.

4. The IBM 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.

```
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, 2025. 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 "5.a" on page 121 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.
 - 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 end of the data set, type FIND ERRORID: PREV, and press Enter. The header of the latest LOGREC record is displayed, for example:

Note the program identifier (if it is not 5655R3600, the problem was not caused by IBM 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.

• To format and display the dump, select option 8. The FORMAT AND DISPLAY THE DUMP panel is displayed:

```
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_
```

• 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 IBM 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 xxxx MSTR, where xxxx is the subsystem name.

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 xxxx MSTR, 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 an IBM MQ for z/OS dump using line mode IPCS" on page 124

Use the IPCS commands to format a dump.

"Processing an IBM MQ for z/OS dump using IPCS in batch" on page 131

Use this topic to understand how IBM MQ for z/OS dumps can be formatted by IPCS commands in batch mode.

"Analyzing the dump and interpreting dump titles on z/OS" on page 132

Use this topic to understand how IBM MQ for z/OS dump titles are formatted, and how to analyze a dump.

Processing an IBM MQ for z/OS 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.DUMP xx')
```

(where SYS1.DUMP xx 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 an IBM MQ for z/OS dump" on page 124
- "Formatting a dump from the channel initiator on z/OS" on page 130

Related concepts

"Processing a dump using the IBM MQ for z/OS dump display panels" on page 120

You can use commands available through IPCS panels to process dumps. Use this topic to understand the IPCS options.

"Processing an IBM MQ for z/OS dump using IPCS in batch" on page 131

Use this topic to understand how IBM MQ for z/OS dumps can be formatted by IPCS commands in batch mode.

"Analyzing the dump and interpreting dump titles on z/OS" on page 132

Use this topic to understand how IBM MQ for z/OS dump titles are formatted, and how to analyze a dump.

Formatting an IBM 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 IBM MQ for z/OS dump formatting program (CSQWDPRD), and enables you to format an SVC dump to display IBM 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 *z/OS MVS IPCS Commands* manual.

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 12 on page 124 explains each keyword that you can specify in the control statement for formatting dumps.

Table 12. Keywords for the IBM 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.	
ALL (default)	All control blocks and the trace table.	

Table 12. Keywords for the IBM MQ for z/OS dump formatting control statement (continued)		
Keyword	Description	
AA	Data is displayed for all IBM 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= nnnnnnn	Only the trace points associated with this EB thread are displayed (the format of this keyword is EB= nnnnnnnn where nnnnnnnn 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= load module name	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= hhhh	The control blocks for a specified address space. Use either of the following formats:	
	• SA= <i>hh</i> or	
	• SA= hhhh	
	where <i>h</i> represents a hexadecimal digit.	
SG	A subset of system-wide control blocks.	
TT	Format trace table	
,HANDLES=x	Indicate threads with greater than x handles	
,LOCKS=x	Indicate threads with greater than x locks	
,INSYNCS=x	Indicate threads with greater than x insync operations	
,URINFO=ALL/LONG	Show UR info for ALL threads or for long-running threads	

 $\underline{\text{Table 13 on page 125}} \text{ details the dump formatting keywords that you can use to format the data relating to individual resource managers.}$

You cannot use these keywords in conjunction with any of the keywords in Table 12 on page 124.

Table 13. Resource manager dump formatting keywords		
Keyword	What is formatted	
BMC=1	Buffer manager data. BMC=1 formats control blocks of all buffers.	
BMC=2(buffer pool number) BMC=3(xx/yyyyyy)	BMC=2 formats data relating to the buffer identified in the 2-digit <i>buffer</i> pool number.	
BMC=4(xx/yyyyyy)	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.)	
BUFL= nnnnnnnnnn	Storage access buffer allocation sz.	

Table 13. Resource manager dump formatting keywords (continued) Keyword What is formatted		
CALLD=Y		
	Show arrow for call depth in TT.	
=W	and indent trace entry.	
CALLTIME=Y	Print call time on exit trace.	
CB=(addr/[strmodel])	Format address as IBM MQ block.	
CBF=1	CBF report level 1.	
CCB=S	Show the Composite Capability Block (CCB) for system EBs in TT.	
CFS=1	CFS report level 1.	
CFS=2	CFS report level 2.	
CHLAUTH=1/2	CHLAUTH report level.	
ONAM=20 chars	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.	
CLUS=1	Cluster report including the cluster repository known on the queue manager.	
CLUS=2	Cluster report showing cluster registrations.	
CLXQ=1	Cluster XMITQ report level 1.	
CLXQ=2	Cluster XMITQ report level 2.	
ONAM=20 chars	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.	
CMD=0/1/2	Command trace table display level.	
D=1/2/3	Detail level for some reports.	
Db2=1	Db2 report level 1.	
DMC=1,	DMC report level 1.	
ONAM=48 chars	The optional ONAM subparameter allows you to specify the object name, up to a maximum of 48 characters, to limit data printed to objects starting with characters in ONAM.	
DMC=2,	DMC report level 2.	
ONAM=48 chars	The optional ONAM subparameter allows you to limit the objects printed to those with names beginning with the characters specified in ONAM (up to a maximum of 48 characters).	
DMC=3,	DMC report level 3.	
ONAM=48 chars	The optional ONAM subparameter allows you to limit the objects printed to those with names beginning with the characters specified in ONAM (up to a maximum of 48 characters).	
GR=1	Group indoubt report level 1.	
IMS=1	IMS report level 1	

Table 14. Resource manager dump	formatting keywords (J -P)
Keyword	What is formatted
JOBNAME= xxxxxxxx	Job name
LKM=1	LKM report level 1.
LKM=2/3,	LKM report level 2/3.
,NAME=up to 48 chars	Name (character)
,NAMEX= xxxxxxxxxxxxxxx	Name (Hex)
,NAMESP=1/2/3/4/5/6/7/8	Namespace
,TYPE=DMCP/QUALNM/TOPIC/	Lock type
STGCLASS	Lock qualification
,QUAL=GET/PUT/CRE/DFXQ/	
PGSYNC/CHGCNT/	LKM report level 3
DELETE/EXPIRE	LKM report level 4
LKM=3	
LKM=4	
,JOBNAME= xxxxxxxx	
,ASID= xxxx	
LMC=1	LMC report level 1.
MAXTR= nnnnnnnn	Max trace entries to format
MHASID= xxxx	Message handle ASID for properties
MMC=1	MMC report level 1
OBJ=MQLO/MQSH/MQRO/	
MQAO/MQMO/MCHL/	Object type
MNLS/MSTC/MPRC/:	The optional ONAM subparameter allows you to limit the objects printed
MAUT	to those with names beginning with the characters specified in ONAM (up to a maximum of 48 characters).
ONAM	
MMC=2	MMC report level 2
ONAM=48 chars	The optional ONAM subparameter allows you to limit the objects printed to those with names beginning with the characters specified in ONAM (up to a maximum of 48 characters).
MSG=nnnnnnnnnnnnnn	Format the message at pointer.
MASID=xxxx	MASID allows storage in other address spaces.
LEN=xxxxxxxx	LEN limits amount of storage to format.
MSGD=S/D	MSGD controls level of detail.
MSGD=S/D	Message details in DMC=3, BMC=3/4, PSID reports.
	The parameter controls level of details, S is summary and D is detailed.
MSGH = nnnnnnnnnnnnnnn	Message handle

Table 14. Resource manager dump formatting keywords (J -P) (continued)		
Keyword	What is formatted	
MT	Message properties trace	
MQVCX	MQCHARVs in hexadecimal format	
PROPS= nnnnnnnnnnnnnn	Message properties pointer	
PSID= nnnnnnn	Pageset to format page	
PSTRX	Properties strings in hex format	

Table 15. Resource manager dump formatting keywords (R -Z)		
Keyword	What is formatted	
RPR= nnnnnnn	Page or record to format	
SHOWDEL	Show deleted records for DMC=3	
SMC=1/2/3	Storage manager	
TC=	TT data char format, concatenated	
*	print all in suitable character set	
A	always print ASCII	
E	always print EBCDIC	
0	never print either	
TFMT=H/M	Time format - human or STCK	
THR= nnnnnnn	Thread address	
THR=*/2/3	Set thread report level	
TOP=1	TOP report level 1	
TOP=2	TOP report level 2	
TOP= nnnnnnnnnnnnn	Tnode 64bit address or	
/TSTR=48 chars	Topic string (wildcard with % at start or end)'	
	This will be converted EBCDIC to ASCII, but only invariant characters	
/TSTRX=hex 1208 str	Hexadecimal of topic string in 1208 always wildcard character at start.	
TOP=3	TOP report level 3	
TOP=4	TOP report level 4	
TSEG=M(RU)/Q(P64)	Search process for 64-bit trace	
I(NTERPOLATE)	Guess missing TSEG address or addresses	
F(WD)	Force forward sort	
D(EBUG)	Debug search process	
TSEG=(M,Q,I,F,D)	Specify multiple TSEG options	
W=0/1/2/3	TT width format	
XA=1	XA report level 1	
ZMH = nnnnnnnnnnnnnn	ZST message handle	

If the dump is initiated by the operator, there is no information in the summary portion of the dump. <u>Table</u> 16 on page 129 shows additional keywords that you can use in the CSQWDMP control statement.

Table 16. Summary dump keywords for the IBM 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 17 on page 129 shows some other commands that are used frequently for analyzing dumps. For more information about these sub commands, see the *z/OS MVS IPCS Commands* manual.

Table 17. 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.	

Table 17. IPCS subcommands used for dump analysis (continued)		
Subcommand	Description	
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.	

"Formatting a dump from the channel initiator on z/OS" on page 130

Use this topic to understand how to format a channel initiator dump for IBM MQ for z/OS using line mode IPCS commands.

Formatting a dump from the channel initiator on z/OS

Use this topic to understand how to format a channel initiator dump for IBM MQ for z/OS 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 18 on page 130 describes the keywords that you can specify with CSQXDPRD.

Table 18. 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:
	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.

Table 18. Keywords for the IPCS VERBEXIT CSQXDPRD (continued)	
Keyword	What is formatted
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,	Select either a short (CTRACE=S) or full (CTRACE=F) CTRACE.
DPRO= nnnnnnn,	The optional DPRO subparameter allows you to specify a CTRACE for the DPRO specified.
TCB= nnnnnn	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
XSMF=1	Format channel initiator SMF data that is available in a dump.

"Formatting an IBM MQ for z/OS dump" on page 124

Use this topic to understand how to format a queue manager dump using line mode IPCS commands.

Processing an IBM MQ for z/OS dump using IPCS in batch

Use this topic to understand how IBM MQ for z/OS 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 <u>Figure 14 on page 132</u>).

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 TERMINÀL DĎNAME(DUMP00) NOCONFIRM
 \star INSERT YOUR IPCS COMMANDS HERE, FOR EXAMPLE: \star VERBEXIT LOGDATA
 VERBEXIT SYMPTOM
 VERBEXIT CSOWDMP 'TT, SUBSYS=OMGR'
 **************
 CLOSE ALL
 END
Figure 14. Sample JCL for printing dumps through IPCS in the z/OS environment
```

"Processing a dump using the IBM MQ for z/OS dump display panels" on page 120

You can use commands available through IPCS panels to process dumps. Use this topic to understand the IPCS options.

"Processing an IBM MQ for z/OS dump using line mode IPCS" on page 124 Use the IPCS commands to format a dump.

"Analyzing the dump and interpreting dump titles on z/OS" on page 132 Use this topic to understand how IBM MQ for z/OS dump titles are formatted, and how to analyze a dump.

z/0S Analyzing the dump and interpreting dump titles on z/OS

Use this topic to understand how IBM MQ for z/OS 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 15 on page 132 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 15. 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 MQ for z/OS messages, completion, and reason codes 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. The value R3600 uniquely identifies IBM MO for z/OS.
- release is a 3-digit code indicating the version, release, and modification level of IBM 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 16 on page 133 illustrates this format.

 ${\tt ssnm, ABN=compltn-reason, U=userid, C=compid.release.comp-function, M=module, PSW=psw_contents, ASID=address_space_id}$

Figure 16. 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.

"Processing a dump using the IBM MQ for z/OS dump display panels" on page 120

You can use commands available through IPCS panels to process dumps. Use this topic to understand the IPCS options.

"Processing an IBM MQ for z/OS dump using line mode IPCS" on page 124 Use the IPCS commands to format a dump.

"Processing an IBM MQ for z/OS dump using IPCS in batch" on page 131
Use this topic to understand how IBM MQ for z/OS dumps can be formatted by IPCS commands in batch mode.

SYSUDUMP information on z/OS

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 *z/OS MVS Diagnosis: Tools and Service Aids* manual.

Figure 17 on page 134 shows a sample of the beginning of a SYSUDUMP dump.

```
JOB MQMBXBA1 STEP TSOUSER TIME 102912 DATE 001019 ID = 000 CPUID = 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
                ASCB
                          FWDP..... 00F60180 BWDP..... 0047800 CMSF..... 019D5A30
 +0000 ASCB....
 SVRB..... 008FE9E0
 +0014 SYNC..... 00000D6F IOSP..... 000000000 TNEW..... 00D18F0 CPUS..... 000000001
 ASID.... 0066
 +0026 R026.... 0000
DP..... 9D
                          LL5..... 00
                                              HLHI.... 01
                                                                 DPHI.... 00
 +002C TRQP..... 80F5D381 LDA..... 7FF154E8 RSMF..... 00
                                                                 R035.... 0000
 TRQI..... 42
 +0038 CSCB..... 00F4D048 TSB..... 00B61938 EJST..... 0000001 8C257E00
               . 9CCDE747 76A09480
 +0048 EWST...
                                              JSTL..... 00141A4 ECB..... 808FEF78
 UBET..... 9CCDE740
 ASSB: 01946600
 +0000 ASSB.... ASSB
                          VAFN..... 00000000 EVST..... 0000000 00000000
 +0010 VFAT.... 00000000 00000000
                                              RSV..... 000
                                                                 XMCC.... 0000
 XMCT.....00000000
 +0020 VSC...... 00000000 NVSC..... 0000004C ASRR..... 0000000 R02C..... 00000000
 00000000 00000000
                00000000 00000000
 +0038
 *** 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
               FF DSP..... FE
 +0022 LMP.....
                                              LLS..... 00D1A88 JLB..... 00011F18
 JPQ.....00000000
 +0030 GPRO-3... 00001000 008A4000 00000000 00000000
 +0040 GPR4-7... 00FDC730 008A50C8 00000002 80E73F04
+0050 GPR8-11.. 81CC4360 008A6754 008A67B4 00000008
Figure 17. Sample beginning of a SYSUDUMP
```

Snap dumps on z/OS

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 on z/OS

Use this topic to understand how the z/OS SYS1.LOGREC information can assist with problem determination.

IBM 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 <u>z/OS MVS Diagnosis: Tools</u> and Service Aids manual.

IBM 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:

- See <u>EREP Selection Parameters</u>, described in the *z/OS 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 IBM MQ for z/OS dump display panels" on page 120).

Only records available in storage when the dump was requested are included. Each formatted record follows the heading *****LOGDATA****.

SVC dumps on z/OS

Use this topic to understand how to suppress SVC dumps on z/OS, 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 <u>IEACMD00</u> in the *z/OS 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 IBM MQ for z/OS dumps using z/OS DAE

You can suppress SVC dumps that duplicate previous dumps. The <u>z/OS MVS Diagnosis: Tools and Service</u> Aids manual gives details about using z/OS dump analysis and elimination (DAE).

To support DAE, IBM 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')

IBM 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 IBM MQ for z/OS 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 MQ and CICS or IMS.

- "IBM MQ for z/OS system considerations" on page 137
- "CICS constraints" on page 137
- "Dealing with applications that are running slowly or have stopped on z/OS" on page 137

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 incorrect output on z/OS" on page 143

Incorrect output can be missing, unexpected, or corrupted information. Read this topic to investigate further.

Related tasks

"Making initial checks" on page 8

There are some initial checks that you can make that may provide answers to common problems that you may have.

IBM 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.

Z/08 CICS constraints

CICS constraints can also have an adverse effect on IBM MQ for z/OS performance. Use this topic for further information about CICS constraints.

Performance of IBM 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 on z/OS.

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 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 on z/OS" on page 138
- "Dealing with waits on z/OS" on page 139
- "Dealing with loops on z/OS" on page 141

Related tasks

"Making initial checks" on page 8

There are some initial checks that you can make that may provide answers to common problems that you may have.

Distinguishing between waits and loops on z/OS

Waits and loops on IBM MQ for z/OS 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 appropriate 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
- IBM 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 IBM 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 on z/OS" on page 139

Waits can occur in batch or TSO applications, CICS transactions, and other components on IBM MQ for z/OS. Use this topic to determine where waits can occur.

"Dealing with loops on z/OS" on page 141

Loops can occur in different areas of a $z/\overline{O}S$ system. Use this topic to help determine where a loop is occurring.

Dealing with waits on z/OS

Waits can occur in batch or TSO applications, CICS transactions, and other components on IBM MQ for z/OS. 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 IBM MQ, take an SVC dump of both the batch or TSO job, and the IBM 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 z/OS" on page 143), or in IBM MQ itself (see "Is IBM MQ waiting for z/OS?" on page 141). Refer to "IBM MQ for z/OS dumps" on page 117 (specifically Figure 8 on page 119) for information about obtaining a dump.

If the problem persists, refer to "Collecting troubleshooting information" on page 41 for further guidance.

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 IBM 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 IBM 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 z/OS" on page 143). 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 z/OS" on page 145).

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 IBM MQ, take an SVC dump of both the CICS region, and the IBM MQ subsystem before canceling the CICS transaction. Refer to "Dealing with loops on z/OS" on page 141 for information about waits. Refer to "IBM MQ for z/OS dumps" on page 117 (specifically Figure 8 on page 119) for information about obtaining a dump.

If the problem persists, refer to "Collecting troubleshooting information" on page 41 for further guidance.

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 IBM MQ waiting for z/OS?

If your investigations indicate that IBM MQ itself is waiting, check the following:

- 1. Use the DISPLAY THREAD(*) command to check if anything is connected to IBM 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 IBM MQ for z/OS stopped processing work?" on page 31).
 - If IBM MQ is using some processor time, reconsider other reasons why IBM MQ might be waiting, or consider whether this is actually a performance problem.
 - If there is no processor activity, check whether IBM MQ responds to commands. If you can get a response, reconsider other reasons why IBM MQ might be waiting.
 - If you cannot get a response, check the console log for messages that might explain the wait (for example, IBM MQ might have run out of active log data sets, and be waiting for offload processing).

If you are satisfied that IBM 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 "Collecting troubleshooting information" on page 41 for further guidance.

Related concepts

"Distinguishing between waits and loops on z/OS" on page 138

Waits and loops on IBM MQ for z/OS can present similar symptoms. Use this topic to help determine if you are experiencing a wait or a loop.

"Dealing with loops on z/OS" on page 141

Loops can occur in different areas of a z/OS system. Use this topic to help determine where a loop is occurring.

Dealing with loops on z/OS

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 Reference* l.

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 IBM MQ for z/OS stopped processing work?" on page 31 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, refer to "Collecting troubleshooting information" on page 41 for further guidance.

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 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 MQ for z/OS performance constraints 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 IBM 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 z/OS" on page 138

Waits and loops on IBM MQ for z/OS can present similar symptoms. Use this topic to help determine if you are experiencing a wait or a loop.

"Dealing with waits on z/OS" on page 139

Waits can occur in batch or TSO applications, CICS transactions, and other components on IBM MQ for z/OS. Use this topic to determine where waits can occur.

Dealing with incorrect output on z/OS

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 <u>"Have you obtained incorrect output?" on page 39.</u>

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 on z/OS" on page 143
- "Problems with missing messages when using distributed queuing on z/OS" on page 145
- "Problems with getting messages when using message grouping on z/OS" on page 146
- "Finding messages sent to a cluster queue on z/OS" on page 147
- "Finding messages sent to the IBM MQ IMS bridge" on page 147
- "Messages contain unexpected or corrupted information on z/OS" on page 148

Related concepts

"Dealing with performance problems on z/OS" on page 136

Use this topic to investigate IBM MQ for z/OS performance problems in more detail.

Related tasks

"Making initial checks" on page 8

There are some initial checks that you can make that may provide answers to common problems that you may have.

Messages do not arrive when expected on z/OS

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 IBM 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 (MsqId 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 IBM 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 z/OS" on page 148.

Problems with missing messages when using distributed queuing on z/OS

Use this topic to understand possible causes of missing messages when using distributed queuing on IBM MQ for z/OS.

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 have been followed correctly.

Are the links available between the two systems?

Check that both systems are available, and connected to IBM 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 for more information about trace-route messaging in a network.

Is the channel running?

• Issue the following command for the transmission queue:

```
DISPLAY QUEUE (gname) 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 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 Distributed queuing and clusters.

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) 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 on z/OS

Use this topic to understand some of the issues with getting messages when using message grouping on IBM MQ for z/OS.

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 z/OS" on page 145. 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 on z/OS

Use this topic to understand some of the issues involved with finding messages sent to a cluster queue on IBM MQ for z/OS.

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 MQ - IMS bridge

Use this topic to understand possible causes for missing messages sent to the IBM MQ - IMS bridge.

If you are using the IBM MQ - IMS bridge, and your message has not arrived as expected, consider the following:

Is the IBM 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 MQ connected to OTMA?

Note: There are two IBM 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 MQ is connected to OTMA. This message also tells you to which IBM MQ system OTMA is connected. For more information about the content of these messages, see <u>IBM MQ for z/OS</u> messages, completion, and reason codes.

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 on z/OS

Use this topic to understand some of the issues that can cause unexpected or corrupted output on z/OS.

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.

Dealing with issues when capturing SMF data for the channel initiator (CHINIT)

Channel accounting and CHINIT statistics SMF data might not be captured for various reasons.

For more information, see:

Related concepts

Layout of SMF records for the channel initiator

Troubleshooting channel accounting data

Checks to carry out if channel accounting SMF data is not being produced for channels.

Procedure

- 1. Check that you have STATCHL set, either at the queue manager or the channel level.
 - A value of OFF at channel level means that data is not collected for this channel.
 - A value of OFF at queue manager level means data is not collected for channels with STATCHL(QMGR).
 - A value of NONE (only applicable at queue manager level) means data is not collected for all channels, regardless of their STATCHL setting.
- 2. For client channels check that STATCHL is set at the queue manager level.
- 3. For automatically defined cluster sender channels, check that the STATACLS is set.
- 4. Issue the display trace command. You need TRACE(A) CLASS(4) enabled for channel accounting data to be collected.
- 5. If the trace is enabled, SMF data is written:
 - On a timed interval depending on the value of the STATIME system parameter. A value of zero
 means that the SMF statistics broadcast is used. Use the DIS SYSTEM command to display the value
 of STATIME.
 - If the SET SYSTEM command is issued to change the value of the STATIME system parameter.
 - · When the CHINIT is shut down.
 - If the STOP TRACE(A) CLASS(4) is issued, any accounting data is written out.
- 6. SMF might hold the data in memory before writing it out to the SMF data sets or the SMF structure. Issue the MVS command **D SMF, 0** and note the MAXDORM value. SMF can keep the data in memory for the MAXDORM period before writing it out.

Related tasks

Planning for channel initiator SMF data
Interpreting IBM MQ performance statistics

Troubleshooting CHINIT statistics data

Checks to carry out if CHINIT statistics SMF data is not being produced.

Procedure

- 1. Issue the display trace command. You need TRACE(S) CLASS(4) enabled for information about the CHINIT.
- 2. If the trace is enabled, SMF data is written:
 - On a timed interval depending on the value of the STATIME system parameter. A value of zero
 means that the SMF statistics broadcast is used. Use the DIS SYSTEM command to display the value
 of STATIME.

- If the SET SYSTEM command is issued to change the value of the STATIME system parameter.
- · When the CHINIT is shut down.
- If the STOP TRACE(S) CLASS(4) is issued, any statistics data is written out.
- 3. SMF can hold the data in memory before writing it out to the SMF data sets or the SMF structure. Issue the MVS command **D SMF**, **O** and note the MAXDORM value. SMF can keep the data in memory for the MAXDORM period before writing it out.

Problem determination in DQM

Aspects of problem determination relating to distributed queue management (DQM) and suggested methods of resolving problems.

Some of the problems that are described are platform and installation specific. Where this is the case, it is made clear in the text.

IBM 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.

Problem determination for the following scenarios is discussed:

- "Error message from channel control" on page 151
- "Ping" on page 151
- "Dead-letter queue considerations" on page 151
- "Validation checks" on page 152
- "In-doubt relationship" on page 152
- "Channel startup negotiation errors" on page 152
- "When a channel refuses to run" on page 153
- "Retrying the link" on page 155
- "Data structures" on page 156
- "User exit problems" on page 156
- "Disaster recovery" on page 156
- "Channel switching" on page 156
- "Connection switching" on page 157
- "Client problems" on page 157
- "Error logs" on page 157
- "Message monitoring" on page 158

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

Related tasks

Configuring distributed queuing

"Making initial checks on UNIX, Linux, and Windows" on page 9

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

"Making initial checks on z/OS" on page 27

Before you start problem determination in detail on z/OS, 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.

"Making initial checks on IBM i" on page 18

Before you start problem determination in detail on IBM i, 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.

Related reference

Messages and reason codes
Communications protocol return codes

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 IBM 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 could be difficult in a network where the problem may 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 will indicate a problem originating from the remote site, but may 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.

Ping

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 makes no use of 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 z/OS z/OS and i5/OS , you can also use the panel interface to select this option.

On UNIX, is is is/OS, and Windows, you can also use the MQSC command PING QMGR to test whether the queue manager is responsive to commands.

Dead-letter queue considerations

In some IBM MQ implementations the dead-letter queue is referred to as an undelivered-message queue.

If a channel ceases to run for any reason, applications will 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 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 application's PUT 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 big 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 may occur when:

- 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 usually resolved automatically on restart, so the system operator does not need to resolve a channel manually in normal circumstances. See <u>In-doubt channels</u> for further information.

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 may 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.

Type of failure:	What happens:
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 refuses to run

If a channel refuses to run, there are a number of potential reasons.

Carry out the following checks:

- Check that DQM and the channels have been set up correctly. This 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 is indicated by messages on the system console, and the status panel may 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-committed unit of work ID (LUWID) committed. Check this 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 should be used with care. Its purpose is to reset the sequence number of messages and you should use it only after using the RESOLVE command to resolve any in-doubt situations.

- Windows Z/OS IBM i UNIX When sequential numbering is being used, and a sender channel starts up after being reset, the sender channel takes two actions:
- It tells the receiver channel that it has been reset.
- It specifies the next message sequence number that is to be used by both the sender and receiver channels.
- If the status of a receiver end of the channel is STOPPED, it can be reset by starting the receiver end.

Note: This does not start the channel, it merely resets the status. The channel must still be started from the sender end.

Triggered channels

If a triggered channel refuses to run, investigate the possibility of in-doubt messages here: <u>"When a channel refuses to run"</u> on page 153

Another possibility is that the trigger control parameter on the transmission queue has been set to NOTRIGGER by the channel. This 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 happens, restart the channel manually.

On IBM 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_KEEPALIVE option to set a system-wide value.

- On IBM MQ for z/OS, you also have the following options:
- Use the Keepalive Interval channel attribute (KAINT) to set channel-specific keepalive values.
- Use the RCVTIME and RCVTMIN channel initiator parameters.

These options are discussed in Checking that the other end of the channel is still available, and Keepalive Interval (KAINT).

Adopting an MCA

The Adopt MCA function enables IBM 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.

Registration time for DDNS

When a group TCP/IP listener is started, it registers with DDNS. But there can 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 is dependent on the name server configuration used.

Dial-up problems

IBM MQ supports connection over dial-up lines but you should be aware that with TCP, some protocol providers assign a new IP address each time you dial in. This 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 an IBM MQ for IBM i, UNIX, or Windows 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 may occur that is difficult to recognize. For example, the link and channel may 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 need to be aware that such situations can arise, often characterized by a system that appears to be busy but is not actually moving messages. You need to 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 IBM i, UNIX, and Windows, to attempt a retry a channel initiator must be running. In platforms other than IBM MQ for IBM i, UNIX, 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 <u>"Shared channel recovery"</u> on page 152, which includes 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 <u>Channel-exit calls and data structures</u> and <u>Developing applications</u> 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 <u>Channel-exit programs for messaging channels</u>. 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 may include the provision of regular system 'snapshot' dumps that are stored safely off-site. These dumps would be available for regenerating the system, should some disaster 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 may have a system log that allows the applications running at the time of failure to be regenerated by replaying the system software from a syncpoint forward to the instant of failure. If this occurs without error, the worst that can happen is that message channel syncpoints to the adjacent system may fail on startup, and that the last batches of messages for the various channels will be sent again. Persistent messages will be recovered and sent again, nonpersistent messages may 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 may be recovered to an earlier state, and the messages held on local queues at the sending and receiving end of channels may be inconsistent.

Messages may have been lost that were put on local queues. The consequence of this happening depends on the particular IBM 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 may need to restart from an earlier message sequence number.

Channel switching

A possible solution to the problem of a channel ceasing to run would be to have two message channels defined for the same transmission queue, but with different communication links. One message channel would be preferred, the other would be 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 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 may omit this step as the queue manager will find 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 would be to switch communication connections from the transmission queues.

To do this:

- If the sender channel is triggered, set the transmission queue attribute NOTRIGGER.
- Ensure 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 may 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 may 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

IBM MQ error messages are placed in different error logs depending on the platform. There are error logs for:

- Windows
- UNIX
- z/OS

Windows Error logs for Windows

IBM MQ for Windows uses a number of error logs to capture messages concerning the operation of IBM 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:

MQ_INSTALLATION_PATH\QMGRS\QMgrName\ERRORS\AMQERR01.LOG

• If the queue manager is not available:

MQ_INSTALLATION_PATH\QMGRS\@SYSTEM\ERRORS\AMQERR01.LOG

• If an error has occurred with a client application:

MQ_INSTALLATION_PATH\ERRORS\AMQERR01.LOG

On Windows, you should also examine the Windows application event log for relevant messages.

Error logs on UNIX and Linux systems

IBM MQ on UNIX and Linux systems uses a number of error logs to capture messages concerning the operation of IBM 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 TLS 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 might be suppressed. See Planning your IBM MQ environment on z/OS.

Message monitoring

If a message does not reach its intended destination, you can use the IBM 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 IBM MQ display route application is described in the IBM MQ display route application section.

Channel authentication records troubleshooting

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('*').

Behavior 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

z/OS

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. 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-

- USERSRC(MAP) for all addresses (*)

Commands troubleshooting

Troubleshooting advice for errors that appear when you use special characters in descriptive text.

• 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.

Distributed publish/subscribe troubleshooting

Use the advice given in the subtopics to help you to detect and deal with problems when you use publish/subscribe clusters or hierarchies.

Before you begin

If your problems relate to clustering in general, rather than to publish/subscribe messaging using clusters, see "Queue manager clusters troubleshooting" on page 193.

There are also some helpful troubleshooting tips in <u>Design considerations for retained publications in</u> publish/subscribe clusters.

Related concepts

Distributed publish/subscribe system queue errors

Related tasks

Configuring a publish/subscribe cluster

Designing publish/subscribe clusters

Routing for publish/subscribe clusters: Notes on behavior

Use the advice given here to help you to detect and deal with routing problems when you are using clustered publish/subscribe messaging.

For information about status checking and troubleshooting for any queue manager cluster, see <u>"Queue manager clusters troubleshooting"</u> on page 193.

 All clustered definitions of the same named topic object in a cluster must have the same CLROUTE setting. You can check the CLROUTE setting for all topics on all hosts in the cluster using the following MQSC command:

display tcluster(*) clroute

- The **CLROUTE** property has no effect unless the topic object specifies a value for the **CLUSTER** property.
- Check that you have spelled the cluster name correctly on your topic. You can define a cluster object such as a topic before defining the cluster. Therefore, when you define a cluster topic, no validation is done on the cluster name because it might not yet exist. Consequently, the product does not alert you to misspelt cluster names.
- When you set the CLROUTE property, if the queue manager knows of a clustered definition of the same object from another queue manager that has a different CLROUTE setting, the system generates an MQRCCF_CLUSTER_TOPIC_CONFLICT exception. However, through near simultaneous object definition on different queue managers, or erratic connectivity with full repositories, differing definitions might be created. In this situation the full repository queue managers arbitrate, accepting one definition and reporting an error for the other one. To get more information about the conflict, use the following MQSC command to check the cluster state of all topics on all queue managers in the cluster:

display tcluster(*) clstate

A state of invalid, or pending (if this does not soon turn to active), indicates a problem. If an invalid topic definition is detected, identify the incorrect topic definition and remove it from the cluster. The full repositories have information about which definition was accepted and which was rejected, and the queue managers that created the conflict have some indication of the nature of the problem. See also CLSTATE in DISPLAY TOPIC.

- Setting the **CLROUTE** parameter at a point in the topic tree causes the entire branch beneath it to route topics in that way. You cannot change the routing behavior of a sub-branch of this branch. For this reason, defining a topic object for a lower or higher node in the topic tree with a different **CLROUTE** setting is rejected with an MQRCCF_CLUSTER_TOPIC_CONFLICT exception.
- You can use the following MQSC command to check the topic status of all the topics in the topic tree:

```
display tpstatus('#')
```

If you have a large number of branches in the topic tree, the previous command might display status for an inconveniently large number of topics. If that is the case, you can instead display a manageably small branch of the tree, or an individual topic in the tree. The information displayed includes the topic string, cluster name and cluster route setting. It also includes the publisher count and subscription count (number of publishers and subscribers), to help you judge whether the number of users of this topic is as you expect.

- Changing the cluster routing of a topic in a cluster is a significant change to the publish/subscribe topology. After a topic object has been clustered (through setting the CLUSTER property) you cannot change the value of the CLROUTE property. The object must be un-clustered (CLUSTER set to ' ') before you can change the value. Un-clustering a topic converts the topic definition to a local topic, which results in a period during which publications are not delivered to subscriptions on remote queue managers; this should be considered when performing this change. See The effect of defining a non-cluster topic with the same name as a cluster topic from another queue manager. If you try to change the value of the CLROUTE property while it is clustered, the system generates an MQRCCF_CLROUTE_NOT_ALTERABLE exception.
- For topic host routing, you can explore alternative routes through the cluster by adding and removing the same cluster topic definition on a range of cluster queue managers. To stop a given queue manager from acting as a topic host for your cluster topic, either delete the topic object, or use the PUB (DISABLED) setting to quiesce message traffic for this topic, as discussed in Special handling for the PUB parameter. Do not un-cluster the topic by setting the CLUSTER property to ' ', because removing the cluster name converts the topic definition to a local topic, and prevents the clustering behavior of the topic when used from this queue manager. See The effect of defining a non-cluster topic with the same name as a cluster topic from another queue manager.
- You cannot change the cluster of a sub-branch of the topic tree when the branch has already been
 clustered to a different cluster and CLROUTE is set to TOPICHOST. If such a definition is detected
 at define time, the system generates an MQRCCF_CLUSTER_TOPIC_CONFLICT exception. Similarly,
 inserting a newly clustered topic definition at a higher node for a different cluster generates an
 exception. Because of the clustering timing issues previously described, if such an inconsistency is
 later detected, the queue manager issues errors to the queue manager log.

Related tasks

Configuring a publish/subscribe cluster
Designing publish/subscribe clusters

Checking proxy subscription locations

A proxy subscription enables a publication to flow to a subscriber on a remote queue manager. If your subscribers are not getting messages that are published elsewhere in the queue manager network, check that your proxy subscriptions are where you expect them to be.

Missing proxy subscriptions can show that your application is not subscribing on the correct topic object or topic string, or that there is a problem with the topic definition, or that a channel is not running or is not configured correctly.

To show proxy subscriptions, use the following MQSC command:

```
display sub(*) subtype(proxy)
```

Proxy subscriptions are used in all distributed publish/subscribe topologies (hierarchies and clusters). For a topic host routed cluster topic, a proxy subscription exists on each topic hosts for that topic.

For a direct routed cluster topic, the proxy subscription exists on every queue manager in the cluster. Proxy subscriptions can also be made to exist on every queue manager in the network by setting the proxysub(force) attribute on a topic.

See also Subscription performance in publish/subscribe networks.

Resynchronization of proxy subscriptions

Under normal circumstances, queue managers automatically ensure that the proxy subscriptions in the system correctly reflect the subscriptions on each queue manager in the network. Should the need arise, you can manually resynchronize a queue manager's local subscriptions with the proxy subscriptions that it propagated across the network using the **REFRESH QMGR TYPE(PROXYSUB)** command. However, you should do so only in exceptional circumstances.

When to manually resynchronize proxy subscriptions

When a queue manager is receiving subscriptions that it should not be sent, or not receiving subscriptions that it should receive, you should consider manually resynchronizing the proxy subscriptions. However, resynchronization temporarily creates a sudden additional proxy subscription load on the network, originating from the queue manager where the command is issued. For this reason, do not manually resynchronize unless IBM MQ service, IBM MQ documentation, or error logging instructs you to do so.

You do not need to manually resynchronize proxy subscriptions if automatic revalidation by the queue manager is about to occur. Typically, a queue manager revalidates proxy subscriptions with affected directly-connected queue managers at the following times:

- When forming a hierarchical connection
- When modifying the PUBSCOPE or SUBSCOPE or CLUSTER attributes on a topic object
- When restarting the queue manager

Sometimes a configuration error results in missing or extraneous proxy subscriptions:

- Missing proxy subscriptions can be caused if the closest matching topic definition is specified with Subscription scope set to Queue Manager or with an empty or incorrect cluster name. Note that Publication scope does not prevent the sending of proxy subscriptions, but does prevent publications from being delivered to them.
- Extraneous proxy subscriptions can be caused if the closest matching topic definition is specified with **Proxy subscription behavior** set to Force.

When configuration errors cause these problems, manual resynchronization does not resolve them. In these cases, amend the configuration.

The following list describes the exceptional situations in which you should manually resynchronize proxy subscriptions:

- After issuing a **REFRESH CLUSTER** command on a queue manager in a publish/subscribe cluster.
- When messages in the queue manager error log tell you to run the **REFRESH QMGR TYPE(REPOS)** command.
- When a queue manager cannot correctly propagate its proxy subscriptions, perhaps because a channel has stopped and all messages cannot be queued for transmission, or because operator error has caused messages to be incorrectly deleted from the SYSTEM.CLUSTER.TRANSMIT.QUEUE queue.
- When messages are incorrectly deleted from other system queues.
- When a **DELETE SUB** command is issued in error on a proxy subscription.
- · As part of disaster recovery.

How to manually resynchronize proxy subscriptions

First rectify the original problem (for example by restarting the channel), then issue the following command on the queue manager:

REFRESH OMGR TYPE (PROXYSUB)

When you issue this command, the queue manager sends, to each of its directly-connected queue managers, a list of its own topic strings for which proxy subscriptions should exist. The directly-connected queue managers then update their held proxy subscriptions to match the list. Next, the directly-connected queue managers send back to the originating queue manager a list of their own topic strings for which proxy subscriptions should exist, and the originating queue manager updates its held proxy subscriptions accordingly.

Important usage notes:

- Publications missed due to proxy subscriptions not being in place are not recovered for the affected subscriptions.
- Resynchronization requires the queue manager to start channels to other queue managers. If you are
 using direct routing in a cluster, or you are using topic host routing and this command is issued on a
 topic host queue manager, the queue manager will start channels to all other queue managers in the
 cluster, even those that have not performed publish/subscribe work. Therefore the queue manager that
 you are refreshing must have enough capability to cope with communicating with every other queue
 manager in the cluster.
- If this command is issued on z/OS when the CHINIT is not running, the command is queued up and processed when the CHINIT starts.

Related concepts

REFRESH CLUSTER considerations for publish/subscribe clusters

Related tasks

Checking that async commands for distributed networks have finished

Loop detection in a distributed publish/subscribe network

In a distributed publish/subscribe network, it is important that publications and proxy subscriptions cannot loop, because this would result in a flooded network with connected subscribers receiving multiple copies of the same original publication.

The proxy subscription aggregation system described in <u>Proxy subscriptions in a publish/subscribe</u> <u>network</u> does not prevent the formation of a loop, although it will prevent the perpetual looping of proxy subscriptions. Because the propagation of publications is determined by the existence of proxy subscriptions, they can enter a perpetual loop. IBM MQ uses the following technique to prevent publications from perpetually looping:

As publications move around a publish/subscribe topology each queue manager adds a unique fingerprint to the message header. Whenever a publish/subscribe queue manager receives a publication from another publish/subscribe queue manager, the fingerprints held in the message header are checked. If its own fingerprint is already present, the publication has fully circulated around a loop, so the queue manager discards the message, and adds an entry to the error log.

Note: Within a loop, publications are propagated in both directions around the loop, and each queue manager within the loop receives both publications before the originating queue manager discards the looped publications. This results in subscribing applications receiving duplicate copies of publications until the loop is broken.

Loop detection fingerprint format

The loop detection fingerprints are inserted into an RFH2 header or flow as part of the 8.0 protocol. An RFH2 programmer needs to understand the header and pass on the fingerprint information intact. earlier versions of IBM Integration Bus use RFH1 headers which do not contain the fingerprint information.

```
<ibm>
<Rfp>uuid1</Rfp>
<Rfp>uuid2</Rfp>
<Rfp>uuid3</Rfp>
```

</ibm>

<ibm> is the name of the folder that holds the list of routing fingerprints containing the unique user identifier (uuid) of each queue manager that has been visited.

Every time that a message is published by a queue manager, it adds its uuid into the <ibm> folder using the <Rfp> (routing fingerprint) tag. Whenever a publication is received, IBM MQ uses the message properties API to iterate through the <Rfp> tags to see if that particular uuid value is present. Because of the way that the WebSphere Platform Messaging component of IBM MQ attaches to IBM Integration Bus through a channel and RFH2 subscription when using the queued publish/subscribe interface, IBM MQ also creates a fingerprint when it receives a publication by that route.

The goal is to not deliver any RFH2 to an application if it is not expecting any, simply because we have added in our fingerprint information.

Whenever an RFH2 is converted into message properties, it will also be necessary to convert the <ibm> folder; this removes the fingerprint information from the RFH2 that is passed on or delivered to applications that have used the IBM MQ 7.0, or later, API.

JMS applications do not see the fingerprint information, because the JMS interface does not extract that information from the RFH2, and therefore does not hand it on to its applications.

The Rfp message properties are created with propDesc.CopyOptions = MQCOPY_FORWARD and MQCOPY_PUBLISH. This has implications for applications receiving and then republishing the same message. It means that such an application can continue the chain of routing fingerprints by using PutMsgOpts.Action = MQACTP_FORWARD, but must be coded appropriately to remove its own fingerprint from the chain. By default the application uses PutMsgOpts.Action = MQACTP_NEW and starts a new chain.

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.

About this task

The subtopics in this section provide advice to help you to detect and deal with problems that you might encounter under the following circumstances:

- When using the IBM MQ resource adapter
- When connecting to a queue manager with a specified provider version

Related concepts

"Tracing IBM MQ classes for JMS applications" on page 87

The trace facility in IBM MQ classes for JMS is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

"Tracing the IBM MQ resource adapter" on page 95

The ResourceAdapter object encapsulates the global properties of the IBM MQ resource adapter. To enable trace of the IBM MQ resource adapter, properties need to be defined in the ResourceAdapter object.

"Tracing additional IBM MQ Java components" on page 97

For Java components of IBM MQ, for example the IBM MQ Explorer and the Java implementation of IBM MQ Transport for SOAP, diagnostic information is output using the standard IBM MQ diagnostic facilities or by Java diagnostic classes.

Related tasks

"Tracing IBM MQ classes for Java applications" on page 91

The trace facility in IBM MQ classes for Java is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

Using IBM MQ classes for JMS
Using the IBM MQ resource adapter
Using IBM MQ classes for Java

Troubleshooting IBM MQ classes for JMS problems

You can investigate problems by running the installation verification programs, and by using the trace and log facilities.

If a program does not complete successfully, run one of the installation verification programs, as described in The point-to-point IVT for IBM MQ classes for JMS and The publish/subscribe IVT for IBM MQ classes for JMS, and follow the advice given in the diagnostic messages.

Related concepts

"Tracing IBM MQ classes for JMS applications" on page 87

The trace facility in IBM MQ classes for JMS is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

Logging errors for IBM MQ classes for JMS

By default, log output is sent to the mqjms.log file. You can redirect it to a specific file or directory.

The IBM MQ classes for JMS log facility is provided to report serious problems, particularly problems that might indicate configuration errors rather than programming errors. By default, log output is sent to the mq jms.log file in the JVM working directory.

You can redirect log output to another file by setting the property com.ibm.msg.client.commonservices.log.outputName. The value for this property can be:

- A single path name.
- A comma-separated list of path names (all data is logged to all files).

Each path name can be:

- Absolute or relative.
- stderr or System.err to represent the standard error stream.
- sttdout or System.out to represent the standard output stream.

If the value of the property identifies a directory, log output is written to mqjms.log in that directory. If the value of the property identifies a specific file, log output is written to that file.

You can set this property in the IBM MQ classes for JMS configuration file or as a system property on the **java** command. In the following example, the property is set as a system property and identifies a specific file:

```
java -Djava.library.path= library_path
-Dcom.ibm.msg.client.commonservices.log.outputName=/mydir/mylog.txt
MyAppClass
```

In the command, *library_path* is the path to the directory containing the IBM MQ classes for JMS libraries (see Configuring the Java Native Interface (JNI) libraries).

You can disable log output by setting the property com.ibm.msg.client.commonservices.log.status to OFF. The default value of this property is ON.

The values System.err and System.out can be set to send log output to the System.err and System.out streams.

JMS provider version troubleshooting

Use the advice that is given here to help you to resolve common problems that can arise when you are connecting to a queue manager with a specified provider version.

JMS 2.0 function is not supported with this connection error

- Error code: JMSCC5008
- Scenario: You receive a JMS 2.0 function is not supported with this connection error.
- **Explanation:** The use of the JMS 2.0 functionality is only supported when connecting to an IBM MQ 8.0 or later queue manager that is using IBM MQ messaging provider Version 8 mode.
- **Solution:** Change the application to not use the JMS 2.0 function, or ensure that the application connects to an IBM MQ 8.0 queue manager that is using IBM MQ messaging provider Version 8 mode.

JMS 2.0 API is not supported with this connection error

- Error code: JMSCC5007
- Scenario: You receive aJMS 2.0 API is not supported with this connection error.
- Explanation: The use of the JMS 2.0 API is only supported when you are connecting to an IBM WebSphere MQ 7 or 8 queue manager that is using IBM MQ messaging provider Normal or Version 8 mode. You might, for example, receive this error if you are attempting to connect to an IBM WebSphere MQ 6 queue manager or if you are connecting by using migration mode. This typically happens if SHARECNV(0) or PROVIDER_VERSION=6 is specified.
- **Solution:** Change the application to not use the JMS 2.0 API, or ensure that the application connects to an IBM WebSphere MQ 7 or 8 queue manager by using IBM MQ messaging provider Normal or Version 8 mode.

Queue manager command level did not match the requested provider version error

- Error code: JMSFMQ0003
- Scenario: You receive a queue manager command level did not match the requested provider versionerror.
- **Explanation:** The queue manager version that is specified in the provider version property on the connection factory is not compatible with the requested queue manager. For example, you might have specified PROVIDER_VERSION=8, and attempted to connect to a queue manager with a command level less than 800, such as 750.
- **Solution:** Modify the connection factory to connect to a queue manager that can support the provider version required.

For more information about provider version, see Configuring the JMS **PROVIDERVERSION** property.

PCF processing in JMS

IBM MQ Programmable Change Format (PCF) messages are a flexible, powerful way in which to query and modify attributes of a queue manager, and the PCF classes that are provided in the IBM MQ classes for Java provide a convenient way of accessing their functionality in a Java application. The functionality can also be accessed from IBM MQ classes for JMS, but there is a potential problem.

The common model for processing PCF responses in JMS

A common approach to processing PCF responses in JMS is to extract the bytes payload of the message, wrap it in a DataInputStream and pass it to the com.ibm.mq.headers.pcf.PCFMessage constructor.

Message m = consumer.receive(10000);
//Reconstitute the PCF response.

```
ByteArrayInputStream bais =
   new ByteArrayInputStream(((BytesMessage)m).getBody(byte[].class));
DataInput di = new DataInputStream(bais);
PCFMessage pcfResponseMessage = new PCFMessage(di);
```

See Using the IBM MQ Headers package for some examples.

Unfortunately this is not a totally reliable approach for all platforms - in general the approach works for big-endian platforms, but not for little-endian platforms.

What is the problem?

The problem is that in parsing the message headers, the PCFMessage class must deal with issues of numeric encoding - the headers contain length fields that are in some encoding that is big-endian or little-endian.

If you pass a pure DataInputStream to the constructor, the PCFMessage class has no good indication of the encoding, and must assume a default, quite possibly incorrectly.

If this situation arises, you will probably see a "MQRCCF_STRUCTURE_TYPE_ERROR" (reason code 3013) in the constructor:

```
com.ibm.mq.headers.MQDataException: MQJE001: Completion Code '2', Reason '3013'.
    at com.ibm.mq.headers.pcf.PCFParameter.nextParameter(PCFParameter.java:167)
    at com.ibm.mq.headers.pcf.PCFMessage.initialize(PCFMessage.java:854)
    at com.ibm.mq.headers.pcf.PCFMessage.
```

This message almost invariably means that the encoding has been misinterpreted. The probable reason for this is that the data that has been read is little-endian data which has been interpreted as big-endian.

The solution

The way to avoid this problem is to pass the PCFMessage constructor something that tells the constructor the numeric encoding of the data it is working with.

To do this, make an MQMessage from the data received.

The following code is an outline example of the code you might use.



Attention: The code is an outline example only and does not contain any error handling information.

```
// get a response into a JMS Message
Message receivedMessage = consumer.receive(10000);
BytesMessage bytesMessage = (BytesMessage) receivedMessage;
byte[] bytesreceived = new byte[(int) bytesMessage.getBodyLength()];
bytesMessage.readBytes(bytesreceived);

// convert to MQMessage then to PCFMessage
MQMessage mqMsg = new MQMessage();
mqMsg.write(bytesreceived);
mqMsg.write(bytesreceived);
mqMsg.encoding = receivedMessage.getIntProperty("JMS_IBM_Encoding");
mqMsg.format = receivedMessage.getStringProperty("JMS_IBM_Format");
mqMsg.seek(0);

PCFMessage pcfMsg = new PCFMessage(mqMsg);
```

JMS connection pool error handling

Connection pool error handling is carried out by various methods of a purge policy.

The connection pool purge policy comes into operation if an error is detected when an application is using a JMS connection to a JMS provider. The connection manager can either:

• Close only the connection that encountered the problem. This is known as the FailingConnectionOnly purge policy and is the default behavior.

Any other connections created from the factory, that is, those in use by other applications, and those that are in the free pool of the factory, are left alone.

• Close the connection that encountered the problem, throw away any connections in the free pool of the factory, and mark any in-use connections as stale.

The next time the application that is using the connection tries to perform a connection-based operation, the application receives a StaleConnectionException. For this behavior, set the purge policy to Entire Pool.

Purge policy - failing connection only

Use the example described in <u>How MDB listener ports use the connection pool</u>. Two MDBs are deployed into the application server, each one using a different listener port. The listener ports both use the jms/CF1 connection factory.

After 600 seconds, you stop the first listener, and the connection that this listener port was using is returned to the connection pool.

If the second listener encounters a network error while polling the JMS destination, the listener port shuts down. Because the purge policy for the jms/CF1 connection factory is set to FailingConnectionOnly, the connection manager throws away only the connection that was used by the second listener. The connection in the free pool remains where it is.

If you now restart the second listener, the connection manager passes the connection from the free pool to the listener.

Purge policy - entire pool

For this situation, assume that you have three MDBs installed into your application server, each one using its own listener port. The listener ports have created connections from the jms/CF1 factory. After a period of time you stop the first listener, and its connection, c1, is put into the jms/CF1 free pool.

When the second listener detects a network error, it shuts itself down and closes c2. The connection manager now closes the connection in the free pool. However, the connection being used by third listener remains.

What should you set the purge policy to?

As previously stated, the default value of the purge policy for JMS connection pools is FailingConnectionOnly.

However, setting the purge policy to EntirePool is a better option. In most cases, if an application detects a network error on its connection to the JMS provider, it is likely that all open connections created from the same connection factory have the same problem.

If the purge policy is set to FailingConnectionOnly, the connection manager leaves all of the connections in the free pool. The next time an application tries to create a connection to the JMS provider, the connection manager returns one from the free pool if there is one available. However, when the application tries to use the connection, it encounters the same network problem as the first application.

Now, consider the same situation with the purge policy set to EntirePool. As soon as the first application encounters the network problem, the connection manager discards the failing connection and closes all connections in the free pool for that factory.

When a new application starts up and tries to create a connection from the factory, the connection manager tries to create a new one, as the free pool is empty. Assuming that the network problem has been resolved, the connection returned to the application is valid.

Connection pool errors while trying to create a JMS Context

If an error occurs while you are trying to create a JMS Context, it is possible to determine from the error message if the top-level pool or lower-level pool had the issue.

How pools are used for Contexts

When using Connection and Sessions, there are pools for each type of object; a similar model is followed for Contexts.

A typical application that uses distributed transactions involves both messaging and non-messaging workloads in the same transaction.

Assuming that no work is currently working, and the application makes its first createConnection method call, a context facade or proxy is created in the equivalent of the connection pool (the top-level pool). Another object is created in the equivalent of the session pool. This second object encapsulates the underlying JMS Context (lower-level pool).

Pooling, as a concept, is used to permit an application to scale. Many threads are able to access a constrained set of resources. In this example, another thread will execute the createContext method call to get a context from the pool. Should other threads still be doing messaging work, then the top-level pool is expanded to provide an additional context for the requesting thread.

In the case where a thread requests a context and the messaging work has completed but the non-messaging work has not, so the transaction is not complete, the lower-level pool is expanded. The top-level context proxy remains assigned to the transaction until that transaction is resolved, so cannot be assigned to another transaction.

In the case of the lower pool becoming full, this means that the non-messaging work is taking potentially a long time.

In the case of the top-level pool becoming full, this means that the overall messaging work is taking a while and the pool should be expanded.

Identifying which pool an error originated from

You can determine the pool in which an error originated from the error message text:

- For the top-level pool, the message text is Failed to create context. This message means that the top-level pool is full of Context-proxy objects, all of which have currently running transactions that are performing messaging.
- For the lower-level pool, the message text is Failed to set up new JMSContext. This message means that although a connect-proxy is available, it is still necessary to wait for non-messaging work to complete.

Top-level pool example

```
at com.ibm.mq.jms.MQSession.run(MQSession.java:937)
at com.ibm.mq.connector.inbound.ASFWorkImpl.doDelivery(ASFWorkImpl.java:104)
at com.ibm.mq.connector.inbound.AbstractWorkImpl.run(AbstractWorkImpl.java:233)
at com.ibm.ejs.j2c.work.WorkProxy.run(WorkProxy.java:668)
at com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1892)
Caused by: com.ibm.websphere.ce.j2c.ConnectionWaitTimeoutException:
CWTE_NORMAL_J2CA1009
at com.ibm.ejs.j2c.FreePool.createOrWaitForConnection(FreePool.java:1783)
at com.ibm.ejs.j2c.PoolManager.reserve(PoolManager.java:3896)
at com.ibm.ejs.j2c.PoolManager.reserve(PoolManager.java:3116)
at com.ibm.ejs.j2c.ConnectionManager.allocateMCWrapper(ConnectionManager.java:1548)
at com.ibm.ejs.j2c.ConnectionManager.allocateConnection(ConnectionManager.java:1031)
at
com.ibm.ejs.jms.JMSConnectionFactoryHandle.createContextInternal(JMSConnectionFactoryHandle.java:4
3)
... 12 more
```

Lower-level pool example

```
[8/19/16 9:44:44:754 UTC] 000000ac SibMessage W
                                                  [:] CWSJY0003W: MQJCA4004: Message delivery to
an MDB
     sib.test.svt.lite.mdb.xa.MDBProxyQueueReceiver_37ea5ce9@505d4b68
(BeanId(SibSVTLiteMDB#SibSVTLiteMDBXA_RecoveryEJB_undeployed.jar#QueueReceiver, null))' failed
with exception:
'nested exception is: javax.jms.JMSRuntimeException: Failed to set up new JMSContext'.
^C[root@username-instance-2 server1]# vi SystemOut.log
                  :com.ibm.ejs.j2c.work.WorkProxy.run(WorkProxy.java:668)
                  : com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1892)
    Caused by [1] --> Message : javax.jms.JMSRuntimeException: Failed to set up new
JMSContext
               Class : class javax.jms.JMSRuntimeException
               Stack
com.ibm.ejs.jms.JMSCMUtils.mapToJMSRuntimeException(JMSCMUtils.java:522)
com.ibm.ejs.jms.JMSContextHandle.setupInternalContext(JMSContextHandle.java:241)
com.ibm.ejs.jms.JMSManagedConnection.getConnection(JMSManagedConnection.java:783)
com.ibm.ejs.j2c.MCWrapper.getConnection(MCWrapper.java:2336)
com.ibm.ejs.j2c.ConnectionManager.allocateConnection(ConnectionManager.java:1064)
com.ibm.ejs.jms.JMSConnectionFactoryHandle.createContextInternal(JMSConnectionFactoryHandle.java:4
43)
sib.test.svt.lite.mdb.xa.SVTMDBBase.sendReplyMessage(SVTMDBBase.java:554)
sib.test.svt.lite.mdb.xa.QueueReceiverBean.onMessage(QueueReceiverBean.java:128)
sib.test.svt.lite.mdb.xa.MDBProxyQueueReceiver_37ea5ce9.onMessage(MDBProxyQueueReceiver_37ea5ce9.j
ava:-1)
com.ibm.mq.connector.inbound.MessageEndpointWrapper.onMessage(MessageEndpointWrapper.java:151)
com.ibm.mq.jms.MQSession$FacadeMessageListener.onMessage(MQSession.java:129)
com.ibm.msg.client.jms.internal.JmsSessionImpl.run(JmsSessionImpl.java:3236)
                   com.ibm.mq.jms.MQSession.run(MQSession.java:937)
com.ibm.mq.connector.inbound.ASFWorkImpl.doDelivery(ASFWorkImpl.java:104)
com.ibm.mq.connector.inbound.AbstractWorkImpl.run(AbstractWorkImpl.java:233)
                  : com.ibm.ejs.j2c.work.WorkProxy.run(WorkProxy.java:668)
: com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1892)
    Caused by [2]
                 --> Message : com.ibm.websphere.ce.j2c.ConnectionWaitTimeoutException:
CWTE_NORMAL_J2CA1009
               Class : class
\verb|com.ibm.websphere.ce.j2c.ConnectionWaitTimeoutException| \\
               Stack : com.ibm.ejs.j2c.FreePool.createOrWaitForConnection(FreePool.java:1783)
com.ibm.ejs.j2c.PoolManager.reserve(PoolManager.java:3840)
                  : com.ibm.ejs.j2c.PoolManager.reserve(PoolManager.java:3116)
com.ibm.ejs.j2c.ConnectionManager.allocateMCWrapper(ConnectionManager.java:1548)
com.ibm.ejs.j2c.ConnectionManager.allocateConnection(ConnectionManager.java:1031)
```

```
com.ibm.ejs.jms.JMSContextHandle.setupInternalContext(JMSContextHandle.java:222)
com.ibm.ejs.jms.JMSManagedConnection.getConnection(JMSManagedConnection.java:783)
com.ibm.ejs.j2c.MCWrapper.getConnection(MCWrapper.java:2336)
com.ibm.ejs.j2c.ConnectionManager.allocateConnection(ConnectionManager.java:1064)
com.ibm.ejs.jms.JMSConnectionFactoryHandle.createContextInternal(JMSConnectionFactoryHandle.java:4
43)
com.ibm.ejs.jms.JMSConnectionFactoryHandle.createContext(JMSConnectionFactoryHandle.java:335)
sib.test.svt.lite.mdb.xa.SVTMDBBase.sendReplyMessage(SVTMDBBase.java:554)
sib.test.svt.lite.mdb.xa.QueueReceiverBean.onMessage(QueueReceiverBean.java:128)
sib.test.svt.lite.mdb.xa.MDBProxyQueueReceiver_37ea5ce9.onMessage(MDBProxyQueueReceiver_37ea5ce9.j
ava:-1)
com.ibm.mq.connector.inbound.MessageEndpointWrapper.onMessage(MessageEndpointWrapper.java:151)
com.ibm.mq.jms.MQSession$FacadeMessageListener.onMessage(MQSession.java:129)
com.ibm.msg.client.jms.internal.JmsSessionImpl.run(JmsSessionImpl.java:3236)
                   : com.ibm.mq.jms.MQSession.run(MQSession.java:937)
com.ibm.mq.connector.inbound.ASFWorkImpl.doDelivery(ASFWorkImpl.java:104)
com.ibm.mq.connector.inbound.AbstractWorkImpl.run(AbstractWorkImpl.java:233)
                   : com.ibm.ejs.j2c.work.WorkProxy.run(WorkProxy.java:668)
: com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1892)
```

Troubleshooting JMSCC0108 messages

There are a number of steps that you can take to prevent a JMSCC0108 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:

JMSCC0108: The IBM 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 IBM 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 MQ messaging provider normal mode.
- A queue manager on any platform, using IBM MQ messaging provider normal mode with restrictions
- A queue manager running on z/OS, using IBM MQ messaging provider migration mode.

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), 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 IBM MQ for Multiplatforms, using IBM MQ messaging provider migration mode, 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.

CWSJY0003W warning messages in WebSphere Application Server SystemOut.log file

A CWSJY0003W warning message is logged in the WebSphere Application Server SystemOut.log file when an MDB processes JMS messages from IBM WebSphere MQ.

Symptom

CWSJY0003W: IBM WebSphere MQ classes for JMS attempted to get a message for delivery to a message listener, that had previously been marked using browse-with-mark, however, the message is not available.

Cause

Activation specifications, and listener ports running in Application Server Facilities (ASF) mode, are used to monitor queues or topics hosted on IBM WebSphere MQ queue managers. Initially messages are browsed on either the queue or topic. When a message is found, a new thread is started which destructively gets the message and passes the message to an instance of a message-driven bean application for processing.

When the message is browsed, the queue manager marks the message for a period of time, and effectively hides the message from other application server instances. The time period that the message

is marked for is determined by the queue manager attribute **MARKINT**, which by default is set to 5000 milliseconds (5 seconds). This means that, after an activation specification or listener port has browsed a message, the queue manager will wait for 5 seconds for the destructive get of the message to occur before allowing another application server instance to see that message and process it.

The following situation can occur:

- An activation specification running on Application Server 1 browses message A on a queue.
- The activation specification starts a new thread to process message A.
- An event occurs on Application Server 1, which means that message A is still on the queue after 5 seconds.
- An activation specification running on Application Server 2 now browses message A and starts a new thread to process message A.
- The new thread running on Application Server 2 destructively gets message A, and passes it to a message-driven bean instance.
- The thread running on Application Server 1 attempts to get message A, only to find that message A is no longer on the queue.
- At this point, Application Server 1 reports the CWSJY0003W message.

Resolving the problem

There are two ways that you can resolve this issue:

- Increase the value of queue manager attribute MARKINT to a higher value. The default value for MARKINT is 5000 milliseconds (5 seconds). Increasing this value gives an application server more time to destructively get a message after it is detected. Changing the MARKINT value affects all applications that connect to the queue manager, and browse messages before applications destructively get messages.
- Change the value to true for the
 com.ibm.msg.client.wmq.suppressBrowseMarkMessageWarning property in WebSphere
 Application Server to suppress the CWSJY0003W warning message. To set the variable in WebSphere
 Application Server, open the administrative console and navigate to Servers -> Application Servers ->
 Java and Process Management -> Process Definition -> Java Virtual Machine -> Custom Properties

```
Name = com.ibm.msg.client.wmq.suppressBrowseMarkMessageWarning
Value = true
```

Note: If an activation specification or listener port is connecting to IBM WebSphere MQ using IBM WebSphere MQ messaging provider migration mode the messages can be ignored. The design of this mode of operation means that this message can occur during normal operation.

Related concepts

-> New

Activation specifications

Listener ports running in Application Server Facilities (ASF) mode

Listener ports running in non Application Server Facilities (non-ASF) mode

Related information

Avoiding repeated delivery of browsed messages

ALTER QMGR

J2CA0027E messages containing the error The method 'xa_end' has failed with errorCode '100'

J2CA0027E messages appear in the WebSphere Application Server SystemOut.log containing the error The method 'xa_end' has failed with errorCode '100'.

Introduction

The following errors appear in the WebSphere Application Server SystemOut.log file when applications using the WebSphere Application Server IBM WebSphere MQ messaging provider try to commit a transaction:

```
J2CA0027E: An exception occurred while invoking end on an XA Resource Adapter from DataSource JMS_Connection_Factory, within transaction ID Transaction_Identifier: javax.transaction.xa.XAException: The method 'xa_end' has failed with errorCode '100'.

J2CA0027E: An exception occurred while invoking rollback on an XA Resource Adapter from DataSource JMS_Connection_Factory, within transaction ID Transaction_Identifier: javax.transaction.xa.XAException: The method 'xa_rollback' has failed with errorCode '-7'
```

Cause

The cause of these errors can be the result of a IBM WebSphere MQ messaging provider JMS connection being closed off by WebSphere Application Server because the aged timeout for the connection has expired.

JMS connections are created from a JMS connection factory. There is a connection pool associated with each connection factory, which is divided into two parts - the active pool and the free pool.

When an application closes off a JMS connection that it has been using, that connection is moved into the free pool of the connection pool for the connection factory unless the aged timeout for that connection has elapsed, in which case the connection is destroyed. If the JMS connection is still involved in an active transaction when it is destroyed, the application server flows an xa_end() to IBM WebSphere MQ, indicating that all of the transactional work on that connection has completed.

This causes issues if the JMS connection had been created inside a transactional message-driven bean that was using either an activation specification or a listener port to monitor a JMS Destination on a IBM WebSphere MQ queue manager.

In this situation, there is a single transaction that is using 2 connections to IBM WebSphere MQ:

- A connection which is used to get a message from IBM WebSphere MQ and deliver it to the messagedriven bean instance for processing.
- A connection that is created within the message-driven bean's onMessage() method.

If the second connection is closed by the message-driven bean, and then destroyed as a result of the aged timeout expiring, then an xa_end() is flown to IBM WebSphere MQ indicating that all of the transactional work has completed.

When the message-driven bean application finishes processing the message it has been given, the application server needs to complete the transaction. It does this by flowing xa_end() to all of the resources that were involved in the transaction, including IBM WebSphere MQ.

However, IBM WebSphere MQ has already received an xa_end() for this particular transaction, and so returns an XA_RBROLLBACK (100) error back to WebSphere Application Server, indicating that the transaction has ended and all of the work IBM WebSphere MQ has been rolled back. This causes the application server to report the following error:

```
J2CA0027E: An exception occurred while invoking end on an XA Resource Adapter from DataSource JMS_Connection_Factory, within transaction ID Transaction_Identifier: javax.transaction.xa.XAException: The method 'xa_end' has failed with errorCode '100'.
```

and then roll back the entire transaction by flowing xa_rollback() to all of the resources enlisted in the transaction. When the application server flows xa_rollback() to IBM WebSphere MQ, the following error occurs:

J2CA0027E: An exception occurred while invoking rollback on an XA Resource Adapter

from DataSource $JMS_Connection_Factory$, within transaction ID $Transaction_Identifier$: javax.transaction.xa.XAException: The method 'xa_rollback' has failed with errorCode '-7'.

Environment

Message-driven bean applications that use activation specifications or listener ports to monitor JMS Destinations hosted on a IBM WebSphere MQ queue manager, and then create a new connection to IBM WebSphere MQ using a JMS connection factory from within its onMessage() method, can be affected by this issue.

Resolving the problem

To resolve this issue, ensure that the JMS connection factory being used by the application has the connection pool property aged timeout set to zero. This will prevent JMS Connections being closed when they are returned to the free pool, and so ensures that any outstanding transactional work can be completed.

2035 MQRC_NOT_AUTHORIZED when connecting to IBM MQ from WebSphere Application Server

The 2035 MQRC_NOT_AUTHORIZED error can occur when an application connects to IBM WebSphere MQ from WebSphere Application Server.

This topic covers the most common reasons why an application that is running in WebSphere Application Server receives a 2035 MQRC_NOT_AUTHORIZED error when connecting to IBM MQ. Quick steps to work around the 2035 MQRC_NOT_AUTHORIZED errors during development are provided in the Resolving the problem section, as well as considerations for implementing security in production environments. A summary is also provided of behavior for outbound scenarios with container-managed and component-managed security and inbound behavior for listener ports and activation specifications.

The cause of the problem

The most common reasons for why the connection is refused by IBM MQ are described in the following list:

- The user identifier that is passed across the client connection from the application server to IBM MQ is either; not known on the server where the IBM MQ queue manager is running, is not authorized to connect to IBM MQ, or is longer than 12 characters and was truncated. There is more information about how this user identifier is obtained and passed over in *Diagnosing the problem*.
 - Windows For queue managers that are running on Windows, the following error might be seen in the IBM MQ error logs for this scenario: AMQ8075: Authorization failed because the SID for entity 'wasuser' cannot be obtained.
 - UNIX For UNIX, no entry in the IBM MQ error logs would be seen.
- The user identifier that is passed across the client connection from the application server to IBM MQ is a member of the *mqm* group on the server that hosts the IBM MQ queue manager and a channel authentication record (CHLAUTH) exists that blocks administrative access to the queue manager. IBM MQ configures a CHLAUTH record by default in IBM WebSphere MQ 7.1 and later that blocks all IBM MQ administrators from connecting as a client to the queue manager. The following error in the IBM MQ error logs would be seen for this scenario: AMQ9777: Channel was blocked.
- The presence of an Advanced Message Security security policy.

For the location of the IBM MQ error logs, see Error log directories.

Diagnosing the problem

To understand the cause of the 2035 MQRC_NOT_AUTHORIZED reason code, you must understand which user name and password is being used by IBM MQ to authorize the application server.

Note: The understanding that is provided in this topic is helpful for development environments, solving the security requirements of production environments usually requires one of the following approaches:

Mutual SSL/TLS authentication

IBM MQ provides features to authenticate a remotely connecting client using the digital certificate that is provided for the SSL/TLS connection.

A custom, or third party supplied, IBM MQ security exit

A security exit can be written for IBM MQ that performs user name and password authentication against a repository, such as the local operating system, an IBM MQ server, or an LDAP repository. When you use a security exit for authentication it is important that SSL/TLS transport security is still configured, to ensure that passwords are not sent in plain text.

MCA user ID configured on the server connection channel

If an MCA user ID configured on the server connection channel that the application server is using to connect, and no security exit or mapping channel authentication record is installed, then the MCA user ID overrides the user name that is provided by the application server. It is common practice for many customers to set an MCA user ID on every server connection channel and use mutual SSL/TLS authentication exclusively for authentication.

Default behavior when no credentials are supplied from the application server

If no credentials are supplied by the application on the **createConnection** call, and neither of the component managed or container managed security systems are configured, then WebSphere Application Server provides a blank user name to IBM MQ. This causes IBM MQ to authorize the client based on the user ID that the IBM MQ listener is running under. In most cases the user ID is mqm on UNIX or Linux systems and $MUSR_MQADMIN$ on Windows. As these users are administrative IBM MQ users, they are blocked by default in IBM WebSphere MQ 7.1 and later, with an AMQ9777 error logged in the error logs of the queue manager.

Container-managed security for outbound connections

The recommended way to configure the user name and password that is passed to IBM MQ by the application server for outbound connections, is to use container-managed security. Outbound connections are those created by using a connection factory, rather than a listener port or activation specification.

User names of 12 characters or less are passed to v by the application server. User names longer than 12 characters in length are truncated, either during authorization (on UNIX), or in the *MQMD* of messages that are sent. Container-managed security means that the deployment descriptor, or EJB 3.0 annotations, of the application declare a resource reference with authentication type set to Container. Then, when the application looks up the connection factory in JNDI, it does so indirectly through the resource reference. For example, an EJB 2.1 application would perform a JNDI lookup as follows, where jms/MyResourceRef is declared as a resource reference in the deployment descriptor:

```
ConnectionFactory myCF = (ConnectionFactory)ctx.lookup("java:comp/env/jms/MyResourceRef")
```

An EJB 3.0 application might declare an annotated object property on the bean as follows:

When the application is deployed by an administrator, they bind this authentication alias to an actual connection factory that has been created in JNDI, and assign it a J2C authentication alias on deployment. It is the user name and password that is contained in this authentication alias that is then passed to IBM MQ or JMS by the application server when the application connects. This approach puts the administrator in control of which user name and password is used by each application, and prevents a different application from looking up the connection factory in JNDI directly to connect with the same user name and password. A default container-managed authentication alias can be supplied on the configuration panels in the administrative console for IBM MQ connection factories. This default is only used in the case

that an application uses a resource reference that is configured for container-managed security, but the administrator has not bound it to an authentication alias during deployment.

Default component-managed authentication alias for outbound connection

For cases where it is impractical to change the application to use container-managed security, or to change it to supply a user name and password directly on the createConnection call, it is possible to supply a default. This default is called the component-managed authentication alias and cannot be configured in the administrative console (since WebSphere Application Server 7.0 when it was removed from the panels for IBM MQ connection factories). The following scripting samples show how to configure it using wsadmin:

• JACL

```
wsadmin>set cell [ $AdminConfig getid "/Cell:mycell" ]
mycell(cells/mycell|cell.xml#Cell_1)
wsadmin>$AdminTask listWMQConnectionFactories $cell
MyCF(cells/mycell|resources.xml#MQConnectionFactory_1247500675104)
wsadmin>$AdminTask modifyWMQConnectionFactory MyCF(cells/mycell|
resources.xml#MQConnectionFactory_1247500675104) { -componentAuthAlias myalias }
MyCF(cells/mycell|resources.xml#MQConnectionFactory_1247500675104)
```

Jython

```
wsadmin>cell = AdminConfig.getid("/Cell:mycell")
wsadmin>AdminTask.listWMQConnectionFactories(cell)
'MyCF(cells/mycell|resources.xml#MQConnectionFactory_1247500675104)'
wsadmin>AdminTask.modifyWMQConnectionFactory('MyCF(cells/mycell|resos
urces.xml#MQConnectionFactory_1247500675104)', "-componentAuthAlias myalias")
'MyCF(cells/mycell|resources.xml#MQConnectionFactory_1247500675104)'
```

Authentication alias for inbound MDB connections using an activation specification

For inbound connections that use an activation specification, an authentication alias can be specified by the administrator when the application is deployed, or a default authentication alias can be specified on the activation specification in the administrative console.

Authentication alias for inbound MDB connections using a listener port

For inbound connections that use a listener port, the value that is specified in the container-managed authentication alias setting of the connection factory is used. On z/OS, first the container-managed authentication alias is checked and used if set, then the component-managed authentication alias is checked and used it set.

Resolving the problem

The simplest steps to resolve the 2035 MQRC_NOT_AUTHORIZED errors in a development environment, where full transport security is not required, are as follows:

- Choose a user that you want WebSphere Application Server to be authenticated as. Typically the user chosen should have authority relevant to the context of the operations required by the application running in WebSphere Application Server and no more. For example, *mqm* or other super user is not appropriate.
- If this user is an IBM MQ administrative user, then relax the channel authentication record (CHLAUTH) security in IBM WebSphere MQ 7.1 or later so that administrative connections are not blocked on the server connection channel you want to use. An example MQSC command for a server connection channel called WAS.CLIENTS is, SET CHLAUTH('WAS.CLIENTS') TYPE(BLOCKUSER) USERLIST(ALLOWANY).
- Configure the server connection channel to set the MCA user ID (MCAUSER) to the user you are using. An example MQSC command to configure a server connection channel to use myuser as the MCA user ID is, ALTER CHL('WAS.CLIENTS') CHLTYPE(SVRCONN) MCAUSER('myuser').

Important extra considerations for production environments

For all production environments where transport security is required, SSL/TLS security must be configured between the application server and IBM MQ.

To configure SSL/TLS transport security, you must establish the appropriate trust between the IBM MQ queue manager and WebSphere Application Server. The application server initiates the SSL/TLS handshake and must always be configured to trust the certificate that is provided by the IBM MQ queue manager. If the application server is configured to send a certificate to the IBM MQ queue manager, then the queue manager must also be configured to trust it. If trust is not correctly configured on both sides, you will encounter 2393 MQRC_SSL_INITIALIZATION_ERROR reason code after SSL/TLS is enabled on the connection.

If you do not have a security exit that performs username and password authentication, then you should configure mutual SSL/TLS authentication on your server connection channel to cause the queue manager to require a trusted certificate is provided by the application server. To do this you set *SSL Authentication* to Required in IBM MQ Explorer or SSLCAUTH(REQUIRED) in MQSC.

If you do have a security exit that performs user name and password authentication that is installed in your IBM MQ server, then configure your application to supply a username and password for validation by that security exit. The details of how to configure the user name and password that is passed to IBM MQ by the application server are described previously in the *Diagnosing the problem* section.

All server connection channels that do not have SSL/TLS security should be disabled. Example MQSC commands to disable the *SYSTEM.DEF.SVRCONN* channel are provided as follows (assuming no user exists on the IBM MQ server called ('NOAUTH'), ALTER CHL(SYSTEM.DEF.SVRCONN) CHLTYPE(SVRCONN) MCAUSER('NOAUTH') STOP CHL(SYSTEM.DEF.SVRCONN).

For instructions to configure the private certificate and trust of an IBM MQ queue manager and to enable SSL security on a server connection channel, see <u>Configuring SSL</u> on queue managers and <u>Configuring SSL</u> channels.

For information about using SSL/TLS from WebSphere Application Server and whether the application server sends a certificate to IBM MQ for authentication, see the following information:

- To create or modify an SSL configuration to contain the appropriate SSL/TLS configuration for connection to IBM MQ, see <u>SSL configurations</u> in the WebSphere Application Server product documentation.
- It is required by IBM MQ that you must specify a matching CipherSpec on both ends of the connection. For more information about CipherSpecs and CipherSuites that can be used with IBM MQ, see CipherSuite and CipherSpec name mappings for connections to a WebSphere® MQ queue manager.
- For more information about enabling SSL/TLS on a client connect and choosing which SSL configuration
 to use, see WebSphere MQ messaging provider connection factory settings and WebSphere MQ
 messaging provider activation specification settings in the WebSphere Application Server product
 documentation.

Related reference

"Return code= 2035 MQRC_NOT_AUTHORIZED" on page 199

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 MQ JMS application, and opening a queue on a cluster. MQS_REPORT_NOAUTH and MQSAUTHERRORS can be used to further diagnose RC2035.

MQRC_NOT_AUTHORIZED

Problem determination for the IBM MQ resource adapter

When using the IBM 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 IBM MQ queue manager, the following exception might be thrown:

javax.jms.JMSException: MQJCA0001: An exception occurred in the JMS layer. See the linked exception for details.

Linked to this exception is a second exception:

javax.jms.JMSException: MQJMS2005: failed to create an MQQueueManager for 'localhost:ExampleQM'

This exception is thrown by IBM MQ classes for JMS and has a further linked exception:

 ${\tt com.ibm.mq.MQException}\colon {\tt MQJE001}\colon {\tt An MQException occurred}\colon {\tt Completion Code 2}, {\tt Reason 2059}$

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 <u>Configuration of</u> the IBM MQ resource adapter.

Configuration problems commonly occur in the following areas:

- Deploying the resource adapter
- Deploying MDBs
- Creating connections for outbound communication

Related tasks

Using the IBM MQ resource adapter

Problems in deploying the resource adapter

If the resource adapter fails to deploy, check that Java EE Connector Architecture (JCA) resources are configured correctly. If IBM MQ is already installed, check that the correct versions of the JCA and IBM 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 Java Naming Directory Interface (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 IBM MQ classes for JMS classes from JAR files in the class path. This type of failure can commonly occur on a system where IBM MQ is already installed. On such a system, the application server might find existing copies of the IBM MQ classes for JMS JAR files and load classes from them in preference to the classes supplied in the IBM MQ resource adapter RAR file.

Related concepts

What is installed for IBM MQ classes for JMS

Related tasks

Configuring the application server to use the latest resource adapter maintenance level

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.

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.

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 IBM 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 an IBM 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 ConnectionFactory object has a property called ccdtURL and a property called channel. These properties are mutually exclusive.
- On a TLS connection, the TLS-related properties, or the TLS-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 IBM MQ resource adapter.

Related tasks

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client

Related reference

Federal Information Processing Standards (FIPS) for UNIX, Linux, and Windows

Using IBM MQ connection property override

Connection property override allows you to change the details that are used by a client application to connect to a queue manager, without modifying the source code.

About this task

Sometimes, it is not possible to modify the source code for an application, for example, if the application is a legacy application and the source code is no longer available.

In this situation, if an application needs to specify different properties when it is connecting to a queue manager, or is required to connect to a different queue manager, then you can use the connection override functionality to specify the new connection details or queue manager name.

The connection property override is supported for two clients:

- IBM MQ classes for JMS
- IBM MQ classes for Java

You can override the properties that you want to change by defining them in a configuration file that is then read by the IBM MQ classes for JMS or IBM MQ classes for Java at startup.

When the connection override functionality is in use, all applications that are running inside the same Java runtime environment pick up and use the new property values. If multiple applications that are using either the IBM MQ classes for JMS or the IBM MQ classes for Java are running inside the same Java runtime environment, it is not possible to just override properties for individual applications.

Important: This functionality is only supported for situations where it is not possible to modify the source code for an application. It must not be used for applications where the source code is available and can be updated.

Related concepts

"Tracing IBM MO classes for JMS applications" on page 87

The trace facility in IBM MQ classes for JMS is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

Related tasks

"Tracing IBM MQ classes for Java applications" on page 91

The trace facility in IBM MQ classes for Java is provided to help IBM Support to diagnose customer issues. Various properties control the behavior of this facility.

Using IBM MQ classes for JMS

Using IBM MQ classes for Java

Using connection property override in IBM MQ classes for JMS

If a connection factory is created programmatically, and it is not possible to modify the source code for the application that creates it, then the connection override functionality can be used to change the properties that the connection factory uses when a connection is created. However, the use of the connection override functionality with connection factories defined in JNDI is not supported.

About this task

In the IBM MQ classes for JMS, details about how to connect to a queue manager are stored in a connection factory. Connection factories can either be defined administratively and stored in a JNDI repository, or created programmatically by an application by using Java API calls.

If an application creates a connection factory programmatically, and it is not possible to modify the source code for that application, the connection override functionality allows you to override the connection factory properties in the short term. In the long term, though, you must put plans in place to allow the connection factory used by the application to be modified without using the connection override functionality.

If the connection factory that is created programmatically by an application is defined to use a Client Channel Definition Table (CCDT), then the information in the CCDT is used in preference to the overridden properties. If the connection details that the application uses need to be changed, then a new version of the CCDT must be created and made available to the application.

The use of the connection override functionality with connection factories defined in JNDI is not supported. If an application uses a connection factory that is defined in JNDI, and the properties of that connection factory need to be changed, then the definition of the connection factory must be updated in JNDI. Although the connection override functionality is applied to these connection factories (and the

overridden properties take precedence over the properties in the connection factory definition that is looked up in JNDI), this use of the connection override functionality is not supported.

Important: The connection override functionality affects all of the applications that are running inside of a Java runtime environment, and applies to all of the connection factories used by those applications. It is not possible to just override properties for individual connection factories or applications.

When an application uses a connection factory to create a connection to a queue manager, the IBM MQ classes for JMS look at the properties that have been overridden and use those property values when creating the connection, rather than the values for the same properties in the connection factory.

For example, suppose a connection factory has been defined with the PORT property set to 1414. If the connection override functionality has been used to set the PORT property to 1420, then when the connection factory is used to create a connection, the IBM MQ classes for JMS use a value of 1420 for the PORT property, rather than 1414.

To modify any of the connection properties that are used when creating a JMS connection from a connection factory, the following steps need to be carried out:

- 1. Add the properties to be overridden to an IBM MQ classes for JMS configuration file.
- 2. Enable the connection override functionality.
- 3. Start the application, specifying the configuration file.

Procedure

- 1. Add the properties to be overridden to an IBM MQ classes for JMS configuration file.
 - a) Create a file containing the properties and values that need to be overridden in the standard Java properties format.
 - For details about how you create a properties file, see <u>The IBM MQ classes for JMS configuration</u> file.
 - b) To override a property, add an entry to the properties file.
 - Any IBM MQ classes for JMS connection factory property can be overridden. Add each required entry in the following format:

```
jmscf.property name=value
```

where *property name* is the JMS administration property name or XMSC constant for the property that needs to be overridden. For a list of connection factory properties, see <u>Properties</u> of IBM MQ classes for JMS objects.

For example, to set the name of the channel that an application should use to connect to a queue manager, you can add the following entry to the properties file:

```
jmscf.channel=MY.NEW.SVRCONN
```

2. Enable the connection override functionality.

To enable connection override, set the

com.ibm.msg.client.jms.overrideConnectionFactory property to be true so that the properties that are specified in the properties file are used to override the values that are specified in the application. You can either set the extra property as another property in the configuration file itself, or pass the property as a Java system property by using:

```
-Dcom.ibm.msg.client.jms.overrideConnectionFactory=true
```

3. Start the application, specifying the configuration file.

Pass the properties file that you created to the application at run time by setting the Java system property:

```
-Dcom.ibm.msg.client.config.location
```

Note that the location of the configuration file must be specified as a URI, for example:

Results

When the connection override functionality is enabled, the IBM MQ classes for JMS write an entry to the jms log whenever a connection is made. The information in the log shows the connection factory properties that were overridden when the connection was created, as shown in the following example entry:

```
Overriding ConnectionFactory properties:
Overriding property channel:
Original value = MY.OLD.SVRCONN
New value = MY.NEW.SVRCONN
```

Related tasks

"Using connection property override in IBM MQ classes for Java" on page 184

In the IBM MQ classes for Java, connection details are set as properties using a combination of different values. The connection override functionality can be used to override the connection details that an

"Overriding connection properties: example with IBM MQ classes for JMS" on page 186
This example shows how to override properties when you are using the IBM MQ classes for JMS.

Creating and configuring connection factories and destinations in an IBM MQ classes for JMS application Configuring connection factories and destinations in a JNDI namespace

Using connection property override in IBM MQ classes for Java

application uses if it is not possible to modify the source code for the application.

In the IBM MQ classes for Java, connection details are set as properties using a combination of different values. The connection override functionality can be used to override the connection details that an application uses if it is not possible to modify the source code for the application.

About this task

The different values that are used to set the connection properties are a combination of:

- Assigning values to static fields on the MQEnvironment class.
- Setting property values in the properties Hashtable in the MQEnvironment class.
- Setting property values in a Hashtable passed into an MQQueueManager constructor.

These properties are then used when an application constructs an MQQueueManager object, which represents a connection to a queue manager.

If it is not possible to modify the source code for an application that uses the IBM MQ classes for Java to specify different properties that must be used when creating a connection to a queue manager, the connection override functionality allows you to override the connection details in the short term. In the long term, though, you must put plans in place to allow the connection details used by the application to be modified without using the connection override functionality.

When an application creates an MQQueueManager, the IBM MQ classes for Java look at the properties that have been overridden and use those property values when creating a connection to the queue manager, rather than the values in any of the following locations:

- The static fields on the MQEnvironment class
- The properties Hashtable stored in the MQEnvironment class
- The properties Hashtable that is passed into an MOOueueManager constructor

For example, suppose an application creates an MQQueueManager, passing in a properties Hashtable that has the CHANNEL property set to MY.OLD.CHANNEL. If the connection override functionality has been used to set the CHANNEL property to MY.NEW.CHANNEL, then when the MQQueueManager is

constructed, the IBM MQ classes for Java attempt to create a connection to the queue manager by using the channel MY.NEW.CHANNEL rather than MY.OLD.CHANNEL.

Note: If an MQQueueManager is configured to use a Client Channel Definition Table (CCDT), then the information in the CCDT is used in preference to the overridden properties. If the connection details that the application creating the MQQueueManager uses need to be changed, then a new version of the CCDT must be created and made available to the application.

To modify any of the connection properties that are used when creating an MQQueueManager, the following steps need to be carried out:

- 1. Create a properties file called mqclassesforjava.config.
- 2. Enable the connection property override functionality by setting the **OverrideConnectionDetails** property to true.
- 3. Start the application, specifying the configuration file as part of the Java invocation.

Procedure

1. Create a properties file called mqclassesforjava.config containing the properties and values that need to be overridden.

It is possible to override 13 properties that are used by the IBM MQ classes for Java when connecting to a queue manager as part of the MQQueueManager constructor. The names of these properties, and the keys that must be specified when you are overriding them, are shown in the following table:

Table 19. Properties that can be overridden		
Property	Property key	
CCSID	\$CCSID_PROPERTY	
Channel	\$CHANNEL_PROPERTY	
Connect options	\$CONNECT_OPTIONS_PROPERTY	
Hostname	\$HOST_NAME_PROPERTY	
SSL key reset	\$SSL_RESET_COUNT_PROPERTY	
Local address	\$LOCAL_ADDRESS_PROPERTY	
Queue manager name	qmgr	
Password	\$PASSWORD_PROPERTY	
Port	\$PORT_PROPERTY	
Cipher suite	\$SSL_CIPHER_SUITE_PROPERTY	
FIPS required	\$SSL_FIPS_REQUIRED_PROPERTY	
SSL peer name	\$SSL_PEER_NAME_PROPERTY	
User ID	\$USER_ID_PROPERTY	

Note: All of the property keys start with the \$ character, except for the queue manager name. The reason for this is because the queue manager name is passed in to the MQQueueManager constructor as an argument, rather than being set as either a static field on the MQEnvironment class, or a property in a Hashtable, and so internally this property needs to be treated slightly differently from the other properties.

To override a property, add an entry in the following format to the properties file:

mqj.property key=value

For example, to set the name of the channel to be used when creating MQQueueManager objects, you can add the following entry to the properties file:

```
mqj.$CHANNEL_PROPERTY=MY.NEW.CHANNEL
```

To change the name of the queue manager that an MQQueueManager object connects to, you can add the following entry to the properties file:

```
mqj.qmgr=MY.OTHER.QMGR
```

2. Enable the connection override functionality by setting the **com.ibm.mq.overrideConnectionDetails** property to be true.

Setting the property **com.ibm.mq.overrideConnectionDetails** to be true means that the properties that are specified in the properties file are used to override the values specified in the application. You can either set the extra property as another property in the configuration file itself, or pass the property as a system property, by using:

```
-Dcom.ibm.mq.overrideConnectionDetails=true
```

3. Start the application.

Pass the properties file you created to the client application at run time by setting the Java system property:

```
-Dcom.ibm.msg.client.config.location
```

Note that the location of the configuration file must be specified as a URI, for example:

```
-Dcom.ibm.msg.client.config.location=file:///classesforjava/mqclassesforjava.config
```

Overriding connection properties: example with IBM MQ classes for JMS

This example shows how to override properties when you are using the IBM MQ classes for JMS.

About this task

The following code example shows how an application creates a Connection Factory programmatically:

The ConnectionFactory is configured to connect to the queue manager QM_V80 using the CLIENT transport and channel MY.CHANNEL.

You can override the connection details by using a properties file, and force the application to connect to a different channel, by using the following procedure.

Procedure

1. Create an IBM MQ classes for JMS configuration file that is called jms.config in the /userHome directory (where userHome is your home directory).

Create this file with the following contents:

```
jmscf.CHANNEL=MY.TLS.CHANNEL
jmscf.SSLCIPHERSUITE=TLS_RSA_WITH_AES_128_CBC_SHA256
```

2. Run the application, passing the following Java system properties into the Java runtime environment that the application is running in:

```
-Dcom.ibm.msg.client.config.location=file:///userHome/jms.config
-Dcom.ibm.msg.client.jms.overrideConnectionFactory=true
```

Results

Carrying out this procedure overrides the ConnectionFactory that was created programmatically by the application, so that when the application creates a connection, it tries to connect by using the channel MY.TLS.CHANNEL and the cipher suite TLS RSA WITH AES 128 CBC SHA256.

Related tasks

"Using IBM MQ connection property override" on page 181

Connection property override allows you to change the details that are used by a client application to connect to a queue manager, without modifying the source code.

"Using connection property override in IBM MQ classes for JMS" on page 182

If a connection factory is created programmatically, and it is not possible to modify the source code for the application that creates it, then the connection override functionality can be used to change the properties that the connection factory uses when a connection is created. However, the use of the connection override functionality with connection factories defined in JNDI is not supported.

"Using connection property override in IBM MQ classes for Java" on page 184

In the IBM MQ classes for Java, connection details are set as properties using a combination of different values. The connection override functionality can be used to override the connection details that an application uses if it is not possible to modify the source code for the application.

Troubleshooting IBM MQ.NET problems

You can use the .NET sample applications to help with troubleshooting problems.

Using the sample applications

If a program does not complete successfully, run one of the .NET sample applications, and follow the advice given in the diagnostic messages. These sample applications are described in <u>Sample applications</u> for .NET.

If the problems continue and you need to contact the IBM service team, you might be asked to turn on the trace facility. For information on using the trace facility, see "Tracing IBM MQ .NET applications" on page 102.

Error messages

You might see the following common error message:

An unhandled exception of type System.IO.FileNotFoundException occurred in unknown module

If this error occurs for either amqmdnet.dll or amqmdxcs.dll, either ensure that both are registered in the Global Assembly Cache or create a configuration file that points to the amqmdnet.dll and amqmdxcs.dll assemblies. You can examine and change the contents of the assembly cache using mscorcfg.msc, which is supplied as part of the .NET framework.

If the .NET framework was unavailable when IBM MQ was installed, the classes might not be registered in the global assembly cache. You can manually rerun the registration process using the command

```
amqidnet -c MQ_INSTALLATION_PATH\bin\amqidotn.txt -l logfile.txt
```

MQ_INSTALLATION_PATH represents the high-level directory in which IBM MQ is installed.

Information about this installation is written to the specified log file (logfile.txt in this example).

Resolving problems with IBM MQ MQI clients

This collection of topics contains information about techniques for solving problems in IBM MQ MQI client applications.

An application running in the IBM MQ MQI client environment receives MQRC_* reason codes in the same way as IBM MQ server applications. However, there are additional reason codes for error conditions associated with IBM 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 MQ MQI client is linked with the correct library file.

IBM 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 MQ MQI client issues an MQCONN or MQCONNX call to a server, socket and port information is exchanged between the IBM 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 MQ MQI client application.

If the connection is successful, IBM MQ protocol messages are exchanged and further checking takes place. During the IBM 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 completion and reason codes.

Stopping IBM MQ MQI clients

Even though an IBM 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 MQ MQI clients

When an error occurs with an IBM MQ MQI client system, error messages are put into the IBM 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 MQ MQI client installation. Usually this directory is C:\Program Files\IBM\MQ\errors.
- On IBM i, these files are found in the /QIBM/UserData/mgm/errors directory

Certain client errors can also be recorded in the IBM MQ error files associated with the server to which the client was connected.

ECONNRESET error over an IBM MQ MQI client channel connection

You are receiving message AMQ9208 or AMQ9209 intermittently from TCP/IP on an IBM MQ MQI client channel connection to a local server, which is an ECONNRESET error.

An ECONNRESET TCP/IP error is caused by a connection reset by peer. This occurs when an established connection is shut down for some reason by the remote computer.

Symptom

The most common scenario shows the error:

AMQ9208I:

Error on receive from host <hostname>.

Explanation

An error occurred receiving data from <hostname> over TCP/IP. This might be due to a communications failure.

Action

The return code from the TCP/IP recv() call was <xxxxx>. Record these values and tell the systems administrator.

You might receive a different message, for example:

AMQ9209I

Connection to host < hostname > for channel < channelname > closed.

The following table shows the return codes for different operating systems for the error *ECONNRESET* Connection Reset by Peer:

Operating system	Decimal	Hexadecimal
AIX	73	x49
IBM i	3426	xD62
Linux	104	x68
Windows	10054	x2746
z/OS	1121	x461

Diagnosing the problem

An ECONNRESET error usually indicates a problem in the TCP/IP network.

There are numerous reasons TCP/IP sends a reset:

- A connection termination that is not orderly, such as a rebooting the client box, can cause a reset.
- An application request a connect to a port and IP address for which no server is listening.
- An application closes a socket with data still in the application receive buffer. The connection is reset to allow the remote partner to know that the data was not delivered.
- Any data that arrives for a connection that has been closed can cause a reset.
- An application closes a socket and sets the linger socket option to zero. This notifies TCP/IP that the connection should not linger.

Note: IBM MQ does not code the linger socket option, therefore IBM MQ does not cause a reset.

• A TCP segment that is not valid arrives for a connection, for example, a bad acknowledge or sequence number can cause a reset.

- The connect request times out. TCP gives up trying to connect to an particular port and IP address and resets the connection.
- A firewall can reset connections if the packet does not adhere to the firewall rules and policies. For example, a source or destination port, or IP address does not match the firewall rule or policy.
- The retransmit timer expires. TCP gives up trying to retransmit a packet and resets the connection.
- A bad hardware device can cause resets

Diagnostic hints and tips:

Consult with your network administrator, who can use TCP/IP packet and a sniffer traces to determine why the reset occurred.

For z/OS, see :

- Instructions for setting a SLIP on an IBM MQ error message and including TCP/IP CTRACE and TCP/IP PACKET trace
- z/OS UNIX reason codes for the last two bytes of the reason code found in the CSQX208E error message.

Resolving the problem

These types of errors are not generated by IBM MQ. IBM MQis simply informing you that the network is having a problem.

Important: This problem is beyond the scope of IBM MQ Support and there is nothing that can be done from the IBM MQ perspective to resolve this network problem. You need to work with your network support team.

See <u>Automatic client reconnection</u> for more information, which could be useful for your development team to use in IBM MQ MQI client applications.

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 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 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 mgclient.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 IBM 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 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 IBM MQ Multicast topic topologies are not recommended.

As was mentioned in <u>IBM MQ Multicast topic topology</u>, IBM 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 https://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 <u>IBM MQ Multicast topic topology</u>, and shows 2 possible multicast data streams. Although it is a simple representation, it demonstrates the kind of situation that IBM MQ Multicast was designed for, and is shown here to contrast the second example:

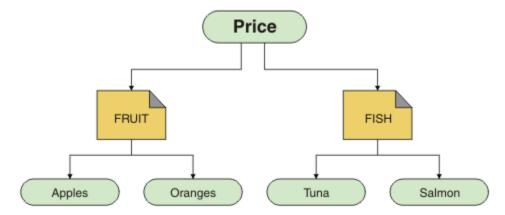
```
DEF COMMINFO(MC1) GRPADDR(
227.20.133.1)
DEF COMMINFO(MC2) GRPADDR(227.20.133.2)
```

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.

IBM 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)
```

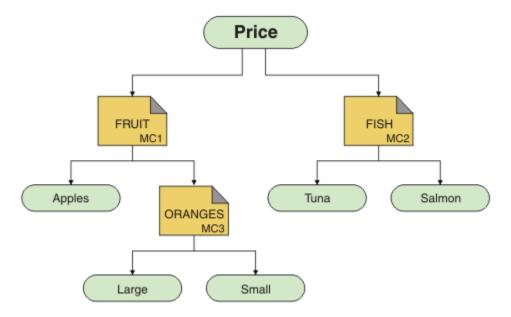
where 227.20.133.1, 227.20.133.2, and 227.20.133.3 are valid multicast addresses.

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.

Queue manager clusters troubleshooting

Use the checklist given here, and the advice given in the subtopics, to help you to detect and deal with problems when you use queue manager clusters.

Before you begin

If your problems relate to publish/subscribe messaging using clusters, rather than to clustering in general, see "Routing for publish/subscribe clusters: Notes on behavior" on page 160.

Procedure

• Check that your cluster channels are all paired.

Each cluster sender channel connects to a cluster receiver channel of the same name. If there is no local cluster receiver channel with the same name as the cluster sender channel on the remote queue manager, then it won't work.

• Check that your channels are running. No channels should be in RETRYING state permanently. Show which channels are running using the following command:

```
runmqsc display chstatus(*)
```

If you have channels in RETRYING state, there might be an error in the channel definition, or the remote queue manager might not be running. While channels are in this state, messages are likely to build up on transmit queues. If channels to full repositories are in this state, then the definitions of cluster objects (for example queues and queue managers) become out-of-date and inconsistent across the cluster.

· Check that no channels are in STOPPED state.

Channels go into STOPPED state when you stop them manually. Channels that are stopped can be restarted using the following command:

```
runmqsc start channel(xyz)
```

A clustered queue manager auto-defines cluster channels to other queue managers in a cluster, as required. These auto-defined cluster channels start automatically as needed by the queue manager, unless they were previously stopped manually. If an auto-defined cluster channel is stopped manually, the queue manager remembers that it was manually stopped and does not start it automatically in the future. If you need to stop a channel, either remember to restart it again at a convenient time, or else issue the following command:

```
stop channel(xyz) status(inactive)
```

The status(inactive) option allows the queue manager to restart the channel at a later date if it needs to do so.

• Check that all queue managers in the cluster are aware of all the full repositories.

You can do this using the following command:

```
runmqsc display clusqmgr(*) qmtype
```

Partial repositories might not be aware of all other partial repositories. All full repositories should be aware of all queue managers in the cluster. If cluster queue managers are missing, this might mean that certain channels are not running correctly.

• Check that every queue manager (full repositories and partial repositories) in the cluster has a manually defined cluster receiver channel running and is defined in the correct cluster.

To see which other queue managers are talking to a cluster receiver channel, use the following command:

```
runmqsc display chstatus(*) rqmname
```

Check that each manually defined cluster receiver has a **conname** parameter defined to be ipaddress(port). Without a correct connection name, the other queue manager does not know the connection details to use when connecting back.

• Check that every partial repository has a manually defined cluster sender channel running to a full repository, and defined in the correct cluster.

The cluster sender channel name must match the cluster receiver channel name on the other queue manager.

• Check that every full repository has a manually defined cluster sender channel running to every other full repository, and defined in the correct cluster.

The cluster sender channel name must match the cluster receiver channel name on the other queue manager. Each full repository does not keep a record of what other full repositories are in the cluster. It assumes that any queue manager to which it has a manually defined cluster sender channel is a full repository.

• Check the dead letter queue.

Messages that the queue manager cannot deliver are sent to the dead letter queue.

• Check that, for each partial repository queue manager, you have defined a single cluster-sender channel to one of the full repository queue managers.

This channel acts as a "bootstrap" channel through which the partial repository queue manager initially joins the cluster.

• Check that the intended full repository queue managers are actual full repositories and are in the correct cluster.

You can do this using the following command:

runmqsc display qmgr repos reposnl

• Check that messages are not building up on transmit queues or system queues.

You can check transmit queues using the following command:

```
runmqsc display ql(*) curdepth where (usage eq xmitq)
```

You can check system queues using the following command:

```
display ql(system*) curdepth
```

Related tasks

Configuring a queue manager cluster

"Making initial checks on UNIX, Linux, and Windows" on page 9

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

"Making initial checks on z/OS" on page 27

Before you start problem determination in detail on z/OS, 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.

"Making initial checks on IBM i" on page 18

Before you start problem determination in detail on IBM i, 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.

Related reference

Messages and reason codes

Application issues seen when running REFRESH CLUSTER

Issuing **REFRESH CLUSTER** is disruptive to the cluster. It might make cluster objects invisible for a short time until the **REFRESH CLUSTER** processing completes. This can affect running applications. These notes describe some of the application issues you might see.

Reason codes that you might see from MOOPEN, MOPUT, or MOPUT1 calls

During **REFRESH CLUSTER** the following reason codes might be seen. The reason why each of these codes appears is described in a later section of this topic.

- 2189 MQRC_CLUSTER_RESOLUTION_ERROR
- 2085 MQRC_UNKNOWN_OBJECT_NAME
- 2041 MQRC_OBJECT_CHANGED
- 2082 MQRC_UNKNOWN_ALIAS_BASE_Q
- 2270 MQRC_NO_DESTINATIONS_AVAILABLE

All these reason codes indicate name lookup failures at one level or another in the IBM MQ code, which is to be expected if apps are running throughout the time of the **REFRESH CLUSTER** operation.

The **REFRESH CLUSTER** operation might be happening locally, or remotely, or both, to cause these outcomes. The likelihood of them appearing is especially high if full repositories are very busy. This happens if **REFRESH CLUSTER** activities are running locally on the full repository, or remotely on other queue managers in the cluster or clusters that the full repository is responsible for.

In respect of cluster queues that are absent temporarily, and will shortly be reinstated, then all of these reason codes are temporary retry-able conditions (although for 2041 MQRC_OBJECT_CHANGED it can be a little complicated to decide whether the condition is retry-able). If consistent with application rules

(for example maximum service times) you should probably retry for about a minute, to give time for the **REFRESH CLUSTER** activities to complete. For a modest sized cluster, completion is likely to be much quicker than that.

If any of these reason codes is returned from **MQOPEN**, then no object handle is created, but a later retry should be successful in creating one.

If any of these reason codes is returned from **MQPUT**, then the object handle is not automatically closed, and retrying should eventually succeed without a need first to close the object handle. However, if the application opened the handle using bind-on-open options, and so requires all messages to go to the same channel, then (contrary to the application's expectations) it is not guaranteed that the retried *put* would go to the same channel or queue manager as before. It is therefore wise to close the object handle and open a new one, in that case, to regain the bind-on-open semantics.

If any of these reason codes is returned from **MQPUT1**, then it is unknown whether the problem happened during the *open* or the *put* part of the operation. Whichever it is, the operation can be retried. There are no bind-on-open semantics to worry about in this case, because the **MQPUT1** operation is an *open-put-close* sequence that is performed in one continuous action.

Multi-hop scenarios

If the message flow incorporates a multi-hop, such as that shown in the following example, then a name lookup failure caused by **REFRESH CLUSTER** can occur on a queue manager that is remote from the application. In that case, the application receives a success (zero) return code, but the name lookup failure, if it occurs, prevents a **CLUSRCVR** channel program from routing the message to any proper destination queue. Instead, the **CLUSRCVR** channel program follows normal rules to write the message to a dead letter queue, based on the persistence of the message. The reason code associated with that operation is this:

• 2001 MORC ALIAS BASE Q TYPE ERROR

If there are persistent messages, and no dead letter queues have been defined to receive them, you will see channels ending.

Here is an example multi-hop scenario:

- MOOPEN on queue manager QM1 specifies Q2.
- Q2 is defined in the cluster on a remote queue manager QM2, as an alias.
- A message reaches QM2, and finds that Q2 is an alias for Q3.
- Q3 is defined in the cluster on a remote queue manager QM3, as a qlocal.
- The message reaches QM3, and is put to Q3.

When you test the multi-hop, you might see the following queue manager error log entries:

• On the sending and receiving sides, when dead letter queues are in place, and there are persistent messages:

AMQ9544: Messages not put to destination queue

During the processing of channel 'CHLNAME' one or more messages could not be put to the destination queue and attempts were made to put them to a dead letter queue. The location of the queue is \$, where 1 is the local dead letter queue and 2 is the remote dead letter queue.

• On the receiving side, when a dead letter queue is not in place, and there are persistent messages:

AMQ9565: No dead letter queue defined

AMQ9599: Program could not open a queue manager object

AMQ9999: Channel program ended abnormally

• On the sending side, when a dead letter queue is not in place, and there are persistent messages:

AMO9506: Message receipt confirmation failed

AMQ9780: Channel to remote machine 'a.b.c.d(1415)' is ending because of an error

More details about why each of these reason codes might be displayed when running REFRESH CLUSTER

2189 (088D) (RC2189): MQRC_CLUSTER_RESOLUTION_ERROR

The local queue manager asked its full repositories about the existence of a queue name. There was no response from the full repositories within a hard-coded timeout of 10 seconds. This is because the request message or the response message is on a queue for processing, and this condition will be cleared in due course. At the app, the condition is retry-able, and will succeed when those internal mechanisms have completed.

MQRC_UNKNOWN_OBJECT_NAME (2085, X'825')

The local queue manager asked (or has previously asked) its full repositories about the existence of a queue name. The full repositories have responded, saying that they did not know about the queue name. In the context of **REFRESH CLUSTER** taking place on full and partial repositories, the owner of the queue might not yet have told the full repositories about the queue. Or it might have done so, but the internal messages carrying this information are on a queue for processing, in which case this condition will be cleared in due course. At the app, the condition is retry-able, and will succeed when those internal mechanisms have completed.

2041 (07F9) (RC2041): MQRC_OBJECT_CHANGED

Most likely to be seen from bind-on-open MQPUT. The local queue manager knows about the existence of a queue name, and about the remote queue manager where it resides. In the context of REFRESH CLUSTER taking place on full and partial repositories, the record of the queue manager has been deleted and is in the process of being queried from the full repositories. At the app, it is a little complicated to decide whether the condition is retry-able. In fact, if the MQPUT is retried, it will succeed when those internal mechanisms have completed the job of learning about the remote queue manager. However there is no guarantee that the same queue manager will be used. It is safer to follow the approach usually recommended when MQRC_OBJECT_CHANGED is received, which is to close the object handle and re-open a new one.

MQRC UNKNOWN ALIAS BASE Q (2082, X'822')

Similar in origin to the 2085 MQRC_UNKNOWN_OBJECT_NAME condition, this reason code is seen when a local alias is used, and its TARGET is a cluster queue that is inaccessible for the reasons previously described for reason code 2085.

MQRC_ALIAS_BASE_Q_TYPE_ERROR (2001, X'7D1')

This reason code is not usually seen at applications. It is only likely to be seen in the queue manager error logs, in relation to attempts to send a message to a dead letter queue. A **CLUSRCVR** channel program has received a message from its partner **CLUSSDR** and is deciding where to put it. This scenario is just a variation of the same condition previously described for reason codes 2082 and 2085. In this case, the reason code is seen when an alias is being processed at a different point in the MQ product, compared to where it is processed during an application **MOPUT** or **MOOPEN**.

2270 (08DE) (RC2270): MQRC_NO_DESTINATIONS_AVAILABLE

Seen when an application is using a queue that it opened with MQOO_BIND_NOT_FIXED, and the destination objects are unavailable for a short time until the **REFRESH CLUSTER** processing completes.

Further remarks

If there is any clustered publish/subscribe activity in this environment, then **REFRESH CLUSTER** can have additional unwanted effects. For example temporarily losing subscriptions for subscribers, that then find they missed a message. See REFRESH CLUSTER considerations for publish/subscribe clusters.

Related concepts

REFRESH CLUSTER considerations for publish/subscribe clusters

Clustering: Using REFRESH CLUSTER best practices

Related reference

MQSC Commands reference: REFRESH CLUSTER

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)
CONNAME(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 conname parameter might be incorrect:

```
1 : alter chl(demo.qm2) chltype(clusrcvr) conname('newhost(1414)')
AMQ8016: IBM 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

Multi

```
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.<yourdomain>(1414))
CLUSTER(DEMO) CHANNEL(DEMO.QM2)
```

z/0S

```
CSQM201I +CSQ2 CSQMDRTC DISPLAY CLUSQMGR DETAILS
CLUSQMGR(SYSTEM.TEMPQMGR.<HOSTNAME>(1716))
CLUSTER(DEMO)
CHANNEL(TO.CSQ1.DEMO)
END CLUSQMGR DETAILS
```

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.

You also need to give the correct authority to the SYSTEM.CLUSTER.TRANSMIT.QUEUE by issuing the following command:

```
setmqaut -m <QMGR Name> -n SYSTEM.CLUSTER.TRANSMIT.QUEUE -t q -g mqm +all
```

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 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 MQRC_NOT_AUTHORIZED for information on:

- JMSWMQ2013 invalid security authentication
- MORC NOT AUTHORIZED on a gueue 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.
                        CLUSTER(DEMO)
CLUSQMGR (QM1)
CHANNEL (DEMO.QM1)
                        QMTYPE(NORMAL)
AMQ8441: Display Cluster Queue Manager details. CLUSQMGR(QM2) CLUSTER(DEMO)
CHANNEL (DEMO.OM2)
                        QMTYPE (REPOS)
STATUS(RUNNING)
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM3)
                        CLUSTER(DEMO)
CHANNEL (DEMO.OM3)
                        OMTYPE (REPOS)
STATUS(RUNNING)
```

Note: When using IBM MQ with WebSphere Application Server, you might also see this issue if you have a JMS application which connects to an IBM MQ queue manager belonging to an IBM MQ cluster and your JMS application tries to access a cluster queue which somewhere else in the cluster. Your application needs to leave the queue manager blank if it wants to open a cluster queue located in the cluster, or specify the name of a queue manager in the cluster which hosts the cluster queue.

Related reference

MQRC_UNKNOWN_OBJECT_NAME (2085, X'825')

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)
```

Related reference

2189 (088D) (RC2189): MQRC_CLUSTER_RESOLUTION_ERROR

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

Leave the **ObjectQmgrName** parameter blank, so that the clustering decides which queue manager to route to.

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.

Related reference

MORC UNKNOWN ALIAS BASE Q (2082, X'822')

Messages are not arriving on the destination queues

Make sure that the corresponding cluster transmission 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

1. Identify the transmission queue that is sending messages to the destination and the status of the channel.

```
1 : dis clusqmgr(QM1) CHANNEL(*) STATUS DEFTYPE QMTYPE XMITQ
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1) CLUSTER(DEMO)
CHANNEL(DEMO.QM1) DEFTYPE(CLUSSDRA)
QMTYPE(NORMAL) STATUS(RUNNING)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM1)
```

2. Make sure that the cluster transmission queue is empty.

```
1 : display ql(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM1) curdepth
AMQ8409: Display Queue details.
QUEUE(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM1) CURDEPTH(0)
```

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 MQ00_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 MQ00_BIND_ON_OPEN. If the alias queue had been opened with MQ00_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 MQ00_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
QTIMX
AMQ8441: Display Cluster Queue Manager details.
CLUSQMGR(QM1)
                    CLUSTER (DEMO)
CHANNEL (DEMO.QM1) DEFTYPE (CLUSSDRA)
QMTYPE(NORMAL)
                    STATUS(RUNNING)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM1)
AMO8441: Display Cluster Queue Manager details. CLUSQMGR(QM2) CLUSTER(DEMO)
CHANNEL(DEMO.QM2) DEFTYPE(CLUSRCVR)
QMTYPE(REPOS)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM2)
AMQ8441: Display Cluster Queue Manager details. CLUSQMGR(QM3) CLUSTER(DEMO)
CHANNEL(DEMO.QM3) DEFTYPE(CLUSSDRB)
OMTYPE(REPOS)
                    STATUS (RUNNING)
XMITQ(SYSTEM.CLUSTER.TRANSMIT.DEMO.QM3)
AMQ8441: Display Cluster Queue Manager details. CLUSQMGR(QM4) CLUSTER(DEMO)
CHANNEL (DEMO.QM4) DEFTYPE (CLUSSDRA)
OMTYPE (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)
```

- 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. After about 90 days, the cluster's knowledge of the older version of the queue manager expires, and is deleted automatically. However you might prefer to delete this information manually.

Cause

- 1. The queue manager might have been deleted and then re-created and redefined.
- 2. It might have been cold-started on z/OS, without first following the procedure to remove a queue manager from a cluster.

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 information

Clustering: Using REFRESH CLUSTER best practices

Workload balancing set on a cluster-sender channel is not working

Any workload balancing you specify on a cluster-sender channel is likely to be ignored. Instead, specify the cluster workload channel attributes on the cluster-receiver channel at the target queue manager.

Symptom

You have specified one or more cluster workload channel attributes on a cluster-sender channel. The resulting workload balancing is not as you were expecting.

Cause

Any workload balancing you specify on a cluster-sender channel is likely to be ignored. For an explanation of this, see <u>Cluster channels</u>. Note that you still get some form of workload balancing, based either on cluster defaults or on properties set on the matching cluster-receiver channel at the target queue manager.

Solution

Specify the cluster workload channel attributes on the cluster-receiver channel at the target queue manager.

Related reference

CLWLPRTY channel attribute CLWLRANK channel attribute CLWLWGHT channel attribute NETPRTY channel attribute

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.

When you run REFRESH CLUSTER (DEMO) on QM1, you remove all the information QM1 has about the cluster DEMO, except for QM1's knowledge of itself and its own queues, and of how to access the full repositories in the cluster. QM1 then contacts the full repositories, and tells them about itself and its queues. QM1 is a partial repository, so the full repositories don't immediately tell QM1 about all the other partial repositories in the cluster. Instead, QM1 slowly builds up its knowledge of the other partial repositories through information it receives as and when each of the other queues and queue managers is next active in the cluster.

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 <u>Refreshing in a large cluster can affect performance and availability of the cluster.</u>

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 <u>Refreshing in a large cluster can affect performance and availability of the cluster.</u>

QM1 removes all information it has about the cluster DEM0, 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 <u>Refreshing in a large cluster can affect</u> 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 <u>Refreshing in a large cluster can affect performance and availability of the cluster.</u>

- 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.
 - DIED ON IBM MQ for z/OS, the SYSTEM. CLUSTER. COMMAND. QUEUE is disabled for MQGET.
 - 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 MQOO_BIND_ON_OPEN returns a return code of MQRC_CLUSTER_PUT_INHIBITED to the application. If MQOO_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

Potential issues when switching transmission queues

A list of some issues that might be encountered when switching transmission queue, their causes, and most likely solutions.

Insufficient access to transmission queues on z/OS

Symptom

A cluster-sender channel on z/OS might report it is not authorized to open its transmission queue.

Cause

The channel is switching, or has switched, transmission queue and the channel initiator has not been granted authority to access the new queue.

Solution

Grant the channel initiator the same access to the channel's transmission queue that is documented for the transmission queue SYSTEM.CLUSTER.TRANSMIT.QUEUE. When using DEFCLXQ a generic profile for SYSTEM.CLUSTER.TRANSMIT.** avoids this problem occurring whenever a new queue manager joins the cluster.

Moving of messages fails

Symptom

Messages stop being sent by a channel and they remain queued on the channel's old transmission queue.

Cause

The queue manager has stopped moving messages from the old transmission queue to the new transmission queue because an unrecoverable error occurred. For example, the new transmission queue might have become full or its backing storage exhausted.

Solution

Review the error messages written to the queue manager's error log (job log on z/OS) to determine the problem and resolve its root cause. Once resolved, restart the channel to resume the switching process, or stop the channel then use **runswchl** instead (CSQUTIL on z/OS).

A switch does not complete

Symptom

The queue manager repeatedly issues messages that indicate it is moving messages. The switch never completes because there are always messages remaining on the old transmission queue.

Cause 1

Messages for the channel are being put to the old transmission queue faster than the queue manager can move them to the new transmission queue. This is likely to be a transient issue during peak workload because if were commonplace then it is unlikely the channel would be able to transmit the messages over the network fast enough.

Cause 2

There are uncommitted messages for the channel on the old transmission queue.

Solution

Resolve the units of work for any uncommitted messages, and/or reduce or suspend the application workload, to allow the moving message phase to complete.

Accidental deletion of a transmission queue

Symptom 1

Channels unexpectedly switch due to the removal of a matching CLCHNAME value.

Symptom 2

A put to a cluster queue fails with MQRC_UNKNOWN_XMIT_Q.

Symptom 3

A channel abnormally ends because its transmission queue does not exist.

Symptom 4

The queue manager is unable to move messages to complete a switch operation because it cannot open either the old or the new transmission queue.

Cause

The transmission queue currently used by a channel, or its previous transmission queue if a switch has not completed, has been deleted.

Solution

Redefine the transmission queue. If it is the old transmission queue that has been deleted then an administrator may alternatively complete the switch operation using **runswch1** with the **-n** parameter (or CSQUTIL with MOVEMSGS(NO) on z/OS).

Use the -n parameter with caution because, if it is used inappropriately, messages for the channel can complete and finish processing but not be updated on the old transmission queue. In this scenario it is safe because as the queue does not exist there cannot be any messages to complete and finish processing.

Queue managers troubleshooting

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 MQ configuration file references the correct queue manager and log directories. On Windows, check for problems in the qm.ini file.

IBM MQ coordinating with Db2 as the resource manager error

• **Scenario:** You start your queue managers from the IBM MQ Explorer and are having problems when coordinating Db2. When you check your queue manager error logs, you see an error like the one shown in the following example:

```
23/09/2008 15:43:54 - Process(5508.1) User(MUSR_MQADMIN) Program(amqzxma0.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 MQ Service process amqsvc.exe is still running with an access token which does not contain group membership information for the group DB2USERS.
- **Solution:** After you have ensured that the IBM MQ Service user ID is a member of DB2USERS, use the following sequence of commands:
 - 1. Stop the service.
 - 2. Stop any other processes running under the same user ID.
 - 3. Restart these processes.

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

Undelivered messages troubleshooting

Use the advice given here to help you to resolve problems when messages do are 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 OUEUE 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 <u>Working with</u> dead-letter queues.

TLS troubleshooting information

Use the information listed here to help you solve problems with your TLS system.

Overview

For the error caused by *Using non-FIPS cipher with FIPS enabled on client*, you receive the following error message:

JMSCMQ001

IBM MQ call failed with completion code 2 ('MQCC_FAILED') reason 2397 ('MQRC_JSSE_ERROR')

For every other problem documented within this topic you receive either the previous error message, or the following error message, or both:

JMSWMQ0018

Failed to connect to queue manager 'queue_manager_name' with connection mode 'connection_mode' and host name 'host_name'

For each problem documented within this topic, the following information is provided:

- Output from the sample SystemOut.log or Console, detailing the cause of the exception..
- · Queue manager error log information.
- · Solution to the problem.

Note:

- You should always list out the stacks and the cause of the first exception.
- Whether or not the error information is written to the stdout log file depends on how the application is written, and on which framework you are using.
- The sample code includes stacks and line numbers. This information is useful guidance, but the stacks and line numbers are likely to change from one fix pack to another. You should use the stacks and line numbers as a guide to locating the correct section, and not use the information specifically for diagnostic purposes.

Cipher suite not set on client

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9641: Remote CipherSpec error for channel 'SYSTEM.DEF.SVRCONN' to host ''. [3=5YSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9639: Remote channel 'SYSTEM.DEF.SVRCONN' did not specify a CipherSpec.

Solution

Set a CipherSuite on the client so that both ends of the channel have a matching CipherSuite or CipherSpec pair.

Cipher suite not set on server

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9641: Remote CipherSpec error for channel 'SYSTEM.DEF.SVRCONN' to host ''. [3=SYSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9639: Remote channel 'SYSTEM.DEF.SVRCONN' did not specify a CipherSpec.

Solution

Change channel SYSTEM.DEF.SVRCONN to specify a valid CipherSpec.

Cipher Mismatch

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9641: Remote CipherSpec error
for channel 'SYSTEM.DEF.SVRCONN' to host ''. [3=SYSTEM.DEF.SVRCONN]
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection
(RemoteConnectionSpecification.java:409)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession
(RemoteConnectionSpecification.java:305)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9631: The CipherSpec negotiated during the TLS handshake does not match the required CipherSpec for channel 'SYSTEM.DEF.SVRCONN'.

Solution

Change either the SSLCIPH definition of the server-connection channel or the Cipher suite of the client so that the two ends have a matching CipherSuite or CipherSpec pair.

Missing client personal certificate

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2059;AMQ9503: Channel negotiation failed. [3=SYSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9637: Channel is lacking a certificate.

Solution

Ensure that the key database of the queue manager contains a signed personal certificate from the truststore of the client.

Missing server personal certificate

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9771: SSL handshake failed.
[1=javax.net.ssl.SSLHandshakeException[Remote host closed connection during handshake],
3=localhost/127.0.0.1:1418 (localhost),4=SSLSocket.startHandshake,5=default]
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1173)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:835)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection
(RemoteConnectionSpecification.java:409)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession
(RemoteConnectionSpecification.java:305)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
... 12 more
```

Caused by:

```
javax.net.ssl.SSLHandshakeException: Remote host closed connection during handshake at com.ibm.jsse2.qc.a(qc.java:158) at com.ibm.jsse2.qc.h(qc.java:185) at com.ibm.jsse2.qc.a(qc.java:566) at com.ibm.jsse2.qc.startHandshake(qc.java:120) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134) at java.security.AccessController.doPrivileged(AccessController.java:229) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134) ... 17 more
```

Caused by:

```
java.io.EOFException: SSL peer shut down incorrectly
at com.ibm.jsse2.a.a(a.java:19)
at com.ibm.jsse2.qc.a(qc.java:207)
```

Queue manager error logs

AMQ9637: Channel is lacking a certificate.

Solution

Ensure that the key database of the queue manager contains a signed personal certificate from the truststore of the client.

Missing server signer on client

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9771: SSL handshake failed.
[1=javax.net.ssl.SSLHandshakeException[com.ibm.jsse2.util.j:
PKIX path validation failed: java.security.cert.CertPathValidatorException:
The certificate issued by CN=JohnDoe, O=COMPANY, L=YOURSITE, C=XX is not trusted; internal cause is: java.security.cert.CertPathValidatorException: Signature does not match.],3=localhost/127.0.0.1:1418 (localhost),4=SSLSocket.startHandshake,5=default]
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1173) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:835) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868) ...
```

Caused by:

```
javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.j: PKIX path validation failed: java.security.cert.CertPathValidatorException: The certificate issued by CN=JohnDoe, O=COMPANY, L=YOURSITE, C=XX is not trusted; internal cause is: java.security.cert.CertPathValidatorException: Signature does not match. ...
```

Caused by:

```
com.ibm.jsse2.util.j: PKIX path validation failed: java.security.cert.CertPathValidatorException:
The certificate issued by CN=JohnDoe, 0=COMPANY, L=YOURSITE, C=XX is not trusted;
internal cause is: java.security.cert.CertPathValidatorException: Signature does not match.
at com.ibm.jsse2.util.h.a(h.java:99)
at com.ibm.jsse2.util.h.b(h.java:27)
at com.ibm.jsse2.util.g.a(g.java:14)
at com.ibm.jsse2.yc.a(yc.java:68)
at com.ibm.jsse2.yc.a(yc.java:17)
at com.ibm.jsse2.yc.a(yc.java:17)
at com.ibm.jsse2.yc.checkServerTrusted(yc.java:154)
at com.ibm.jsse2.bb.a(bb.java:246)
... 28 more
```

Caused by:

```
java.security.cert.CertPathValidatorException:
The certificate issued by CN=JohnDoe, 0=COMPANY, L=YOURSITE, C=XX is not trusted;
internal cause is: java.security.cert.CertPathValidatorException: Signature does not match.
at com.ibm.security.cert.BasicChecker.(BasicChecker.java:111)
at com.ibm.security.cert.PKIXCertPathValidatorImpl.engineValidate(PKIXCertPathValidatorImpl.java:174)
at java.security.cert.CertPathValidator.validate(CertPathValidator.java:265)
at com.ibm.jsse2.util.h.a(h.java:13)
... 34 more
```

Caused by:

```
java.security.cert.CertPathValidatorException: Signature does not match.
at com.ibm.security.cert.CertPathUtil.findIssuer(CertPathUtil.java:297)
at com.ibm.security.cert.BasicChecker.(BasicChecker.java:108)
```

Queue manager error logs

AMQ9665: SSL connection closed by remote end of channel '?????'.

Solution

Add the certificate used to sign the personal certificate of the queue manager to the truststore of the client.

Missing client signer on server

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9771: SSL handshake failed.
[1=java.net.SocketException[Software caused connection abort: socket write error],
3=localhost/127.0.0.1:1418 (localhost),4=SSLSocket.startHandshake,5=default]
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1173)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:835)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection
(RemoteConnectionSpecification.java:409)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession
(RemoteConnectionSpecification.java:305)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
... 12 more
```

Caused by:

```
java.net.SocketException: Software caused connection abort: socket write error at java.net.SocketOutputStream.socketWrite(SocketOutputStream.java:120) at java.net.SocketOutputStream.write(SocketOutputStream.java:164) at com.ibm.jsse2.c.a(c.java:57) at com.ibm.jsse2.c.a(c.java:34) at com.ibm.jsse2.qc.a(qc.java:527) at com.ibm.jsse2.qc.a(qc.java:635) at com.ibm.jsse2.qc.a(qc.java:743) at com.ibm.jsse2.bb.b(bb.java:743) at com.ibm.jsse2.bb.b(bb.java:194) at com.ibm.jsse2.bb.d(bb.java:194) at com.ibm.jsse2.bb.d(bb.java:162) at com.ibm.jsse2.bb.a(bb.java:7) at com.ibm.jsse2.bb.a(ab.java:560) at com.ibm.jsse2.ab.r(ab.java:332) at com.ibm.jsse2.ab.a(ab.java:332) at com.ibm.jsse2.qc.a(qc.java:435) at com.ibm.jsse2.qc.a(qc.java:435) at com.ibm.jsse2.qc.a(qc.java:185) at com.ibm.jsse2.qc.a(qc.java:566) at com.ibm.jsse2.qc.startHandshake(qc.java:120) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134) at java.security.AccessController.doPrivileged(AccessController.java:229) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134)
```

Queue manager error logs

AMQ9633: Bad SSL certificate for channel '?????'.

Solution

Add the certificate used to sign the personal certificate of the client to the key database of the queue manager.

SSLPEER set on server does not match certificate

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9643: Remote SSL peer name error for channel 'SYSTEM.DEF.SVRCONN' on host ''. [3=SYSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:24176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9636: SSL distinguished name does not match peer name, channel 'SYSTEM.DEF.SVRCONN'.

Solution

Ensure the value of SSLPEER set on the server-connection channel matches the distinguished name of the certificate.

SSLPEER set on client does not match certificate

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2398;AMQ9636: SSL distinguished name does not match peer name,
channel '?'. [CN=JohnDoe, O=COMPANY, L=YOURSITE, C=XX]
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1215)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:835)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection
(RemoteConnectionSpecification.java:409)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession
(RemoteConnectionSpecification.java:305)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9208: Error on receive from host host-name (address).

Solution

Ensure the value of SSLPEER set in the client matches the distinguished name of the certificate.

Using a non-FIPS cipher with FIPS enabled on client

Output

```
Check the queue manager is started and if running in client mode, check there is a listener running. Please see the linked exception for more information.
at com.ibm.msg.client.wmq.common.internal.Reason.reasonToException(Reason.java:578)
at com.ibm.msg.client.wmq.common.internal.Reason.createException(Reason.java:214)
at com.ibm.msg.client.wmq.internal.WMQConnection.getConnectOptions(WMQConnection.java:1423)
at com.ibm.msg.client.wmq.internal.WMQConnection.(WMQConnection.java:339)
at com.ibm.msg.client.wmq.factories.WMQConnectionFactory.createV7ProviderConnection
(WMQConnectionFactory.java:6865)
at com.ibm.msg.client.wmq.factories.WMQConnectionFactory.createProviderConnection
(WMQConnectionFactory.java:6221)
at com.ibm.msg.client.jms.admin.JmsConnectionFactoryImpl._createConnection
(JmsConnectionFactoryImpl.java:285)
at com.ibm.msg.client.jms.admin.JmsConnectionFactoryImpl.createConnection
(JmsConnectionFactoryImpl.java:233)
at com.ibm.mq.jms.MQConnectionFactory.createCommonConnection(MQConnectionFactory.java:6016)
at com.ibm.mq.jms.MQConnectionFactory.createConnection(MQConnectionFactory.java:6041)
at tests.SimpleSSLConn.runTest(SimpleSSLConn.java:46)
at tests.SimpleSSLConn.main(SimpleSSLConn.java:26)
```

Caused by:

```
com.ibm.mq.MQException: JMSCMQ0001: IBM MQ call failed with compcode '2' ('MQCC_FAILED') reason '2400' ('MQRC_UNSUPPORTED_CIPHER_SUITE'). at com.ibm.msg.client.wmq.common.internal.Reason.createException(Reason.java:202)
```

Queue manager error logs

Not applicable.

Solution

Use a FIPS-enabled cipher, or disable FIPS on the client.

Using a non-FIPS cipher with FIPS enabled on the queue manager

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2397;AMQ9771: SSL handshake failed.
[1=javax.net.ssl.SSLHandshakeException[Received fatal alert: handshake_failure],
3=localhost/127.0.0.1:1418 (localhost),4=SSLSocket.startHandshake,5=default]
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1173)
at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:835)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection
```

```
(RemoteConnectionSpecification.java:409)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession
(RemoteConnectionSpecification.java:305)
at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
... 12 more
```

Caused by:

```
javax.net.ssl.SSLHandshakeException: Received fatal alert: handshake_failure
at com.ibm.jsse2.j.a(j.java:13)
at com.ibm.jsse2.qc.lp(gc.java:601)
at com.ibm.jsse2.qc.a(qc.java:601)
at com.ibm.jsse2.qc.a(qc.java:100)
at com.ibm.jsse2.qc.a(qc.java:185)
at com.ibm.jsse2.qc.a(qc.java:566)
at com.ibm.jsse2.qc.startHandshake(qc.java:120)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134)
at java.security.AccessController.doPrivileged(AccessController.java:229)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134)
```

Queue manager error logs

AMQ9616: The CipherSpec proposed is not enabled on the server.

Solution

Use a FIPS-enabled cipher, or disable FIPS on the queue manager.

Can not find client keystore using IBM JRE

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2059;AMQ9204: Connection to host 'localhost(1418)' rejected.
[1=com.ibm.mq.jmqi.JmqiException[CC=2;RC=2059;AMQ9503: Channel negotiation failed.
[3=SYSTEM.DEF.SVRCONN]],3=localhost(1418),5=RemoteConnection.analyseErrorSegment]
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:2450)
at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1396)
at com.ibm.mq.ese.jmqi.InterceptedJmqiImpl.jmqiConnect(InterceptedJmqiImpl.java:376)
at com.ibm.mq.ese.jmqi.ESEJMQI.jmqiConnect(ESEJMQI.java:561)
at com.ibm.mg.sc.lient.wmq.internal.WMQConnection.(WMQConnection.java:342)
... 8 more
```

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2059;AMQ9503: Channel negotiation failed. [3=SYSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9637: Channel is lacking a certificate.

Solution

Ensure the JVM property javax.net.ssl.keyStore specifies the location of a valid keystore.

Can not find client keystore using Oracle JRE

Output

Caused by:

```
java.security.PrivilegedActionException: java.io.FileNotFoundException:
C:\filepath\wrongkey.jks (The system cannot find the file specified)
at java.security.AccessController.doPrivileged(Native Method)
at sun.security.ssl.SSLContextImpl$DefaultSSLContext.getDefaultKeyManager(Unknown Source)
at sun.security.ssl.SSLContextImpl$DefaultSSLContext.(Unknown Source)
at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method)
at sun.reflect.NativeConstructorAccessorImpl.newInstance(Unknown Source)
at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(Unknown Source)
at java.lang.reflect.Constructor.newInstance(Unknown Source)
at java.lang.Class.newInstance0(Unknown Source)
```

```
at java.lang.Class.newInstance(Unknown Source)
... 28 more
```

Caused by:

```
java.io.FileNotFoundException: C:\filepath\wrongkey.jks (The system cannot find the file specified)
at java.io.FileInputStream.open(Native Method)
at java.io.FileInputStream.(Unknown Source)
at java.io.FileInputStream.(Unknown Source)
at sun.security.ssl.SSLContextImpl$DefaultSSLContext$2.run(Unknown Source)
at sun.security.ssl.SSLContextImpl$DefaultSSLContext$2.run(Unknown Source)
```

Queue manager error logs

AMQ9637: Channel is lacking a certificate.

Solution

Ensure the JVM property javax.net.ssl.keyStore specifies the location of a valid keystore.

Keystore password error - IBM JRE

Output

Caused by:

```
com.ibm.mq.jmqi.JmqiException: CC=2;RC=2059;AMQ9503: Channel negotiation failed. [3=SYSTEM.DEF.SVRCONN] at com.ibm.mq.jmqi.remote.impl.RemoteConnection.analyseErrorSegment(RemoteConnection.java:4176) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.receiveTSH(RemoteConnection.java:2969) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.initSess(RemoteConnection.java:1180) at com.ibm.mq.jmqi.remote.impl.RemoteConnection.connect(RemoteConnection.java:838) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSessionFromNewConnection (RemoteConnectionSpecification.java:409) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionSpecification.getSession (RemoteConnectionSpecification.java:305) at com.ibm.mq.jmqi.remote.impl.RemoteConnectionPool.getSession(RemoteConnectionPool.java:146) at com.ibm.mq.jmqi.remote.api.RemoteFAP.jmqiConnect(RemoteFAP.java:1868)
```

Queue manager error logs

AMQ9637: Channel is lacking a certificate.

Solution

Ensure that the value of the JVM property javax.net.ssl.keyStorePassword specifies the password for the keystore specified by javax.net.ssl.keyStore.

Truststore password error - IBM JRE

Output

Caused by:

```
javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException:
No X509TrustManager implementation available
at com.ibm.jsse2.j.a(j.java:13)
at com.ibm.jsse2.qc.a(qc.java:204)
at com.ibm.jsse2.ab.a(ab.java:342)
at com.ibm.jsse2.ab.a(ab.java:222)
at com.ibm.jsse2.bb.a(bb.java:157)
at com.ibm.jsse2.bb.a(bb.java:492)
at com.ibm.jsse2.ab.r(ab.java:529)
at com.ibm.jsse2.ab.r(ab.java:332)
at com.ibm.jsse2.qc.a(qc.java:435)
at com.ibm.jsse2.qc.a(qc.java:435)
at com.ibm.jsse2.qc.a(qc.java:185)
at com.ibm.jsse2.qc.a(qc.java:166)
at com.ibm.jsse2.qc.a(qc.java:566)
at com.ibm.jsse2.qc.startHandshake(qc.java:120)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134)
at java.security.AccessController.doPrivileged(AccessController.java:229)
at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134)
... 17 more
```

Caused by:

```
java.security.cert.CertificateException: No X509TrustManager implementation available
at com.ibm.jsse2.xc.checkServerTrusted(xc.java:2)
at com.ibm.jsse2.bb.a(bb.java:246)
```

Queue manager error logs

AMQ9665: SSL connection closed by remote end of channel '????'.

Solution

Ensure that the value of the JVM property javax.net.ssl.trustStorePassword specifies the password for the keystore specified by javax.net.ssl.trustStore.

Can not find or open queue manager key database

Output

Caused by:

```
javax.net.ssl.SSLHandshakeException: Remote host closed connection during handshake at com.ibm.jsse2.qc.a(qc.java:158) at com.ibm.jsse2.qc.h(qc.java:185) at com.ibm.jsse2.qc.a(qc.java:566) at com.ibm.jsse2.qc.startHandshake(qc.java:120) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134) at java.security.AccessController.doPrivileged(AccessController.java:229) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134) ... 17 more
```

Caused by:

```
java.io.EOFException: SSL peer shut down incorrectly
at com.ibm.jsse2.a.a(a.java:19)
at com.ibm.jsse2.qc.a(qc.java:207)
```

Queue manager error logs

AMQ9657: The key repository could not be opened (channel '????').

Solution

Ensure that the key repository you specify exists and that its permissions are such that the IBM MQ process involved can read from it.

Can not find or use queue manager key database password stash file

Output

Caused by:

```
javax.net.ssl.SSLHandshakeException: Remote host closed connection during handshake at com.ibm.jsse2.qc.a(qc.java:158) at com.ibm.jsse2.qc.a(qc.java:185) at com.ibm.jsse2.qc.a(qc.java:566) at com.ibm.jsse2.qc.startHandshake(qc.java:120) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1142) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection$6.run(RemoteTCPConnection.java:1134) at java.security.AccessController.doPrivileged(AccessController.java:229) at com.ibm.mq.jmqi.remote.impl.RemoteTCPConnection.protocolConnect(RemoteTCPConnection.java:1134) ... 17 more
```

Caused by:

```
ava.io.EOFException: SSL peer shut down incorrectly
at com.ibm.jsse2.a.a(a.java:19)
at com.ibm.jsse2.qc.a(qc.java:207)
```

Queue manager error logs

AMQ9660: SSL key repository: password stash file absent or unusable.

Solution

Ensure that a password stash file has been associated with the key database file in the same directory, and that the user ID, under which IBM MQ is running, has read access to both files.

Troubleshooting RDQM configurations

These topics give information that is useful for troubleshooting RDQM high availability (HA) and disaster recovery (DR) configurations.

RDQM HA architecture

Describes the basic architecture of replicated data queue manager high availability (RDQM HA) configurations to assist with troubleshooting.

Resource names

Various resources are created for each RDQM queue manager and these resources have names based on the Directory name of the queue manager. The name can be found in the file /var/mqm/mqs.ini, and is referred to here as qm. For example, for an RDQM HA queue manager named TMPQM1, qm would be tmpqm1.

Architecture

The architecture of RDQM high availability (HA) involves both DRBD, for data replication, and Pacemaker, for managing where HA RDQM queue managers run.

When you create an RDQM HA queue manager, the following steps are completed:

- 1. A DRBD resource is created to replicate the data for the queue manager.
- 2. A queue manager is created and configured to use the DRBD resource for its storage.
- 3. A set of Pacemaker resources is created to monitor and manage the queue manager.

DRBD

Each RDQM HA queue manager has a DRBD resource file generated for it named /etc/drbd.d/qm.res. For example, when an RDQM HA queue manager named HAQM1 is created, the DRBD resource file is /etc/drbd.d/haqm1.res.

The most important information for troubleshooting purposes in the .res file is the device minor number for this particular DRBD resource. Many of the messages that DRBD logs use this minor number. For the example queue manager, HAQM1, the .res file contains the following information:

```
device minor 100;
```

For this queue manager, you should look for messages such as the following example:

```
Jul 31 00:17:24 mqhavm13 kernel: drbd haqm1/0 drbd100 mqhavm15.gamsworthwilliam.com: drbd_sync_handshake:
```

The presence of the string drbd100 indicates that the message relates to HAQM1. Not all messages logged by DRBD use the device minor number, some use the DRBD resource name, which is the same as the Directory name of the RDQM HA queue manager. For example:

```
Jul 31 00:17:22 mqhavm13 kernel: drbd haqm1 mqhavm15.gamsworthwilliam.com: Connection closed
```

Pacemaker

There are a number of Pacemaker resources generated for an RDOM HA queue manager:

qm

Is the main resource representing the RDQM HA queue manager.

p_rdqmx_*qm*

Is an internal resource.

p_fs_qm

Is a standard filesystem resource that mounts the volume for the queue manager on /var/mqm/vols/am.

ms_drbd_qm

Is the master/slave resource for the DRBD resource for the RDQM.

p_drbd_qm

Is the primitive resource for the DRBD resource for the RDQM.

If a floating IP address is configured for an HA RDQM then an additional resource is configured:

p_ip_qm

Example RDQM HA configurations and errors

An example RDQM HA configuration, complete with example errors and information on how to resolve them.

The example RDQM HA group consists of three nodes:

- mqhavm13.gamsworthwilliam.com (referred to as vm13).
- mqhavm14.gamsworthwilliam.com (referred to as vm14).
- mqhavm15.gamsworthwilliam.com (referred to as vm15).

Three RDQM HA queue managers have been created:

- HAQM1 (created on vm13)
- HAQM2 (created on vm14)
- HAQM3 (created on vm15)

Initial conditions

The initial condition on each of the nodes is given in the following listings:

vm13

```
[midtownjojo@mqhavm13 ~]$ rdqmstatus -m HAQM1
                                            mqhavm13.gamsworthwilliam.com
Queue manager status:
                                            Running
CPU:
                                            0.00%
Memory:
                                            135MB
Queue manager file system:
                                            51MB used, 1.0GB allocated [5%]
HA role:
                                            Primary
HA status:
                                            Normal
HA control:
                                            Enabled
HA current location:
                                            This node
HA preferred location:
HA floating IP interface:
HA floating IP address:
                                            This node
                                            None
                                            None
                                            mqhavm14.gamsworthwilliam.com
HA status:
                                            Normal
Node:
                                            mqhavm15.gamsworthwilliam.com
HA status:
                                            Normal
Command '/opt/mgm/bin/rdgmstatus' run with sudo.
[midtownjojo@mqhavm13 ~]$ rdqmstatus -m HAQM2
                                            mqhavm13.gamsworthwilliam.com
Node:
Queue manager status:
                                            Running elsewhere
HA role:
                                            Secondary
HA status:
                                            Normal
HA control:
                                            Enabled
HA current location:
                                            mqhavm14.gamsworthwilliam.com
HA preferred location:
HA floating IP interface:
                                            mqhavm14.gamsworthwilliam.com
                                            None
HA floating IP address:
                                            None
Node:
                                            mghavm14.gamsworthwilliam.com
HA status:
                                            Normal
Node:
                                            mqhavm15.gamsworthwilliam.com
HA status:
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
[midtownjojo@mqhavm13 ~]$ rdqmstatus -m HAQM3
                                            mqhavm13.gamsworthwilliam.com
Queue manager status:
                                            Running elsewhere
HA role:
                                            Secondary
```

```
HA status:
                                            Normal
HA control:
                                            Enabled
HA current location:
                                            mqhavm15.gamsworthwilliam.com
HA preferred location:
                                            mqhavm15.gamsworthwilliam.com
HA floating IP interface:
HA floating IP address:
                                            None
                                            None
Node:
                                            mqhavm14.gamsworthwilliam.com
HA status:
                                            Normal
Node:
                                            mqhavm15.gamsworthwilliam.com
HA status:
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

vm14

```
[midtownjojo@mqhavm14 ~]$ rdqmstatus -m HAQM1
                                         mqhavm14.gamsworthwilliam.com
Queue manager status:
                                         Running elsewhere
HA role:
                                         Secondary
HA status:
                                         Normal
HA control:
                                         Enabled
HA current location:
                                         mqhavm13.gamsworthwilliam.com
HA preferred location:
                                         mqhavm13.gamsworthwilliam.com
HA floating IP interface:
                                         None
HA floating IP address:
                                         None
Node:
                                         mqhavm13.gamsworthwilliam.com
HA status:
                                         Normal
Node:
                                         mqhavm15.gamsworthwilliam.com
HA status:
                                         Normal
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
[midtownjojo@mghavm14 ~]$ rdgmstatus -m HAQM2
                                         mqhavm14.gamsworthwilliam.com
Node:
Queue manager status:
                                         Running
CPU:
                                         0.00%
Memory:
                                         135MB
Queue manager file system:
                                         51MB used, 1.0GB allocated [5%]
HA role:
                                         Primary
HA status:
                                         Normal
HA control:
                                         Enabled
HA current location:
                                         This node
HA preferred location:
                                         This node
HA floating IP interface:
                                         None
HA floating IP address:
                                         None
Node:
                                         mqhavm13.gamsworthwilliam.com
HA status:
                                         Normal
Node:
                                         mqhavm15.gamsworthwilliam.com
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
[midtownjojo@mqhavm14 ~]$ rdqmstatus -m HAQM3
Node:
                                         mqhavm14.gamsworthwilliam.com
Queue manager status:
                                         Running elsewhere
HA role:
                                         Secondary
                                         Normal
HA status:
HA control:
                                         Enabled
                                         mqhavm15.gamsworthwilliam.com
HA current location:
HA preferred location:
                                         mqhavm15.gamsworthwilliam.com
HA floating IP interface:
                                         None
HA floating IP address:
                                         None
Node:
                                         mqhavm13.gamsworthwilliam.com
HA status:
Node:
                                         mqhavm15.gamsworthwilliam.com
                                         Normal
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

vm15

```
HA control:
                                          Enabled
HA current location:
                                          mqhavm13.gamsworthwilliam.com
HA preferred location:
                                          mqhavm13.gamsworthwilliam.com
HA floating IP interface:
                                          None
HA floating IP address:
                                          None
                                          mqhavm13.gamsworthwilliam.com
HA status:
                                          Normal
Node:
                                          mqhavm14.gamsworthwilliam.com
HA status:
                                          Normal
Command '/opt/mgm/bin/rdgmstatus' run with sudo.
[midtownjojo@mqhavm15 ~]$ rdqmstatus -m HAQM2
                                          mqhavm15.gamsworthwilliam.com
Queue manager status:
                                          Running elsewhere
HA role:
                                          Secondary
HA status:
                                          Normal
                                          Enabled
HA control:
HA current location:
                                          mqhavm14.gamsworthwilliam.com
HA preferred location: HA floating IP interface:
                                          mghavm14.gamsworthwilliam.com
                                          None
HA floating IP address:
                                          None
Node:
                                          mqhavm13.gamsworthwilliam.com
HA status:
                                          Normal
Node:
                                          mqhavm14.gamsworthwilliam.com
HA status:
                                          Normal
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
[midtownjojo@mqhavm15 ~]$ rdqmstatus -m HAQM3
                                          mqhavm15.gamsworthwilliam.com
Node:
Queue manager status:
                                          Running
ČPU:
                                          0.02%
Memory:
Queue manager file system:
                                          135MB
                                          51MB used, 1.0GB allocated [5%]
HA role:
                                          Primary
HA status:
                                          Normal
HA control:
                                          Enabled
HA current location:
                                          This node
HA preferred location:
                                          This node
HA floating IP interface:
                                          None
HA floating IP address:
                                          None
Node:
                                          mqhavm13.gamsworthwilliam.com
HA status:
                                          Normal
                                          mqhavm14.gamsworthwilliam.com
HA status:
                                          Normal
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

DRBD scenarios

RDQM HA configurations use DRBD for data replication. The following scenarios illustrate the following possible problems with DRBD:

- Loss of DRBD quorum
- · Loss of a single DRBD connection
- Synchronization stuck

DRBD Scenario 1: Loss of DRBD quorum

If the node running an RDQM HA queue manager loses the DRBD quorum for the DRBD resource corresponding to the queue manager, DRBD immediately starts returning errors from I/O operations, which will cause the queue manager to start producing FDCs and eventually stop.

If the remaining two nodes have a DRBD quorum for the DRBD resource then Pacemaker chooses one of the two nodes to start the queue manager. Because there were no updates on the original node from the time where the quorum was lost, it is safe to start the queue manager somewhere else.

The two main ways that you can monitor for a loss of DRBD quorum are:

- By using the **rdqmstatus** command.
- By monitoring the syslog of the node where the RDOM HA queue manager is initially running.

rdqmstatus

If you use the **rdqmstatus** command, if the node vm13 loses DRBD quorum for the DRBD resource for HAQM1, you might see status similar to the following example:

```
[midtownjojo@mghavm13 ~]$ rdgmstatus -m HAQM1
                                          mqhavm13.gamsworthwilliam.com
Running elsewhere
Node:
Queue manager status:
HA role:
                                          Secondary
HA status:
                                          Remote unavailable
HA control:
                                          Enabled
HA current location:
                                          mqhavm14.gamsworthwilliam.com
HA preferred location:
                                          This node
HA floating IP interface:
                                          None
HA floating IP address:
                                          None
                                          mqhavm14.gamsworthwilliam.com
Node:
HA status:
                                          Remote unavailable
HA out of sync data:
                                          0KB
                                          mqhavm15.gamsworthwilliam.com
HA status:
                                          Remote unavailable
HA out of sync data:
                                          0KB
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

Notice that the HA status has changed to Remote unavailable, which indicates that both DRBD connections to the other nodes have been lost.

In this case the other two nodes have DRBD quorum for the DRBD resource so the RDQM is running somewhere else, on mqhavm14.gamsworthwilliam.com as shown as the value of HA current location.

monitoring syslog

If you monitor syslog, you will see that DRBD logs a message when it loses quorum for a resource:

```
Jul 30 09:38:36 mqhavm13 kernel: drbd haqm1/0 drbd100: quorum( yes -> no )
```

When quorum is restored a similar message is logged:

```
Jul 30 10:27:32 mqhavm13 kernel: drbd haqm1/0 drbd100: quorum( no -> yes )
```

DRBD Scenario 2: Loss of a single DRBD connection

If only one of the two DRBD connections from a node running an RDQM HA queue manager is lost then the queue manager does not move.

Starting from the same initial conditions as in the first scenario, after blocking just one of the DRBD replication links, the status reported by **rdqmstatus** on vm13 is similar to the following example:

```
mqhavm13.gamsworthwilliam.com
Queue manager status:
                                         Running
CPU:
                                         0.01%
Memory:
                                         133MB
Queue manager file system:
                                         52MB used, 1.0GB allocated [5%]
HA role:
                                         Primary
HA status:
                                         Mixed
HA control:
                                         Enabled
HA current location:
                                         This node
HA preferred location:
                                         This node
HA floating IP interface:
                                         None
HA floating IP address:
                                         None
Node:
                                         mqhavm14.gamsworthwilliam.com
HA status:
                                         Remote unavailable
HA out of sync data:
                                         OKR
Node:
                                         mqhavm15.gamsworthwilliam.com
```

```
HA status:
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

DRBD Scenario 3: Synchronization stuck

Some versions of DRBD had an issue where a synchronization would appear to be stuck and this prevented an RDQM HA queue manager from failing over to a node when the sync to that node is still in progress.

One way to see this is to use the drbdadm status command. When operating normally a response similar to the following example is output:

```
[midtownjojo@mqhavm13 ~]$ drbdadm status
hagm1 role:Primary
  disk:UpToDate
  mqhavm14.gamsworthwilliam.com role:Secondary
   peer-disk:UpToDate
  mghavm15.gamsworthwilliam.com role:Secondary
   peer-disk:UpToDate
haqm2 role:Secondary
  disk:UpToDate
  mghavm14.gamsworthwilliam.com role:Primary
    peer-disk:UpToDate
 mqhavm15.gamsworthwilliam.com role:Secondary
   peer-disk:UpToDate
haqm3 role:Secondary
  disk:UpToDate
  mqhavm14.gamsworthwilliam.com role:Secondary
    peer-disk:UpToDate
  mqhavm15.gamsworthwilliam.com role:Primary
   peer-disk:UpToDate
```

If synchronization gets stuck, the response is similar to the following example:

```
[midtownjojo@mqhavm13 ~]$ drbdadm status
haqm1 role:Primary
  disk:UpToDate
  mqhavm14.gamsworthwilliam.com role:Secondary
   peer-disk:UpToDate
  mqhavm15.gamsworthwilliam.com role:Secondary
   replication:SyncSource peer-disk:Inconsistent done:90.91
hagm2 role:Secondary
  disk:UpToDate
  mqhavm14.gamsworthwilliam.com role:Primary
   peer-disk:UpToDate
 mqhavm15.gamsworthwilliam.com role:Secondary
   peer-disk:UpToDate
haqm3 role:Secondary
  disk:UpToDate
  mqhavm14.gamsworthwilliam.com role:Secondary
    peer-disk:UpToDate
  mqhavm15.gamsworthwilliam.com role:Primary
    peer-disk:UpToDate
```

In this case the RDQM HA queue manager HAQM1 cannot move to vm15 as the disk on vm15 is Inconsistent.

The done value is the percentage complete. If that value is not increasing you could try disconnecting that replica then connecting it again with the following commands (run as root) on vm13:

```
drbdadm disconnect haqm1:mqhavm15.gamsworthwilliam.com
drbdadm connect haqm1:mqhavm15.gamsworthwilliam.com
```

If the replication to both Secondary nodes is stuck, you can do the **disconnect** and **connect** commands without specifying a node and that will disconnect both connections:

```
drbdadm disconnect haqm1
drbdadm connect haqm1
```

Pacemaker scenarios

RDQM HA configurations use Pacemaker to determine where an RDQM HA queue manager runs. The following scenarios illustrate the following possible problems that involve Pacemaker:

- Corosync main process not scheduled
- RDQM HA queue manager not running where it should

Pacemaker scenario 1: Corosync main process not scheduled

If you see a message in the syslog similar to the following example this indicates that the system is either too busy to schedule CPU time to the main Corosync process or, more commonly, that the system is a Virtual Machine and the Hypervisor has not scheduled any CPU time to the entire VM.

```
corosync[10800]: [MAIN ] Corosync main process was not scheduled for 2787.0891 ms (threshold is 1320.0000 ms). Consider token timeout increase.
```

Both Pacemaker (and Corosync) and DRBD have timers that are used to detect loss of quorum, so messages like the example indicate that the node did not run for so long that it would have been dropped from the quorum. The Corosync timeout is 1.65 seconds and the threshold of 1.32 seconds is 80% of that, so the message shown in the example is printed when the delay in the scheduling of the main Corosync process hits 80% of the timeout. In the example the process was not scheduled for nearly three seconds. Whatever is causing such a problem must be resolved. One thing that might help in a similar situation is to reduce the requirements of the VM, for example, reducing the number of vCPUs required, as this makes it easier for the Hypervisor to schedule the VM.

Pacemaker scenario 2: An RDQM HA queue manager is not running where it should be

The main tool to help troubleshooting in this scenario is the **crm status** command. The following example shows a response for the configuration when everything is working as expected:

```
Stack: corosvnc
Current DC: mqhavm13.gamsworthwilliam.com (version 1.1.20.linbit-1+20190404+eab6a2092b71.el7.2-
eab6a2092b) - partition with quorum
Last updated: Tue Jul 30 09:11:29 2019
Last change: Tue Jul 30 09:10:34 2019 by root via crm_attribute on mqhavm14.gamsworthwilliam.com
3 nodes configured
18 resources configured
Online: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com
mqhavm15.gamsworthwilliam.com ]
Full list of resources:
 Master/Slave Set: ms_drbd_haqm1 [p_drbd_haqm1]
     Masters: [ mqhavm13.gamsworthwilliam.com ]
Slaves: [ mqhavm14.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com ]
                (ocf::heartbeat:Filesystem):
                                                     Started mqhavm13.gamsworthwilliam.com
 p_fs_haqm1
                                           Started mqhavm13.gamsworthwilliam.com
 p_rdqmx_haqm1
                    (ocf::ibm:rdqmx):
           (ocf::ibm:rdqm):
                                 Started mqhavm13.gamsworthwilliam.com
Master/Slave Set: ms_drbd_haqm2 [p_drbd_haqm2]
Masters: [ mqhavm14.gamsworthwilliam.com ]
     Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com
 p_fs_haqm2
                (ocf::heartbeat:Filesystem):
                                                     Started mqhavm14.gamsworthwilliam.com
 p_rdqmx_haqm2
                    (ocf::ibm:rdqmx):
                                           Started mghavm14.gamsworthwilliam.com
                                 Started mqhavm14.gamsworthwilliam.com
           (ocf::ibm:rdqm):
 haam2
Master/Slave Set: ms_drbd_haqm3 [p_drbd_haqm3]
Masters: [ mqhavm15.gamsworthwilliam.com ]
     Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com
                (ocf::heartbeat:Filesystem):
                                                    Started mghavm15.gamsworthwilliam.com
 p fs haqm3
                                           Started mqhavm15.gamsworthwilliam.com
 p_rdqmx_haqm3
                    (ocf::ibm:rdqmx):
                                 Started mqhavm15.gamsworthwilliam.com
 haam3
           (ocf::ibm:rdqm):
```

Note the following points:

• All three nodes are shown as Online.

• Each RDQM HA queue manager is running on the node where it was created, for example, HAQM1 is running on vm13 and so on.

This scenario is constructed by preventing HAQM1 from running on vm14, and then attempting to move HAQM1 to vm14. HAQM1 cannot run on vm14 because the file /var/mqm/mqs.ini on vm14 has an invalid value for the Directory of queue manager HAQM1.

The preferred location for HAQM1 is changed to vm14 by running the following command on vm13:

```
rdqmadm -m HAQM1 -n mqhavm14.gamsworthwilliam.com -p
```

This command would normally cause HAQM1 to move to vm14 but in this case checking the status on vm13 returns the following information:

```
[midtonjojo@mqhavm13 ~]$ rdqmstatus -m HAQM1
Node:
                                         mghavm13.gamsworthwilliam.com
                                         Running
Queue manager status:
CPU:
                                         0.15%
                                         133MB
Memory:
Queue manager file system:
                                         52MB used, 1.0GB allocated [5%]
HA role:
                                         Primary
HA status:
                                         Normal
                                         Enabled
HA control:
HA current location:
                                         This node
HA preferred location:
                                         mqhavm14.gamsworthwilliam.com
HA floating IP interface:
                                         None
HA floating IP address:
                                         None
                                         mqhavm14.gamsworthwilliam.com
Node:
HA status:
                                         Normal
Node:
                                         mghavm15.gamsworthwilliam.com
HA status:
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

HAQM1 is still running on vm13, it has not moved to vm14 as requested and the cause needs investigating. Examining the Pacemaker status gives the following response:

```
[midtownjojo@mqhavm13 ~]$ crm status
Stack: corosync
Current DC: mqhavm13.gamsworthwilliam.com (version 1.1.20.linbit-1+20190404+eab6a2092b71.el7.2-
eab6a2092b) - partition with quorum
Last updated: Thu Aug 1 14:16:40 2019
Last change: Thu Aug 1 14:16:35 2019 by hacluster via crmd on mqhavm14.gamsworthwilliam.com
3 nodes configured
18 resources configured
Online: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com
mqhavm15.gamsworthwilliam.com ]
Full list of resources:
 Master/Slave Set: ms_drbd_haqm1 [p_drbd_haqm1]
     Masters: [ mqhavm13.gamsworthwilliam.com
     Slaves: [ mqhavm14.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com ]
p_rs_haqm1 (ocf::heartbeat:Filesystem):
p_rdqmx_haqm1 (ocf::ibm:rdamx):
                                                 Started mqhavm13.gamsworthwilliam.com
                                        Started mqhavm13.gamsworthwilliam.com
          (ocf::ibm:rdqm):
                               Started mqhavm13.gamsworthwilliam.com
 Master/Slave Set: ms_drbd_haqm2 [p_drbd_haqm2]
     Masters: [ mqhavm14.gamsworthwilliam.com ]
     Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com ]
 p_fs_haqm2
               (ocf::heartbeat:Filesystem):
                                                 Started mqhavm14.gamsworthwilliam.com
 p_rdqmx_haqm2
                   (ocf::ibm:rdqmx):
                                        Started mqhavm14.gamsworthwilliam.com
Masters: [ mqhavm15.gamsworthwilliam.com ]
μ_Is_haqm3 (ocf::heartbeat:Filesystem):
p_rdqmx_haqm3 (ocf::ibm:rdqmx):
     Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com ]
                                                Started mghavm15.gamsworthwilliam.com
                                        Started mqhavm15.gamsworthwilliam.com
                               Started mqhavm15.gamsworthwilliam.com
Failed Resource Actions:
* haqm1_monitor_0 on mqhavm14.gamsworthwilliam.com 'not installed' (5): call=372,
status=complete, exitreason='
    last-rc-change='Thu Aug 1 14:16:37 2019', queued=0ms, exec=17ms
```

Take note of the Failed Resource Actions section that has appeared.

The name of the action, haqm1_monitor_0 tells us that it was a monitor action for the RDQM HAQM1 that failed, and it failed on mqhavm14.gamsworthwilliam.com, so it looks like Pacemaker tried to do what we expected and start HAQM1 on vm14, but for some reason it could not.

You can see when Pacemaker tried do this from the value of last-rc-change

Understanding the failure

To understand the failure we need to look at the syslog for vm14 at the time of the failure:

```
Aug 1 14:16:37 mqhavm14 crmd[26377]: notice: Result of probe operation for haqm1 on mqhavm14.gamsworthwilliam.com: 5 (not installed)
```

The entry shows that when Pacemaker tried to check the state of haqm1 on vm14 it got an error because haqm1 is not configured, which is because of the deliberate misconfiguration in /var/mqm/mqs.ini.

Correcting the failure

To correct the failure you must correct the underlying problem (in this case restoring the correct directory value for haqm1 in /var/mqm/mqs.ini on vm14). Then you must clear the failed action by using the command **crm resource cleanup** on the appropriate resource, which in this case is the resource haqm1 as that is the resource mentioned in the failed action. For example:

```
[midtownjojo@mqhavm13 ~]$ crm resource cleanup haqm1
Cleaned up haqm1 on mqhavm15.gamsworthwilliam.com
Cleaned up haqm1 on mqhavm14.gamsworthwilliam.com
Cleaned up haqm1 on mqhavm13.gamsworthwilliam.com
```

Then check the Pacemaker status again:

```
[midtownjojo@mghavm13 ~]$ crm status
Stack: corosync
Current DC: mqhavm13.gamsworthwilliam.com (version 1.1.20.linbit-1+20190404+eab6a2092b71.el7.2-
eab6a2092b) - partition with quorum
Last updated: Thu Aug 1 14:23:17 2019
Last change: Thu Aug 1 14:23:03 2019 by hacluster via crmd on mqhavm13.gamsworthwilliam.com
3 nodes configured
18 resources configured
Online: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com
mqhavm15.gamsworthwilliam.com ]
Full list of resources:
 Master/Slave Set: ms_drbd_haqm1 [p_drbd_haqm1]
     Masters: [ mqhavm14.gamsworthwilliam.com ]
Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com ]
                 (ocf::heartbeat:Filesystem):
                                                      Started mqhavm14.gamsworthwilliam.com
 p_fs_haqm1
 p_rdqmx_haqm1
                     (ocf::ibm:rdqmx):
                                            Started mqhavm14.gamsworthwilliam.com
           (ocf::ibm:rdqm):
                                  Started mqhavm14.gamsworthwilliam.com
 Master/Slave Set: ms_drbd_haqm2 [p_drbd_haqm2]
Masters: [ mqhavm14.gamsworthwilliam.com ]
 p_ts_haqm2 (ocf::heartbeat:Filesystem):
p_rdqmx_haqm2 (ocf::ibm:rdamv):
      Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm15.gamsworthwilliam.com ]
                                                       Started mghavm14.gamsworthwilliam.com
                                             Started mqhavm14.gamsworthwilliam.com
           (ocf::ibm:rdqm):
                                   Started mqhavm14.gamsworthwilliam.com
 haam2
 Master/Slave Set: ms_drbd_haqm3 [p_drbd_haqm3]
      Masters: [ mqhavm15.gamsworthwilliam.com ]
 p_ts_haqm3 (ocf::heartbeat:Filesystem):
p_rdqmx_haqm3 (ocf::ihm:rdcmv):
      Slaves: [ mqhavm13.gamsworthwilliam.com mqhavm14.gamsworthwilliam.com ]
                                                      Started mqhavm15.gamsworthwilliam.com
                                             Started mqhavm15.gamsworthwilliam.com
                                   Started mqhavm15.gamsworthwilliam.com
            (ocf::ibm:rdqm):
```

The failed action has disappeared and HAQM1 is now running on vm14 as expected. The following example shows the RDQM status:

```
[midtownjojo@mqhavm13 ~]$ rdqmstatus -m HAQM1
Node: mqhavm13.gamsworthwilliam.com
```

```
Queue manager status:
                                         Running elsewhere
HA role:
                                         Secondary
HA status:
                                         Normal
HA control:
                                         Enabled
HA current location:
                                         mqhavm14.gamsworthwilliam.com
HA preferred location:
                                         mqhavm14.gamsworthwilliam.com
HA floating IP interface:
HA floating IP address:
                                         None
Node:
                                         mqhavm14.gamsworthwilliam.com
HA status:
                                         Normal
Node:
                                         mqhavm15.gamsworthwilliam.com
HA status:
                                         Normal
Command '/opt/mqm/bin/rdqmstatus' run with sudo.
```

Windows Linux AIX MQ Telemetry troubleshooting

Look for a troubleshooting task to help you solve a problem with running MQ Telemetry applications. **Related concepts**

MQ Telemetry

Location of telemetry logs, error logs, and configuration files

Find the logs, error logs, and configuration files used by 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 telemetry (MQXR) service writes FDC files to the IBM MQ error directory:

WMQ data directory\errors\AMQ nnn.n.FDC

The format of the FDC files is MQXRn.FDC.

It also writes a log for the telemetry (MQXR) service. The log path is:

```
WMQ data directory\Qmgrs\qMgrName\errors\mqxr.log
```

The format of the log file is mqxr_n.log.

The IBM MQ telemetry sample configuration created by IBM MQ Explorer starts the telemetry (MQXR) service using the command **runMQXRService**, which is in WMQ Telemetry installation directory\bin. This command writes to:

```
WMQ data directory \Qmgrs\qMgrName \mqxr.stdout \\ WMQ data directory \Qmgrs\qMgrName \mqxr.stderr
```

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 IBM MQ Explorer, or MQSC commands, to configure telemetry channels.

IBM 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 18. Telemetry configuration directory on Windows

/var/mqm/qmgrs/qMgrName/mqxr

Figure 19. 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 <u>Figure 18 on</u> page 231 and Figure 19 on page 231.

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
```

JSSE

Enable a JSSE debug trace for the MQTT service, so you can see the TLS handshake.

Do this by adding the following property to the java.properties file, as previously described in the "JVM" on page 231 section.

```
javax.net.debug=<value>
```

where *<value>* is one or more of all, true, or ssl, optionally followed by debug specifiers. For more information about these options, see Debugging Utilities in the IBMJSSE2 provider documentation.

The output is sent to the maxr.stdout file.

JAAS

The JAAS configuration file is described in <u>Telemetry channel JAAS configuration</u>, which includes the sample JAAS configuration file, JAAS.config, shipped with 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 an IBM 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
```

⁻Xms1024m

```
CLASSPATH= WMQ Installation Directory\mqxr\samples\samples
```

Note: service.env must not contain any variables. Substitute the actual value of WMQ Installation Directory.

Figure 20. Sample service.env for Windows

Trace

See "Tracing the telemetry (MQXR) service" on page 232. The parameters to configure trace are stored in two files:

```
WMQ data directory\Qmgrs\qMgrName\mqxr\trace.config
WMQ data directory\Qmgrs\qMgrName\mqxr\mqxrtraceOn.properties
```

and there is a corresponding file:

WMQ data directory\Qmgrs\qMgrName\mqxr\mqxrtraceOff.properties

Client-side log files and client-side configuration files

For the latest information and downloads, see the following resources:

- The <u>Eclipse Paho</u> project, and <u>MQTT.org</u>, have free downloads of the latest telemetry clients and samples for a range of programming languages. Use these sites to help you develop sample programs for publishing and subscribing IBM MQ Telemetry Transport, and for adding security features.
- The IBM Messaging Telemetry Clients SupportPac is no longer available for download. If you have a previously downloaded copy, it has the following contents:
 - The MA9B version of the IBM Messaging Telemetry Clients SupportPac included a compiled sample application (mqttv3app.jar) and associated client library (mqttv3.jar). They were provided in the following directories:
 - ma9b/SDK/clients/java/org.eclipse.paho.sample.mqttv3app.jar
 - ma9b/SDK/clients/java/org.eclipse.paho.client.mqttv3.jar
 - In the MA9C version of this SupportPac, the /SDK/ directory and contents was removed:
 - Only the source for the sample application (mqttv3app.jar) was provided. It was in this directory:

ma9c/clients/java/samples/org/eclipse/paho/sample/mqttv3app/*.java

- The compiled client library was still provided. It was in this directory:

ma9c/clients/java/org.eclipse.paho.client.mqttv3-1.0.2.jar

Windows Linux AIX Tracing the telemetry (MQXR) service

The trace facility provided by the IBM MQ telemetry (MQXR) service is provided to help IBM Support diagnose customer issues related to the service.

About this task

There are two ways to control trace for the IBM MO telemetry service:

By using the strmqtrc and endmqtrc commands to start and stop trace. Enabling trace, using the strmqtrc command, generates trace information for the entire queue manager where the IBM MQ telemetry service is running. This includes the IBM MQ telemetry service itself, and the underlying Java Message Queuing Interface (JMQI) that the service uses to communicate with other queue manager components.

• By running the **controlMQXRChannel** command. Note, that turning trace on using the **controlMQXRChannel** command traces only the IBM MQ telemetry service.

If you are unsure which option to use, contact your IBM Support representative and they will be able to advise you on the best way to collect trace for the issue that you are seeing.

Procedure

- 1. Method one
 - a) Bring up a command prompt and navigate to the directory: MQ_INSTALLATION_PATH\bin
 - b) Run the **strmqtrc** command to enable trace.

```
strmqtrc -m qmgr_name
```

where qmgr_name is the name of the queue manager where the IBM MQ MQXR service is running.

- c) Reproduce the issue.
- d) Stop trace, by running the command:

```
endmqtrc -m qmgr_name
```

- 2. Method two.
 - a) Bring up a command prompt and navigate to the directory:
 MQ_INSTALLATION_PATH\mqxx\bin
 - b) Run the following command to enable trace:
 - . Windows

```
\verb|controlMQXRC| controlMQXRC| controlMQXRC
```

. Linux VINIX

```
./{\tt controlMQXRChannel.sh-qmgr=qmgr\_name} \ -{\tt mode=starttrace} \ \ [{\tt clientid=ClientIdentifier}]
```

where *qmgr_name* is the name of the queue manager where the MQXR Service is running. Set *ClientIdentifier* to the client identifier of an MQTT client. If you specify the **clientid** parameter, the IBM MQ telemetry service trace captures activity for only the MQTT client with that client identifier.

If you want to trace the IBM MQ telemetry service activity for more than one specific MQTT client, you can run the command multiple times, specifying a different client identifier each time.

- c) Reproduce the issue.
- d) When the issue occurs, stop trace by running the following command:
 - controlMQXRChannel -qmgr=qmgr_name -mode=stoptrace

```
. Linux VINIX
```

```
./controlMQXRChannel.sh -qmgr=qmgr_name -mode=stoptrace [clientid=ClientIdentifier]
```

where amgr name is the name of the queue manager where the MOXR Service is running.

Results

To view the trace output, go to the following directory:

- Windows MQ_DATA_PATH\trace.
- Linux | UNIX /var/mqm/trace.

The trace files containing the information from the MQXR service are called $mqxr_N.trc$, where N is a number.

Trace information generated by the JMQI is written to a trace file called mqxr_PPPPP.trc, where PPPPP is the process identifier for the MQXR Service.

Related reference

strmqtrc

Windows Linux AIX Additional diagnostics using the controlMOXRChannel command

Using the **controlMQXRChannel** command to provide additional diagnostic information about the MQXR service.

Procedure

Run the following command to provide useful diagnostic information from the MQXR service:

 $< MQ_INSTALLATION_PATH> \\ mqxx\\bin\\controlMQXRChannel -qmgr=< QMGR_NAME> -mode= \\ diagnostics -diagnostics type= < number>$

The diagnostic information generated depends on the value of the **-diagnosticstype**=<number> parameter:

-diagnosticstype= 0

Thread dump written to the console

-diagnosticstype= 1

FDC with some internal service statistics

-diagnosticstype= 2

FDC with internal statistics, plus information about the clients that are currently connected

-diagnosticstype= 3

Heap dump

-diagnosticstype= 4

Javacore

-diagnosticstype= 5

Full system dump

-diagnosticstype= 6

Detailed information about a specific client. Note that you must also supply the **-clientid** parameter for that client as well.

Windows Linux AIX Resolving problem: MQTT client does not

connect

Resolve the problem of an MQTT client program failing to connect to the telemetry (MQXR) 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 IBM MQTT clients?

See <u>Verifying the installation of MQ Telemetry</u> for further information, and check that the telemetry channel and telemetry (MQXR) service are running correctly.

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 (MQXR) 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.
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 (MQXR) 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 MQ Telemetry, look at the CONNACK return code.

From these three errors you can infer that the client has connected to the telemetry (MQXR) 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 (MQXR) service does not respond:

Option	Description
REASON_CODE_CLIENT_EXCEPTION REASON_CODE_CLIENT_TIMEOUT	Look for an FDC file at the server; see <u>"Server-side logs"</u> on page 230. When the telemetry (MQXR) 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 (MQXR) service might not have responded to the client, and the timeout at the client expires. The 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 TCPIP request with a tool such as **tcpmon**, available from (for example) https://code.google.com/p/tcpmon/
- 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 IBM MQ Telemetry creates a folder with the name: <code>clientIdentifier-tcphostNameport</code> or <code>clientIdentifier-sslhostNameport</code> in the client working directory. The folder name tells you the hostName and port used in the connection attempt; see "Client-side log files and client-side configuration files" on page 232.

- 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 (MQXR) service found a problem in the client request.

The telemetry (MQXR) service writes errors it detects into mqxr_n.log, and the queue manager writes errors into AMQERR01.LOG; see

- 4. Attempt to isolate the problem by running another client.
 - See Verifying the installation of MQ Telemetry for further information

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 IBM MQ Explorer. Its default size is 4096. Check that the backlog has not been configured to a low value.
 - b) Are the telemetry (MQXR) 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?



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 IBM 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.

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 230.
- 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.



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. Re-create 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 Support.

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 IBM 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 and client-side configuration files" on page 232
 - 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.

6. Check that the client has a subscription for the publication it expects to receive.

List subscriptions using IBM 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 IBM MQ Explorer. Simulate your application client publishing and subscribing by using the client utility. Start the utility from IBM 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 IBM 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.



does not start

Resolve the problem of the telemetry (MQXR) service failing to start. Check the MQ Telemetry installation and no files are missing, moved, or have the wrong permissions. Check the paths that are used by the telemetry (MQXR) service locate the telemetry (MQXR) service programs.

Before you begin

The MQ Telemetry feature is installed. The IBM MQ Explorer has a Telemetry folder in **IBM MQ > Queue Managers > qMgrName > Telemetry**. If the folder does not exist, the installation has failed.

The Telemetry (MQXR) service must have been created for it to start. If the telemetry (MQXR) service has not been created, then run the **Define sample configuration...** wizard in the Telemetry folder.

If the telemetry (MQXR) 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 IBM MQ 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 (MQXR) 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:

```
IBM 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 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 (MQXR) service is not installed on the standby.

Investigation

Look in error logs; see "Server-side logs" on page 230.

Actions

Install, or uninstall and reinstall the MQ Telemetry feature.

2. Start the service; wait for 30 seconds; refresh the IBM MQ Explorer and check the service status.

Result

The service starts and then stops.

Reasor

SYSTEM. MQXR. SERVICE started the runMQXRService command, but the command failed.

Investigation

Look in error logs; see "Server-side logs" on page 230.

See if the problem occurs with only the sample channel defined. Back up 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 (MQXR) service, and configure JAAS to correct the problem.

Before you begin

You have modified WMQ installation directory\mqxr\samples\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 (MQXR) service, you suspect that your authentication class is not being called by the telemetry (MQXR) service.

Note: Guard against the possibility that your authentication classes might be overwritten by maintenance being applied to IBM MQ. Use your own path for authentication classes, rather than a path within the IBM 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 and AuthCallback MQXR class for help in configuring JAAS for MQ Telemetry. The example, "Example JAAS configuration" on page 241 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 230 for the path to mqxr.log, and Figure 26 on page 242 for an example of the exception listed in the log.
- 2. Correct your JAAS configuration by comparing it with the worked example in <u>"Example JAAS</u> configuration" on page 241.
- 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 mgm 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 IBM MQ Explorer, or **dspmqauth**.

Example JAAS configuration

Use the New telemetry channel wizard, in IBM MQ Explorer, to configure a telemetry channel.

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 21. WMQ Installation directory\data\qmgrs\qMgrName\mqxr\jaas.config

When SYSTEM. MQTT. SERVICE starts, it adds the path in Figure 22 on page 241 to its classpath.

```
CLASSPATH=C:\WMQTelemtryApps;
```

Figure 22. WMQ Installation directory\data\qmgrs\qMgrName\service.env

<u>Figure 23 on page 241</u> shows the additional path in <u>Figure 22 on page 241</u> added to the classpath that is set up for the telemetry (MQXR) service.

```
CLASSPATH=;C:\IBM\MQ\Program\mqxr\bin\\..\lib\MQXRListener.jar;
C:\IBM\MQ\Program\mqxr\bin\\..\lib\bbjectManager.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.mqjms.jar;
C:\IBM\MQ\Program\mqxr\bin\\..\.\java\lib\com.ibm.mq.jar;
C:\IBM\MQ\Program\mqxr\bin\\..\.\java\lib\com.ibm.mq.jar;
C:\IBM\MQ\Program\mqxr\bin\\..\.\java\lib\com.ibm.mq.jar;
```

Figure 23. Classpath output from runMQXRService.bat

The output in Figure 24 on page 241 shows that the telemetry (MQXR) service has started.

```
21/05/2010 15:32:12 [main] com.ibm.mq.MQXRService.MQXRPropertiesFile
AMQXR2011I: Property com.ibm.mq.MQXR.channel/JAASMCAUser 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 24. WMQ Installation directory\data\qmgrs\qMgrName\errors\

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 25 on page 242. The second exception, Client is not connected (32104), was thrown because the client attempted to disconnect when it was not connected.

```
Connecting to tcp://localhost:1883 with client ID SampleJavaV3_publish
reason 5
msg Not authorized to connect
loc Not authorized to connect
cause null
excep Not authorized to connect (5)
Not authorized to connect (5)
at
org.eclipse.paho.client.mqttv3.internal.ExceptionHelper.createMqttException(ExceptionHelper.java:28)
at
org.eclipse.paho.client.mqttv3.internal.ClientState.notifyReceivedAck(ClientState.java:885)
at org.eclipse.paho.client.mqttv3.internal.CommsReceiver.run(CommsReceiver.java:118)
at java.lang.Thread.run(Thread.java:809)
```

Figure 25. Exception thrown when connecting to the Eclipse Paho sample

mqxr.log contains additional output shown in Figure 25 on page 242.

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 26 on page 242, 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.

```
15/06/15 13:49:28.337
AMQXR2050E: Unable to load JAAS config:MQXRWrongConfig.
The following exception occurred javax.security.auth.login.LoginException:
No LoginModules configured for MQXRWrongConfig
```

Figure 26. Error loading JAAS configuration

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. 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" on page 243
- · "Damaged queue manager object" on page 244
- "Damaged single object" on page 245
- "Automatic media recovery failure" on page 245

z/OS

See the following links for instructions on how to recover from different types of failures on IBM MQ for z/OS:

. z/0S

"Shared queue problems" on page 246

> z/0S

"Active log problems" on page 246

. z/0S

"Archive log problems" on page 252

. z/0S

"BSDS problems" on page 254

z/OS

"Page set problems" on page 261

z/0S

"Coupling facility and Db2 problems" on page 262

> z/0S

"Problems with long-running units of work" on page 265

> z/0S

"IMS-related problems" on page 265

z/0S

"Hardware problems" on page 267

Related concepts

"IBM MQ Troubleshooting and support" on page 7

If you are having problems with your queue manager network or IBM MQ applications, use the techniques described to help you diagnose and solve the problems.

"Troubleshooting overview" on page 7

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?"

Related tasks

"Collecting troubleshooting information" on page 41

Collect troubleshooting information (MustGather data) to help with investigating the problem that you are having with IBM MQ. You can also subscribe to notifications about IBM MQ fixes, troubleshooting and other news.

"Making initial checks on UNIX, Linux, and Windows" on page 9

Before you start problem determination in detail on UNIX, Linux, and Windows, 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.

Backing up and restoring IBM MQ

Planning for backup and recovery on z/OS

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 dspmqfls 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 procedures on z/OS

Use this topic as a reference for various recovery procedures.

This topic describes procedures for recovering IBM MQ after various error conditions. These error conditions are grouped in the following categories:

Table 20. Example recovery procedures			
Problem category	Problem	Where to look next	
Shared queue problems	Conflicting definitions for both private and shared queues.	"Shared queue problems" on page 246	
Active log problems	 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 246	
Archive log problems	 Insufficient DASD space to complete offloading active log data sets. Offload task has terminated abnormally. Archive data set allocation problem. 1/2 Read I/O errors on the archive data set during restart. 	"Archive log problems" on page 252	
BSDS problems	 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 254	
Page set problems	Page set full. A page set has an I/O error.	"Page set problems" on page 261	

Table 20. Example recovery procedures (continued)			
Problem category	Problem	Where to look next	
coupling facility and Db2 problems	 Storage medium full. Db2 system fails. Db2 data-sharing group fails. Db2 and the coupling facility fail. 	"Coupling facility and Db2 problems" on page 262	
Unit of work problems	A long-running unit of work is encountered.	"Problems with long-running units of work" on page 265	
IMS problems	 An IMS application terminates abnormally. The IMS adapter cannot connect to IBM MQ. IMS not operational. 	"IMS-related problems" on page 265	
Hardware problems	Media recovery procedures	"Hardware problems" on page 267	

Shared queue problems

Problems occur if IBM MQ discovers that a page set based queue, and a shared queue of the same name are defined.

Symptoms

IBM MQ issues the following message:

 ${\tt CSQI063E\ +CSQ1\ QUEUE}\ queue-name\ IS\ {\tt BOTH\ PRIVATE\ AND\ SHARED}$

During queue manager restart, IBM 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" on page 247
- "Active log stopped" on page 247
- "One or both copies of the active log data set are damaged" on page 248
- "Write I/O errors on an active log data set" on page 248
- "I/O errors occur while reading the active log" on page 249
- "Active log is becoming full" on page 250

· Active log is full

Dual logging is lost

Symptoms

IBM MQ issues the following message:

```
CSQJ004I +CSQ1 ACTIVE LOG COPY n INACTIVE, LOG IN SINGLE MODE, ENDRBA=...
```

Having completed one active log data set, IBM MQ found that the subsequent (COPY n) data sets were not offloaded or were marked stopped.

System action

IBM 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

IBM MQ issues the following message:

```
CSQJ030E +CSQ1 RBA RANGE startrba TO endrba 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 IBM 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 IBM 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.
- 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

IBM 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).

Operator action

None.

Write I/O errors on an active log data set

Symptoms

IBM MQ issues the following message:

```
CSQJ105E +CSQ1 csect-name LOG WRITE ERROR DSNAME=...,
LOGRBA=..., ERROR STATUS=ccccffss
```

System action

IBM 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.)

I/O errors occur while reading the active log

Symptoms

IBM 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 xxxxxxx DUE TO I/O ERROR
```

- If the second copy is satisfactory, the first copy is not stopped.
- If the error occurs during recovery, IBM 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, IBM 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" on page 251

System programmer action

Use the DEFINE LOG command to dynamically add further active log data sets. This permits IBM 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.

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.

IBM MQ issues the following CSQJ111A 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.

If the offload process has stalled because some resource is not available or for some other reason, it might be possible to resolve the problem by canceling the currently running offload task using the ARCHIVE LOG CANCEL OFFLOAD command.

System action

IBM MQ waits for an available active log data set before resuming normal IBM 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 106.

System programmer action

You can provide additional active log data sets before restarting the queue manager. This permits IBM 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.

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).
- 3. Adding additional log data sets dynamically, using the DEFINE LOG command.

When you restart the queue manager, offloading starts automatically during startup, and any work that was in progress when IBM 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 IBM 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" on page 252
- "Offload task terminated abnormally" on page 252
- "Insufficient DASD space to complete offload processing" on page 253
- "Read I/O errors on the archive data set while IBM MQ is restarting" on page 254

Allocation problems

Symptoms

IBM MQ issues message: CSQJ103E

```
CSQJ103E +CSQ1 LOG ALLOCATION ERROR DSNAME=dsname,
ERROR STATUS=eeeeiiii, 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:

```
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 IBM MQ message is issued for write I/O errors.

Only a z/OS error recovery program message appears. If you get IBM MQ message <u>CSQJ128E</u>, the offload task has ended abnormally.

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, message <u>CSQJ114I</u> 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. IBM MQ issues message CSQJ128E:

```
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

IBM MQ de-allocates the data set on which the error occurred. If IBM MQ is running in dual archive mode, IBM 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

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 return code in this message might provide further explanations for the cause of these symptoms.

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 IBM 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 IBM 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 IBM 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 <u>Using CSQ6ARVP</u>.)

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 IBM MQ is restarting

Symptoms

No specific IBM 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 <u>Planning your IBM MQ</u> environment on z/OS .

This topic describes the following BSDS problems:

- "Error occurs while opening the BSDS" on page 255
- "Log content does not agree with the BSDS information" on page 255
- "Both copies of the BSDS are damaged" on page 256
- "Unequal time stamps" on page 256
- "Out of synchronization" on page 257
- "I/O error" on page 258
- "Log range problems" on page 258

Normally, there are two copies of the BSDS, but if one is damaged, IBM 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.

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)

Error occurs while opening the BSDS

Symptoms

IBM MQ issues the following message:

```
CSQJ100E +CSQ1 ERROR OPENING BSDSn DSNAME=..., ERROR STATUS=eeii
```

where *eeii* is the VSAM return code. For information about VSAM codes, see the *DFSMS/MVS Macro Instructions for Data Sets* documentation.

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

IBM 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

IBM 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

IBM 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 earlier 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 IBM MQ data, they are also down-level. The down-level volume has the earlier 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

IBM 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 for details of how to do this.

Out of synchronization

Symptoms

IBM MQ issues the following message during queue manager initialization:

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. IBM 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, IBM MQ compares the copies of the BSDSs. If the copies are different, CSQJ122E is issued.

This message is also issued by the BSDS conversion utility if two input BSDS are specified and a record is found that differs between the two BSDS copies. This situation can arise if the queue manager terminated abnormally prior to the BSDS conversion utility being run.

System action

Queue manager startup or the utility is terminated.

System programmer action

None.

Operator action

If the error occurred during queue manager initialization, 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.

If the error occurred when running the BSDS conversion utility, carry out these steps:

- 1. Attempt to restart the queue manager and shut it down cleanly before attempting to run the BSDS conversion utility again.
- 2. If this does not solve the problem, run the print log map utility on both copies of the BSDS, and compare the lists to determine which copy is accurate or current.
- 3. Change the JCL used to invoke the BSDS conversion utility to specify the current BSDS in the SYSUT1 DD statement, and remove the SYSUT2 DD statement, before submitting the job again.

I/O error

Symptoms

IBM 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 IBM 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.

Log range problems Symptoms

IBM MQ has issued message <u>CSQJ113E</u> when reading its own log, or message <u>CSQJ133E</u> or <u>CSQJ134E</u> when reading the log of a queue manager in the queue sharing group. This can happen when you do not have the archive logs needed to restart the queue manager or recover a CF structure.

System action

Depending upon what log record is being read and why, the requestor might end abnormally with a reason code of X '00D1032A'.

System programmer action

Run the print log map utility (CSQJU004) to determine the cause of the error. When message CSQJ133E or CSQJ134E has been issued, run the utility against the BSDS of the queue manager indicated in the message.

If you have:

- Deleted the entry with the log range (containing the log RBA or LRSN indicated in the message) from the BSDS, and
- · Not deleted or reused the data set

you can add the entry back into the BSDS using the following procedure:

- 1. Identify the data set containing the required RBA or LRSN, by looking at an old copy of the contents of BSDS, or by running CSQJU004 against a backup of the BSDS.
- 2. Add the data set back into the BSDS using the change log inventory utility (CSQJU003).

3. Restart the queue manager.

If an archive log data set has been deleted, you will not be able to recover the page set or CF structure that needs the archive logs. Identify the reason that the queue manager needs to read the log record, then take one of the following actions depending on the page set or CF structure affected.

Page sets

Message CSQJ113E during the recovery phase of queue manager restart indicates that the log is needed to perform media recovery to bring a page set up to date.

Identify the page sets that need the deleted log data set for media recovery, by looking at the media recovery RBA in the <u>CSQI1049I</u> message issued for each page set during queue manager restart, then perform the following actions.

· Page set zero

You can recover the objects on page set zero, by using the following procedure.



Attention: All data in all other page sets will be lost when you carry out the procedure.

- 1. Use function SDEFS of the CSQUTIL utility to produce a file of IBM MQ DEFINE commands.
- 2. Format page set zero using CSQUTIL, then redefine the other page sets as described in the next section.
- 3. Restart the queue manager.
- 4. Use CSQUTIL to redefine the objects using the DEFINE commands produced by the utility in step 1.

• Page sets 1-99

Use the following procedure to redefine the page sets.



Attention: Any data on the page set is lost when you carry out this operation.

- 1. If you can access the page set without any I/O errors, reformat the page set using the CSQUTIL utility with the command FORMAT TYPE(NEW).
- 2. If I/O errors occurred when accessing the page set, delete the page set and re-create it.

If you want the page set to be the same size as before, use the command LISTCAT ENT(dsname) ALLOC to obtain the existing space allocations, and use these in the z/OS <u>DEFINE CLUSTER</u> command.

Format the new page set using the CSQUTIL utility with the command FORMAT TYPE(NEW).

3. Restart the queue manager. You might have to take certain actions, such as resetting channels or resolving indoubt channels.

CF structures

Messages <u>CSQJ113E</u>, <u>CSQJ133E</u>, or <u>CSQJ134E</u>, during the recovery of a CF structure, indicate that the logs needed to recover the structure are not available on at least one member of the queue sharing group.

Take one of the following actions depending on the structure affected:

Application CF structure

Issue the command RECOVER CFSTRUCT(structure-name) TYPE(PURGE).

This process empties the structure, so any messages on the structure are lost.

CSQSYSAPPL structure

Contact your IBM support center.

Administration structure

This structure is rebuilt using log data since the last checkpoint on each queue manager, which should be in active logs.

If you get this error during administration structure recovery, contact your IBM support center as this indicates that the active log is not available.

Once you have recovered the page set or CF structure, perform a backup of the logs, BSDS, page sets, and CF structures.

To prevent this problem from occurring again, increase the:

- Archive log retention (ARCRETN) value to be longer, and
- Increase the frequency of the CF structure backups.

Recovering a CF structure

Conceptually, the data from the previously backed up CF structure is read from the IBM MQ log; the log is read forwards from the backup and any changes are reapplied to the restored structure.

About this task

The log range to use is found from the latest backup of each structure to be recovered, to the current time. The log range is identified by log range sequence number (LRSN) values.

A LRSN uses the six most significant digits of a 'store clock value'.

Note that the whole log (back to the time the structure was created) is read, if you have not done a backup of the structure.

Procedure

1. Check that the logs from each queue manager in the queue sharing group (QSG) are read for records in this LSRN range.

Note that the logs are read backwards.

- 2. Check that a list of changes for each structure to be recovered is built.
- 3. Data from the coupling facility (CF) structure backup is read and the data is restored.

For example, if the backup was done on queue manager A, and the recovery is running on queue manager B, queue manager B reads the logs from queue manager A to restore the structure.

When the start of the backup of the CF structure is read, an internal task is started to take the restored data for the structure and merge it with the changes read from the log.

4. Check that processing continues for each structure being restored.

Example

In the following example, the command RECOVER CFSTRUCT(APP3) has been issued, and the following messages produced:

```
04:00:00 CSQE132I CDL2 CSQERRPB Structure recovery started, using log range from LRSN=CC56D01026CC to LRSN=CC56D0368924
This is the start of reading the logs backwards from each qmgr in the queue sharing group from the time of failure to the to the structure backup. The LRSN values give the ranges being used. Log records for all structures (just one structure in this example) being recovered are processed at the same time.

04:02:00 CSQE133I CDL2 CSQERPLS Structure recovery reading log backwards, LRSN=CC56D0414372 This message is produced periodically to show the process

04:02:22 CSQE134I CDL2 CSQERRPB Structure recovery reading log completed The above process of replaying the logs backwards has finished,

04:02:22 CSQE130I CDL2 CSQERCF2 Recovery of structure APP3 started, using CDL1 log range from RBA=000EE86D902E to RBA=000EF5E8E4DC
The task to process the data for APP3 has been started. The last backup of CF structure APP3 was done on CDL1 within the given RBA range, so this log range has to be read.

04:02:29 CSQE131I CDL2 CSQERCF2 Recovery of structure APP3 completed The data merge has completed. The structure is recovered.
```

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" on page 261 describes what happens if a page set is damaged.
- "Page set full" on page 261 describes what happens if there is not enough space on the page set for any more MOI operations.

Page set I/O errors

Problem

A page set has an I/O error.

Symptoms

This message is issued:

```
{\tt 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. IBM 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. IBM 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 IBM 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 might take a long time if a large amount of 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.

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 IBM 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 <u>Managing page sets</u> 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" on page 262
- "A Db2 system fails" on page 263
- "A Db2 data-sharing group fails" on page 263
- "Db2 and the coupling facility fail" on page 264

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 IBM 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= newsize

where newsize 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 IBM MQ is connected to fails, IBM 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, IBM 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 IBM 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 IBM MQ API operations continue to function as normal for shared queues, and all IBM 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 IBM 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 IBM 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, IBM 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 IBM MQ commands to define any missing objects. To do this, use the following procedure:

- 1. Prevent IBM 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 IBM 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 IBM 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 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.

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

IBM 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 IBM MQ.

This topic includes plans for the following problems that you might encounter in the IMS environment:

- "IMS cannot connect to IBM MQ" on page 265
- "IMS application problem" on page 266
- "IMS is not operational" on page 266

IMS cannot connect to IBM MQ

Problem

The IMS adapter cannot connect to IBM 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 IBM MQ for z/OS messages, completion, and reason codes documentation.

If an IMS application program tries to access IBM 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 IBM 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

IBM 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 IBM MQ when IMS terminates, IBM MQ issues message CSQ3201E.
 This message indicates that IBM MQ end-of-task (EOT) routines have been run to clean up and disconnect any connected threads.

System action

IBM 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 IBM MQ that might be in doubt.

You might need to restart the connection to IBM 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, IBM 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. Reinstate 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 IBM 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.

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®, 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.

IBW.

Part Number: