

7.5

*Developing Applications for IBM
WebSphere MQ*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 1077](#).

This edition applies to version 7 release 5 of IBM® WebSphere® MQ and to all subsequent releases and modifications until otherwise indicated in new editions.

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

© **Copyright International Business Machines Corporation 2007, 2024.**

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

Contents

Developing applications.....	7
Application development concepts.....	8
Application programs using the MQI.....	9
IBM WebSphere MQ messages.....	9
Preparing and running Microsoft Transaction Server applications.....	39
Using IBM WebSphere MQ with WebSphere Application Server.....	39
Transactional support scenarios.....	39
Deciding which language to use.....	75
IBM WebSphere MQ data definition files.....	77
Coding in C.....	79
Coding in COBOL.....	81
Coding in pTAL.....	82
Coding in Visual Basic.....	83
The IBM WebSphere MQ Object Model.....	83
Using JMS or Java.....	85
Designing IBM WebSphere MQ applications.....	85
Designing your messages.....	88
Application design and performance.....	89
Advanced IBM WebSphere MQ techniques.....	90
Sample IBM WebSphere MQ programs.....	91
Sample programs for distributed platforms.....	92
Writing a queuing application.....	186
The Message Queue Interface interface overview.....	187
Connecting to and disconnecting from a queue manager.....	198
Opening and closing objects.....	206
Putting messages on a queue.....	216
Getting messages from a queue.....	230
Writing publish/subscribe applications.....	266
Inquiring about and setting object attributes.....	306
Committing and backing out units of work.....	309
Starting IBM WebSphere MQ applications using triggers.....	315
Working with the MQI and clusters.....	331
Writing client applications.....	335
Using the message queue interface (MQI) for client applications.....	336
Building applications for IBM WebSphere MQ MQI clients.....	341
Running applications in the IBM WebSphere MQ MQI client environment.....	343
Preparing and running CICS and Tuxedo applications.....	354
Preparing and running Microsoft Transaction Server applications.....	356
Preparing and running IBM WebSphere MQ JMS applications.....	357
User exits, API exits and installable services.....	357
Writing and compiling exits and installable services.....	357
Building an IBM WebSphere MQ application.....	409
Building your application on AIX.....	409
Building your application on HP Integrity NonStop Server.....	416
Building your application on HP-UX.....	421
Building your application on Linux.....	427
Building your application on Solaris.....	433
Building your application on Windows systems.....	439
Using lightweight directory access protocol services with IBM WebSphere MQ for Windows.....	446
Developing IBM WebSphere MQ Telemetry applications.....	452
IBM WebSphere MQ Telemetry sample programs.....	452
Creating your first publisher using Java.....	455

Creating an asynchronous publisher using Java.....	460
Creating a recoverable asynchronous publisher using Java.....	464
Creating a subscriber using Java.....	469
Authenticating an MQTT client using JAAS.....	474
Authenticating an SSL connection using self-signed certificates.....	480
Authenticating an SSL connection using a certificate chain.....	484
Creating your first publisher using C.....	489
Creating an asynchronous publisher using C.....	493
Creating a subscriber using C.....	496
Client programming concepts.....	501
C client programming concepts.....	520
Handling program errors.....	523
Locally determined errors.....	523
Using report messages for problem determination.....	525
Remotely determined errors.....	525
Multicast programming.....	528
Multicast and the Message Queue Interface.....	528
Multicast connection to a queue manager.....	530
Programming data conversion for Multicast messaging.....	531
Multicast exception reporting.....	531
Using .NET.....	534
Getting started with IBM WebSphere MQ classes for .NET.....	535
Writing and deploying IBM WebSphere MQ.NET programs.....	549
IBM WebSphere MQ custom channel for Microsoft Windows Communication Foundation (WCF).....	568
Introduction to using the IBM WebSphere MQ custom channel for WCF with .NET 3.....	568
Using IBM WebSphere MQ custom channels for WCF.....	572
Using the WCF samples.....	588
Problem determination on the WCF custom channel for IBM WebSphere MQ.....	595
Using C++.....	601
Sample programs.....	604
C++ language concepts.....	608
Messaging in C++.....	612
Building IBM WebSphere MQ C++ programs.....	618
Using IBM WebSphere MQ classes for Java.....	625
Getting started with IBM WebSphere MQ classes for Java.....	625
Installation and configuration of IBM WebSphere MQ classes for Java.....	627
Introduction for programmers.....	638
Writing IBM WebSphere MQ classes for Java applications.....	639
Using IBM WebSphere MQ classes for JMS.....	684
Getting started with IBM WebSphere MQ classes for JMS.....	686
Installation and configuration of IBM WebSphere MQ classes for JMS.....	687
Introduction for programmers.....	761
Writing IBM WebSphere MQ classes for JMS applications.....	770
Application Server Facilities (ASF).....	886
Using the IBM WebSphere MQ JMS administration tool.....	893
Using IBM WebSphere MQ Explorer for JMS configuration.....	901
Using the WebSphere MQ Headers package.....	901
Using with WebSphere MQ classes for Java	902
Using with WebSphere MQ classes for JMS.....	903
Using Web services in IBM WebSphere MQ.....	904
IBM WebSphere MQ transport for SOAP.....	905
IBM WebSphere MQ bridge for HTTP.....	977
Using the Component Object Model Interface (IBM WebSphere MQ Automation Classes for ActiveX).....	987
Designing and programming using IBM WebSphere MQ Automation Classes for ActiveX.....	988
IBM WebSphere MQ Automation Classes for ActiveX reference.....	993
Troubleshooting.....	1058
ActiveX interface to the MQAI.....	1063

About the IBM WebSphere MQ Automation Classes for ActiveX Starter samples.....	1071
Notices.....	1077
Programming interface information.....	1078
Trademarks.....	1078

Developing applications

IBM WebSphere MQ provides several ways in which you can develop applications to send and receive messages that you need to support your business processes. You can also develop applications to manage your queue managers and related resources.

Before you develop applications for IBM WebSphere MQ, ensure you are familiar with the concepts in [IBM WebSphere MQ Technical overview](#) [IBM WebSphere MQ Technical overview](#).

You can develop applications for IBM WebSphere MQ in a number of different programming languages. For information about the programming languages supported and their features, see [“Deciding which programming language to use” on page 75](#).

See the following sections for the types of applications that you can write for IBM WebSphere MQ on different platforms.

Types of application you can write for IBM WebSphere MQ

This information is about the types of applications that can be written on IBM WebSphere MQ.

IBM WebSphere MQ products are queue managers and application enablers. They support the IBM Message Queue Interface (MQI) through which programs can put messages on a queue and get messages from a queue.

With IBM WebSphere MQ for non-z/OS platforms, you can write applications that:

- Send messages to other applications running under the same operating systems. The applications can be on either the same or another system.
- Send messages to applications that run on other IBM WebSphere MQ platforms.
- Use message queuing from within CICS® for TXSeries® for AIX®, TXSeries for HP-UX, TXSeries for Solaris, and TXSeries for Windows systems applications.
- Use message queuing from within Encina for AIX, HP-UX, Solaris, and Windows systems.
- Use message queuing from within Tuxedo for AIX, AT&T, HP-UX, Solaris, and Windows systems.
- Use IBM WebSphere MQ as a transaction manager, coordinating updates made by external resource managers within IBM WebSphere MQ units of work. The following external resource managers are supported and comply with the X/OPEN XA interface
 - DB2®
 - Informix®
 - Oracle
 - Sybase
- Process several messages together as a single unit of work that can be committed or backed out.
- Run from a full IBM WebSphere MQ environment, or run from a IBM WebSphere MQ MQI client environment on the following platforms:
 - UNIX and Linux® systems
 - Windows

Related concepts

[Security](#)

Application development concepts

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

Before you start to design and write your IBM WebSphere MQ applications, familiarize yourself with the basic IBM WebSphere MQ concepts, see the topics in [Technical overview](#). For information about the types of application you can write for IBM WebSphere MQ, see [“Developing applications” on page 7](#).

Use the following links to find out about IBM WebSphere MQ concepts specific to application development:

- [“IBM WebSphere MQ messages” on page 9](#)
- [Point-to-point messaging](#)
- [Introduction to WebSphere MQ publish/subscribe messaging](#)
- [“Using the message queue interface \(MQI\) in a client application” on page 336](#)
- [“Using Web services in WebSphere MQ” on page 904](#)
- [“Channel-exit programs for messaging channels” on page 378](#)
- [“Transactional support scenarios” on page 39](#)

Before you can run applications that use the MQI you must create certain IBM WebSphere MQ objects. For more information, see [“Application programs using the MQI” on page 9](#).

Related concepts

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Sample WebSphere MQ programs” on page 91](#)

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

[“Writing a queuing application” on page 186](#)

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Writing client applications” on page 335](#)

What you need to know to write client applications on WebSphere MQ.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Using WebSphere MQ classes for JMS” on page 684](#)

WebSphere MQ classes for Java Message Service (WebSphere MQ classes for JMS) is the JMS provider that is supplied with WebSphere MQ. As well as implementing the interfaces defined in the javax.jms package, WebSphere MQ classes for JMS provides two sets of extensions to the JMS API.

[“Using the Component Object Model Interface \(WebSphere MQ Automation Classes for ActiveX\)” on page 987](#)

The WebSphere MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access WebSphere MQ.

[“Using WebSphere MQ classes for Java” on page 625](#)

WebSphere MQ classes for Java enable you to use WebSphere MQ in a Java environment. A Java application can use either WebSphere MQ classes for Java or WebSphere MQ classes for JMS to access WebSphere MQ resources.

[“Using .NET” on page 534](#)

WebSphere MQ classes for .NET allow a program written in the .NET programming framework to connect to WebSphere MQ as a WebSphere MQ MQI client or to connect directly to a WebSphere MQ server.

[“Using C++” on page 601](#)

WebSphere MQ provides C++ classes equivalent to WebSphere MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

Application programs using the MQI

IBM WebSphere MQ application programs need certain objects before they can run successfully.

Figure 1 on [page 9](#) shows an application that removes messages from a queue, processes them, and then sends some results to another queue on the same queue manager.




Figure 1. Queues, messages, and applications

Whereas applications can put messages onto local or remote queues (using MQPUT), they can only get messages directly from local queues (using MQGET).

Before this application can run, the following conditions must be satisfied:

- The queue manager must exist and be running.
- The first application queue, from which the messages are to be removed, must be defined.
- The second queue, on which the application puts the messages, must also be defined.
- The application must be able to connect to the queue manager. To do this it must be linked to IBM WebSphere MQ. See [“Building an IBM WebSphere MQ application” on page 409](#).
- The applications that put the messages on the first queue must also connect to a queue manager. If they are remote, they must also be set up with transmission queues and channels. This part of the system is not shown in [Figure 1 on page 9](#).

IBM WebSphere MQ messages

This information introduces the IBM WebSphere MQ message concept, message parts, and the message descriptor.

IBM WebSphere MQ messages consist of two parts:

- Message properties
- Application data

Figure 2 on [page 10](#) represents a message and shows how it is logically divided into message properties and application data.




Figure 2. Representation of a message

The application data carried in a WebSphere MQ message is not changed by a queue manager unless data conversion is carried out on it. Also, WebSphere MQ does not put any restrictions on the content of this data. The length of the data in each message cannot exceed the value of the *MaxMsgLength* attribute of both the queue and queue manager.

On WebSphere MQ for AIX, WebSphere MQ for HP-UX, WebSphere MQ for Linux, WebSphere MQ for Solaris, and WebSphere MQ for Windows, the *MaxMsgLength* defaults to 100 MB (104 857 600 bytes).

Make your messages slightly shorter than the value of the *MaxMsgLength* attribute in some circumstances. See [“The data in your message”](#) on page 220 for more information.

You create a message when you use the MQPUT or MQPUT1 MQI calls. As input to these calls, you supply the control information (such as the priority of the message and the name of a reply queue) and your data, and the call then puts the message on a queue. See [MQPUT](#) and [MQPUT1](#) for more information about these calls.

Message descriptor

You can access message control information using the MQMD structure, which defines the *message descriptor*.

For a full description of the MQMD structure, see [MQMD - Message descriptor](#).

See [“Message context”](#) on page 37 for a description of how to use the fields within the MQMD that contain information about the origin of the message.

There are different versions of the message descriptor. Additional information for grouping and segmenting messages (see [“Message groups”](#) on page 34) is provided in Version 2 of the message descriptor (or the MQMDE). This is the same as the Version 1 message descriptor but has additional fields. These are described in the [MQMDE - Message descriptor extension](#).

Types of message

There are four types of messages defined by IBM WebSphere MQ.

These four messages are:

- [Datagram](#)
- [Request messages](#)
- [Reply messages](#)
- [Report messages](#)
 - [Types of report message](#)
 - [Report message options](#)

Applications can use the first three types of messages to pass information between themselves. The fourth type, report, is for applications and queue managers to use to report information about events such as the occurrence of an error.

Each type of message is identified by an MQMT_* value. You can also define your own types of message. For the range of values you can use, see [MsgType](#).

Datagrams

Use a *datagram* when you do not require a reply from the application that receives the message (that is, gets the message from the queue).

An example of an application that might use datagrams is one that displays flight information in an airport lounge. A message might contain the data for a whole screen of flight information. Such an application is unlikely to request an acknowledgment for a message because it probably does not matter if a message is not delivered. The application sends an update message after a short time.

Request messages

Use a *request message* when you want a reply from the application that receives the message.

An example of an application that could use request messages is one that displays the balance of a checking account. The request message could contain the number of the account, and the reply message would contain the account balance.

If you want to link your reply message with your request message, there are two options:

- Make the application that handles the request message responsible for ensuring that it puts information into the reply message that relates to the request message.
- Use the report field in the message descriptor of your request message to specify the content of the *MsgId* and *CorrelId* fields of the reply message:
 - You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the reply message (the default action is to copy *MsgId*).
 - You can request that either a new *MsgId* is generated for the reply message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the reply message (the default action is to generate a new message identifier).

Reply messages

Use a *reply message* when you reply to another message.

When you create a reply message, respect any options that were set in the message descriptor of the message to which you are replying. Report options specify the content of the message identifier (*MsgId*) and correlation identifier (*CorrelId*) fields. These fields allow the application that receives the reply to correlate the reply with its original request.

Report messages

Report messages inform applications about events such as the occurrence of an error when processing a message.

They can be generated by:

- A queue manager,
- A message channel agent (for example, if they cannot deliver the message), or
- An application (for example, if it cannot use the data in the message).

Report messages can be generated at any time, and might arrive on a queue when your application is not expecting them.

Types of report message

When you put a message on a queue, you can select to receive:

- An *exception report message*. This is sent in response to a message with the exceptions flag set. It is generated by the message channel agent (MCA) or the application.
- An *expiry report message*. This indicates that an application attempted to retrieve a message that had reached its expiry threshold; the message is marked to be discarded. This type of report is generated by the queue manager.
- A *confirmation of arrival (COA) report message*. This indicates that the message has reached its target queue. It is generated by the queue manager.
- A *confirmation of delivery (COD) report message*. This indicates that the message has been retrieved by a receiving application. It is generated by the queue manager.
- A *positive action notification (PAN) report message*. This indicates that a request has been successfully serviced (that is, the action requested in the message has been performed successfully). This type of report is generated by the application.
- A *negative action notification (NAN) report message*. This indicates that a request has not been successfully serviced (that is, the action requested in the message has not been performed successfully). This type of report is generated by the application.

Note: Each type of report message contains one of the following:

- The entire original message
- The first 100 bytes of data in the original message
- No data from the original message

You can request more than one type of report message when you put a message on a queue. If you select the delivery confirmation report message and the exception report message options, if the message fails to be delivered, you receive an exception report message. However, if you select only the delivery confirmation report message option and the message fails to be delivered, you do not get an exception report message.

The report messages that you request, when the criteria for generating a particular message are met, are the only ones that you receive.

Report message options

You can *discard* a message after an exception has arisen. If you select the discard option, and have requested an exception report message, the report message goes to the *ReplyToQ* and *ReplyToQMgr*, and the original message is discarded.

Note: A benefit of this is that you can reduce the number of messages going to the dead-letter queue. However, it does mean that your application, unless it sends only datagram messages, has to deal with returned messages. When an exception report message is generated, it inherits the persistence of the original message.

If a report message cannot be delivered (if the queue is full, for instance), the report message is placed on the dead-letter queue.

If you want to receive a report message, specify the name of your reply-to queue in the *ReplyToQ* field; otherwise the MQPUT or MQPUT1 of your original message fails with MQRC_MISSING_REPLY_TO_Q.

You can use other report options in the message descriptor (MQMD) of a message to specify the content of the *MsgId* and *CorrelId* fields of any report messages that are created for the message:

- You can request that either the *MsgId* or the *CorrelId* of the original message is to be copied into the *CorrelId* field of the report message. The default action is to copy the message identifier. Use MQRO_COPY_MSG_ID_TO_CORRELID because it enables the sender of a message to correlate the reply or report message with the original message. The correlation identifier of the reply or report message is identical to the message identifier of the original message.
- You can request that either a new *MsgId* is generated for the report message, or that the *MsgId* of the original message is to be copied into the *MsgId* field of the report message. The default action is to generate a new message identifier. Use MQRO_NEW_MSG_ID because it ensures that each

message in the system has a different message identifier, and can be distinguished unambiguously from all other messages in the system.

- Specialized applications might need to use MQRO_PASS_MSG_ID or MQRO_PASS_CORREL_ID. However, you need to design the application that reads the messages from the queue to ensure that it works correctly when, for example, the queue contains multiple messages with the same message identifier.

Server applications must check the settings of these flags in the request message, and set the *MsgId* and *CorrelId* fields in the reply or report message appropriately.

Applications that act as intermediaries between a requester application and a server application do not need to check the settings of these flags. This is because these applications typically need to forward the message to the server application with the *MsgId*, *CorrelId*, and *Report* fields unchanged. This allows the server application to copy the *MsgId* from the original message in the *CorrelId* field of the reply message.

When generating a report about a message, server applications must test to see if any of these options have been set.

For more information about how to use report messages, see [Report](#).

To indicate the nature of the report, queue managers use a range of feedback codes. They put these codes in the *Feedback* field of the message descriptor of a report message. Queue managers can also return MQI reason codes in the *Feedback* field. IBM WebSphere MQ defines a range of feedback codes for applications to use.

For more information about feedback and reason codes, see [Feedback](#).

An example of a program that could use a feedback code is one that monitors the workloads of other programs serving a queue. If there is more than one instance of a program serving a queue, and the number of messages arriving on the queue no longer justifies this, such a program can send a report message (with the feedback code MQFB_QUIT) to one of the serving programs to indicate that the program should terminate its activity. (A monitoring program could use the MQINQ call to find out how many programs are serving a queue.)

Reports and segmented messages

Not supported on WebSphere MQ for z/OS®.

If a message is segmented (see [“Message segmentation” on page 252](#) for a description of segmented messages) and you ask for reports to be generated, you might receive more reports than you would have done had the message not been segmented.

For reports generated by WebSphere MQ

If you segment your messages or allow the queue manager to do so, there is only one case in which you can expect to receive a single report for the entire message. This is when you have requested only COD reports, and you have specified MQGMO_COMPLETE_MSG on the getting application.

In other cases your application must be prepared to deal with several reports; usually one for each segment.

Note: If you segment your messages, and you need only the first 100 bytes of the original message data to be returned, change the setting of the report options to ask for reports with no data for segments that have an offset of 100 or more. If you do not do this, and you leave the setting so that each segment requests 100 bytes of data, and you retrieve the report messages with a single MQGET specifying MQGMO_COMPLETE_MSG, the reports assemble into a large message containing 100 bytes of read data at each appropriate offset. If this happens, you need a large buffer or you need to specify MQGMO_ACCEPT_TRUNCATED_MSG.

For reports generated by applications

If your application generates reports, always copy the WebSphere MQ headers that are present at the start of the original message data to the report message data.

Then add none, 100 bytes, or all of the original message data (or whatever other amount you would usually include) to the report message data.

You can recognize the WebSphere MQ headers that must be copied by looking at the successive *Format* names, starting with the MQMD and continuing through any headers present. The following *Format* names indicate these WebSphere MQ headers:

- MQMDE
- MQDLH
- MQXQH
- MQIIH
- MQH*

MQH* means any name that starts with the characters MQH.

The *Format* name occurs at specific positions for MQDLH and MQXQH, but for the other WebSphere MQ headers it occurs at the same position. The length of the header is contained in a field that also occurs at the same position for MQMDE, MQIMS, and all MQH* headers.

If you are using a Version 1 MQMD, and you are reporting on a segment, or a message in a group, or a message for which segmentation is allowed, the report data must start with an MQMDE. Set the *OriginalLength* field to the length of the original message data excluding the lengths of any WebSphere MQ headers that you find.

Retrieving reports

If you ask for COA or COD reports, you can ask for them to be reassembled for you with MQGMO_COMPLETE_MSG.

An MQGET with MQGMO_COMPLETE_MSG is satisfied when enough report messages (of a single type, for example COA, and with the same *GroupId*) are present on the queue to represent one complete original message. This is true even if the report messages themselves do not contain the complete original data; the *OriginalLength* field in each report message gives the length of original data represented by that report message, even if the data itself is not present.

You can use this technique even if there are several different report types present on the queue (for example, both COA and COD), because an MQGET with MQGMO_COMPLETE_MSG reassembles report messages only if they have the same *Feedback* code. However, you cannot usually use this technique for exception reports, because, in general, these have different *Feedback* codes.

You can use this technique to get a positive indication that the entire message has arrived. However, in most circumstances you need to cater for the possibility that some segments arrive while others might generate an exception (or expiry, if you have allowed this). You cannot use MQGMO_COMPLETE_MSG in this case, because, in general, you might get different *Feedback* codes for different segments and, you might get more than one report for a segment. You can, however, use MQGMO_ALL_SEGMENTS_AVAILABLE.

To allow for this you might need to retrieve reports as they arrive, and build up a picture in your application of what happened to the original message. You can use the *GroupId* field in the report message to correlate reports with the *GroupId* of the original message, and the *Feedback* field to identify the type of each report message. The way in which you do this depends on your application requirements.

One approach is as follows:

- Ask for COD reports and exception reports.

- After a specific time, check whether a complete set of COD reports has been received using MQGMO_COMPLETE_MSG. If so, your application knows that the entire message has been processed.
- If not, and exception reports relating to this message are present, handle the problem as for unsegmented messages, but ensure that you clean up orphan segments at some point.
- If there are segments for which there are no reports of any kind, the original segments (or the reports) might be waiting for a channel to be reconnected, or the network might be overloaded at some point. If no exception reports at all have been received (or if you think that the ones you have might be temporary only), you might decide to let your application wait a little longer.

As before, this is similar to the considerations you have when dealing with unsegmented messages, except that you must also consider the possibility of cleaning up orphan segments.

If the original message is not critical (for example, if it is a query, or a message that can be repeated later), set an expiry time to ensure that orphan segments are removed.

Back-level queue managers

When a report is generated by a queue manager that supports segmentation, but is received on a queue manager that does *not* support segmentation, the MQMDE structure (which identifies the *Offset* and *OriginalLength* represented by the report) is always included in the report data, in addition to zero, 100 bytes, or all of the original data in the message.

However, if a segment of a message passes through a queue manager that does not support segmentation, if a report is generated there, the MQMDE structure in the original message is treated purely as data. It is not therefore included in the report data if zero bytes of the original data have been requested. Without the MQMDE, the report message might not be useful.

Request at least 100 bytes of data in reports if there is a possibility that the message might travel through a back-level queue manager.

Format of message control information and message data

The queue manager is only interested in the format of the control information within a message, whereas applications that handle the message are interested in the format of both the control information and the data.

Format of message control information

Control information in the character-string fields of the message descriptor must be in the character set used by the queue manager.

The *CodedCharSetId* attribute of the queue manager object defines this character set. Control information must be in this character set because, when applications pass messages from one queue manager to another, message channel agents that transmit the messages use the value of this attribute to determine what data conversion to perform.

Format of message data

You can specify any of the following things:

- The format of the application data
- The character set of the character data
- The format of numeric data

To do this, use these fields:

Format

This indicates to the receiver of a message the format of the application data in the message.

When the queue manager creates a message, in some circumstances it uses the *Format* field to identify the format of that message. For example, when a queue manager cannot deliver a message,

it puts the message on a dead-letter (undelivered message) queue. It adds a header (containing more control information) to the message, and changes the *Format* field to show this.

The queue manager has a number of *built-in formats* with names beginning MQ, for example MQFMT_STRING. If these do not meet your needs, you can define your own formats (*user-defined formats*), but you must not use names beginning with MQ for these.

When you create and use your own formats, you must write a data-conversion exit to support a program getting the message using MQGMO_CONVERT.

CodedCharSetId

This defines the character set of character data in the message. If you want to set this character set to that of the queue manager, you can set this field to the constant MQCCSI_Q_MGR or MQCCSI_INHERIT.

When you get a message from a queue, compare the value of the *CodedCharSetId* field with the value that your application is expecting. If the two values differ, you might need to convert any character data in the message or use a data-conversion message exit if one is available.

Encoding

This describes the format of numeric message data that contains binary integers, packed-decimal integers, and floating point numbers. It is typically encoded according to the particular machine on which the queue manager is running.

When you put a message on a queue, you typically specify the constant MQENC_NATIVE in the *Encoding* field. This means that the encoding of your message data is the same as that of the machine on which your application is running.

When you get a message from a queue, compare the value of the *Encoding* field in the message descriptor with the value of the constant MQENC_NATIVE on your machine. If the two values differ, you might need to convert any numeric data in the message or use a data-conversion message exit if one is available.

Application data conversion

Application data might need to be converted to the character set and the encoding required by another application where different platforms are concerned.

It can be converted at the sending queue manager, or at the receiving queue manager. If the library of built-in formats does not meet your needs, you can define your own. The type of conversion depends on the message format that is specified in the format field of the message descriptor, MQMD.

Note: Messages with MQFMT_NONE specified are not converted.

Conversion at the sending queue manager

Set the CONVERT channel attribute to YES if you need the sending message channel agent (MCA) to convert the application data.

The conversion is performed at the sending queue manager for certain built-in formats and for user-defined formats if a suitable user exit is supplied.

Built-in formats

These include:

- Messages that are all characters (using the format name MQFMT_STRING)
- WebSphere MQ defined messages, for example Programmable Command Formats

WebSphere MQ uses Programmable Command Format messages for administration messages and events (the format name used is MQFMT_ADMIN in this case). You can use the same format (using the format name MQFMT_PCF) for your own messages, and take advantage of the built-in data conversion.

The queue manager built-in formats all have names beginning with MQFMT. They are listed and described in [Format](#).

Application-defined formats

For user-defined formats, application data conversion must be performed by a data-conversion exit program (for more information, see [“Writing data-conversion exits”](#) on page 397). In a client-server environment, the exit is loaded at the server and conversion takes place there.

Conversion at the receiving queue manager

Application message data can be converted by the receiving queue manager for both built-in and user-defined formats.

The conversion is performed during the processing of an MQGET call if you specify the MQGMO_CONVERT option. For details, see the [Options](#)

Coded character sets

WebSphere MQ products support the coded character sets that are provided by the underlying operating system.

When you create a queue manager, the queue manager coded character set ID (CCSID) used is based on that of the underlying environment. If this is a mixed code page, WebSphere MQ uses the SBCS part of the mixed code page as the queue manager CCSID.

For general data conversion, if the underlying operating system supports DBCS code pages, WebSphere MQ can use it.

See the documentation for your operating system for details of the coded character sets that it supports.

You need to consider application data conversion, format names, and user exits when writing applications that span multiple platforms. See [“Writing data-conversion exits”](#) on page 397 for information about invoking and writing data-conversion exits.

Message priorities

You set the priority of a message (in the *Priority* field of the MQMD structure) when you put the message on a queue. You can set a numeric value for the priority, or you can let the message take the default priority of the queue.

The *MsgDeliverySequence* attribute of the queue determines whether messages on the queue are stored in FIFO (first in, first out) sequence, or in FIFO within priority sequence. If this attribute is set to MQMDS_PRIORITY, messages are enqueued with the priority specified in the *Priority* field of their message descriptors; but if it is set to MQMDS_FIFO, messages are enqueued with the default priority of the queue. Messages of equal priority are stored on the queue in order of arrival.

The *DefPriority* attribute of a queue sets the default priority value for messages being put on that queue. This value is set when the queue is created, but it can be changed afterward. Alias queues, and local definitions of remote queues, can have different default priorities from the base queues to which they resolve. If there is more than one queue definition in the resolution path (see [“Name resolution”](#) on page 208), the default priority is taken from the value (at the time of the put operation) of the *DefPriority* attribute of the queue specified in the open command.

The value of the *MaxPriority* attribute of the queue manager is the maximum priority that you can assign to a message processed by that queue manager. You cannot change the value of this attribute. In WebSphere MQ, the attribute has the value 9; you can create messages having priorities between 0 (the lowest) and 9 (the highest).

Message properties

Use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing MQMD or MQRFH2 headers. They also facilitate communication between WebSphere MQ and JMS applications.

A message property is data associated with a message, consisting of a textual name and a value of a particular type. Message properties are used by message selectors to filter publications to topics or to selectively get messages from queues. Message properties can be used to include business data or state information without having to store it in the application data. Applications do not have to access data in the MQ Message Descriptor (MQMD) or MQRFH2 headers because fields in these data structures can be accessed as message properties using Message Queue Interface (MQI) function calls.

The use of message properties in WebSphere MQ mimics the use of properties in JMS. This means that you can set properties in a JMS application and retrieve them in a procedural WebSphere MQ application, or the other way round. To make a property available to a JMS application, assign it the prefix "usr"; it is then available (without the prefix) as a JMS message user property. For example, the WebSphere MQ property *usr.myproperty* (a character string) is accessible to a JMS application using the JMS call `message.getStringProperty('myproperty')`. Note that JMS applications are unable to access properties with the prefix "usr" if they contain two or more U+002E (".") characters. A property with no prefix and no U+002E (".") character is treated as if it had the prefix "usr". Conversely, a user property set in a JMS application can be accessed in a WebSphere MQ application by adding the "usr." prefix to the property name inquired on in an MQINQMP call.

Message properties and message length

Use the queue manager attribute *MaxPropertiesLength* to control the size of the properties that can flow with any message in a WebSphere MQ queue manager.

In general, when you use MQSETMP to set properties, the size of a property is the length of the property name in bytes, plus the length of the property value in bytes as passed into the MQSETMP call. It is possible for the character set of the property name and the property value to change during transmission of the message to its destination because these can be converted into Unicode; in this case the size of the property might change.

On an MQPUT or MQPUT1 call, properties of the message do not count toward the length of the message for the queue and the queue manager, but they do count toward the length of the properties as perceived by the queue manager (whether they were set using the message property MQI calls or not).

If the size of the properties exceeds the maximum properties length, the message is rejected with MQRC_PROPERTIES_TOO_BIG. Because the size of the properties is dependent on its representation, you should set the maximum properties length at a gross level.

It is possible for an application to successfully put a message with a buffer that is larger than the value of *MaxMsgLength*, if the buffer includes properties. This is because, even when represented as MQRFH2 elements, message properties do not count toward the length of the message. The MQRFH2 header fields add to the properties length only if one or more folders are contained and every folder in the header contains properties. If one or more folders are contained in the MQRFH2 header and any folder does not contain properties, the MQRFH2 header fields count toward the message length instead.

On an MQGET call, properties of the message do not count toward the length of the message as far as the queue and the queue manager are concerned. However, because the properties are counted separately it is possible that the buffer returned by an MQGET call is larger than the value of the *MaxMsgLength* attribute.

Do not have your applications query the value of *MaxMsgLength* and then allocate a buffer of this size before calling MQGET; instead, allocate a buffer you consider large enough. If the MQGET fails, allocate a buffer guided by the size of the *DataLength* parameter.

The *DataLength* parameter of the MQGET call returns the length in bytes of the application data and any properties returned in the buffer you have provided, if a message handle is not specified in the MQGMO structure.

The *Buffer* parameter of the MQPUT call contains the application message data to be sent and any properties represented in the message data.

When flowing to a queue manager that is earlier than Version 7.0 of the product, properties of the message, except those in the message descriptor, count toward the length of the message. Therefore, you should either raise the value of the *MaxMsgLength* attribute of channels going to a system earlier than Version 7.0 as necessary, to compensate for the fact that more data might be sent for each message. Alternatively, you can lower the queue or queue manager *MaxMsgLength*, so that the overall level of data being sent around the system remains the same.

There is a length limit of 100 MB for message properties, excluding the message descriptor or extension for each message.

The size of a property in its internal representation is the length of the name, plus the size of its value, plus some control data for the property. There is also some control data for the set of properties after one property is added to the message.

Property names

A property name is a character string. Certain restrictions apply to its length and the set of characters that can be used.

A property name is a case-sensitive character string, limited to +4095 characters unless otherwise restricted by the context. This limit is contained in the MQ_MAX_PROPERTY_NAME_LENGTH constant.

If you exceed this maximum length when using a message property MQI call, the call fails with reason code MQRC_PROPERTY_NAME_LENGTH_ERR.

Because there is no maximum property name length in JMS, it is possible for a JMS application to set a valid JMS property name that is not a valid WebSphere MQ property name when stored in an MQRFH2 structure.

In this case, when parsed, only the first 4095 characters of the property name are used; the following characters are truncated. This could cause an application using selectors to fail to match a selection string, or to match a string when not expecting to, since more than one property might truncate to the same name. When a property name is truncated, WebSphereMQ issues an error log message.

All property names must follow the rules defined by the Java Language Specification for Java Identifiers, with the exception that Unicode character U+002E (.) is permitted as part of the name - but not the start. The rules for Java Identifiers equate to those contained in the JMS specification for property names.

White space characters and comparison operators are prohibited. Embedded nulls are allowed in a property name but not recommended. If you use embedded nulls, this prevents the use of the MQVS_NULL_TERMINATED constant when used with the MQCHARV structure to specify variable length strings.

Keep property names simple because applications can select messages based on the property names and the conversion between the character set of the name and of the selector might cause the selection to fail unexpectedly.

WebSphere MQ property names use character U+002E (.) for logical grouping of properties. This divides up the namespace for properties. Properties with the following prefixes, in any mixture of lowercase or uppercase are reserved for use by the product:

- mcd
- jms
- usr
- mq
- sib
- wmq
- Root
- Body

- Properties

A good way to avoid name clashes is to ensure that all applications prefix their message properties with their Internet domain name. For example, if you are developing an application using domain name "ourcompany.com" you could name all properties with the prefix "com.ourcompany". This naming convention also allows for easy selection of properties; for example, an application can inquire on all message properties starting "com.ourcompany.%".

See [Property name restrictions](#) for further information about the use of property names.

Property name restrictions

When you name a property, you must observe certain rules.

The following restrictions apply to property names:

1. A property must not begin with the following strings:
 - "JMS" - reserved for use by WebSphere MQ classes for JMS.
 - "usr.JMS" - not valid.

The only exceptions are the following properties providing synonyms for JMS properties:

Property	Synonym for
JMSCorrelationID	Root.MQMD.CorrelId or jms.Cid
JMSDeliveryMode	Root.MQMD.Persistence or jms.Dlv
JMSDestination	jms.Dst
JMSExpiration	Root.MQMD.Expiry or jms.Exp
JMSMessageID	Root.MQMD.MsgId
JMSPriority	Root.MQMD.Priority or jms.Pri
JMSRedelivered	Root.MQMD.BackoutCount
JMSReplyTo (a string encoded as a URI)	Root.MQMD.ReplyToQ or Root.MQMD.ReplyToQMgr or jms.Rto
JMSTimestamp	Root.MQMD.PutDate or Root.MQMD.PutTime or jms.Tms
JMSType	mcd.Type or mcd.Set or mcd.Fmt
JMSXAppID	Root.MQMD.PutApplName
JMSXDeliveryCount	Root.MQMD.BackoutCount
JMSXGroupID	Root.MQMD.GroupId or jms.Gid
JMSXGroupSeq	Root.MQMD.MsgSeqNumber or jms.Seq
JMSXUserID	Root.MQMD.UserIdentifier

These synonyms allow an MQI application to access JMS properties in a similar fashion to WebSphere MQ classes for JMS client application. Of these properties, only JMSCorrelationID, JMSReplyTo, JMSType, JMSXGroupID, and JMSXGroupSeq can be set using the MQI.

Note that the JMS_IBM_* properties available from within WebSphere MQ classes for JMS are not available using the MQI. The fields that the JMS_IBM_* properties reference can be accessed in other ways by MQI applications.

2. A property must not be called, in any mixture of lower or uppercase, "NULL", "TRUE", "FALSE", "NOT", "AND", "OR", "BETWEEN", "LIKE", "IN", "IS" and "ESCAPE". These are the names of SQL keywords used in selection strings.

3. A property name beginning "mq " in any mixture of lowercase or uppercase and not beginning "mq_usr" can contain only one "." character (U+002E). Multiple "." characters are not allowed in properties with those prefixes.
4. Two "." characters must contain other characters in between; you cannot have an empty point in the hierarchy. Similarly a property name cannot end in a "." character.
5. If an application sets the property "a.b" and then the property "a.b.c", it is unclear whether in the hierarchy "b" contains a value or another logical grouping. Such a hierarchy is "mixed content" and this is not supported. Setting a property that causes mixed content is not allowed.

These restrictions are enforced by the validation mechanism as follows:

- Property names are validated when setting a property using the `MQSETMP` - Set message property call, if validation was requested when the message handle was created. If an attempt to validate a property is undertaken and fails due to an error in the specification of the property name, the completion code is `MQCC_FAILED` with reason:
 - `MQRC_PROPERTY_NAME_ERROR` for reasons 1-4.
 - `MQRC_MIXED_CONTENT_NOT_ALLOWED` for reason 5.
- The names of properties specified directly as `MQRFH2` elements are not guaranteed to be validated by the `MQPUT` call.

Message descriptor fields as properties

Most message descriptor fields can be treated as properties. The property name is constructed by adding a prefix to the message descriptor field's name.

If an MQI application wants to identify a message property contained in a message descriptor field, for example, in a selector string or using the message property APIs, use the following syntax:

Property name	Message descriptor field
Root.MQMD.<Field>	<Field>

Specify <Field> with the same case as for the `MQMD` structure fields in the C language declaration. For example, the property name `Root.MQMD.AccountingToken` accesses the `AccountingToken` field of the message descriptor.

The `StrucId` and `Version` fields of the message descriptor are not accessible using the syntax shown.

Message descriptor fields are never represented in an `MQRFH2` header as for other properties.

If the message data starts with an `MQMDE` that is honored by the queue manager, the `MQMDE` fields can be accessed using the `Root.MQMD.<Field>` notation described. In this case the `MQMDE` fields are treated as logically part of the `MQMD` from a properties perspective. See the section "MQMDE specified on `MQPUT` and `MQPUT1` calls" in [Overview of MQMDE](#).

Property data types and values

A property can be a boolean, a byte string, a character string, or a floating-point or integer number. The property can store any valid value in the range of the data type unless otherwise restricted by the context.

The data type of a property value must be one of the following values:

- `MQBOOL`
- `MQBYTE[]`
- `MQCHAR[]`
- `MQFLOAT32`
- `MQFLOAT64`
- `MQINT8`
- `MQINT16`
- `MQINT32`

- MQINT64

A property can exist but have no defined value; it is a null property. A null property is different from a byte property (MQBYTE[]) or character string property (MQCHAR[]) in that it has a defined but empty value, that is, one with a zero-length value.

Byte string is not a valid property data type in JMS or XMS. You are advised not to use byte string properties in the <usr> folder.

Selecting messages from queues

You can select messages from queues using the MsgId and CorrelId fields on an MQGET call, or by using a SelectionString on an MQOPEN or MQSUB call.

Selectors

A message selector is a variable-length string used by an application to register its interest in only those messages that have properties that satisfy the Structured Query Language (SQL) query that the selection string represents.

Selection using the MQSUB and MQOPEN function calls

You use the *SelectionString*, which is a structure of type MQCHARV, to make selections using the MQSUB and MQOPEN calls.

The *SelectionString* structure is used to pass a variable-length selection string to the queue manager.

The CCSID associated with the selector string is set via the VSCCSID field of the MQCHARV structure. The value used must be a CCSID that is supported for selector strings. See [Code page conversion](#) for a list of supported code pages.

Specifying a CCSID for which there is no WebSphere MQ supported Unicode conversion, results in an error of MQRC_SOURCE_CCSID_ERROR. This error is returned at the time that the selector is presented to the queue manager, that is, on the MQSUB, MQOPEN, or MQPUT1 call.

The default value for the VSCCSID field is MQCCSI_APPL, which indicates that the CCSID of the selection string is equal to the queue manager CCSID, or the client CCSID if connected through a client. The MQCCSI_APPL constant can however be overridden by an application redefining it before compiling.

If the MQCHARV selector represents a NULL string, no selection takes place for that message consumer and messages are delivered as if a selector had not been used.

The maximum length of a selection string is limited only by what can be described by the MQCHARV field *VSLength*.

The SelectionString is returned on the output from an MQSUB call using the MQSO_RESUME subscribe option, if you have provided a buffer and there is a positive buffer length in VSBufSize. If you do not provide a buffer, only the length of the selection string is returned in the VSLength field of the MQCHARV. If the buffer provided is smaller than the space required to return the field, only VSBufSize bytes are returned in the provided buffer.

An application cannot alter a selection string without first closing either the handle to the queue (for MQOPEN), or subscription (for MQSUB). A new selection string can then be specified on a subsequent MQOPEN or MQSUB call.

MQOPEN

Use MQCLOSE to close the opened handle, then specify a new selection string on a subsequent MQOPEN call.

MQSUB

Use MQCLOSE to close the returned subscription handle (hSub), then specify a new selection string on a subsequent MQSUB call.

[Figure 3 on page 23](#) shows the process of selection using the MQSUB call.

MQOPEN

(APP 1)

ObjectName = "MyDestQ"

hObj



MQSUB

(APP 1)


SelectionString = "Sport = 'Football'"

hObj

TopicString = "ResultsTopic"



ResultsTopic



MQGET

(APP 1)

hObj




Figure 3. Selection using MQSUB call


A selector can be passed in on the call to MQSUB by using the *SelectionString* field in the MQSD structure. The effect of passing in a selector on the MQSUB is that only those messages published to the topic being subscribed to, that match a supplied selection string, are made available on the destination queue.

Figure 4 on page 24 shows the process of selection using the MQOPEN call.

MQOPEN

(APP 1)

SelectorString = "League = 'Premiership'"
ObjectName = "SportQ"
hObj



MQGET

(APP 1) hObj

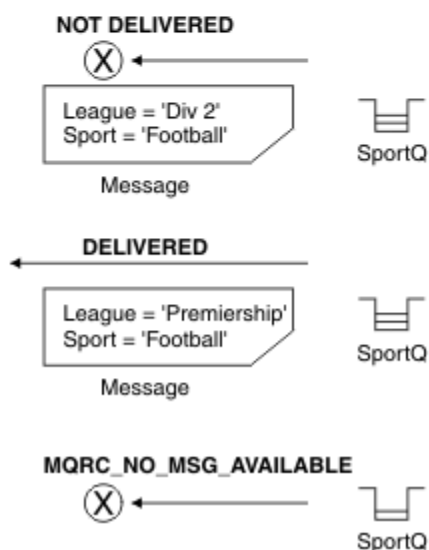


Figure 4. Selection using MQOPEN call

A selector can be passed in on the call to MQOPEN by using the *SelectorString* field in the MQOD structure. The effect of passing in a selector on the MQOPEN call is that only those messages on the opened queue, that match a selector, are delivered to the message consumer.

The main use for the selector on the MQOPEN call is for the point-to-point case where an application can elect to receive only those messages on a queue that match a selector. The previous example shows a simple scenario where two messages are put to a queue opened by MQOPEN but only one is received by the application getting it, as it is the only one that matches a selector.

Note that subsequent MQGET calls result in MQRC_NO_MSG_AVAILABLE as no further messages exist on the queue that match the given selector.

Selection behavior

Overview of IBM WebSphere MQ selection behavior.

The fields in an MQMDE structure are considered to be the message properties for the corresponding message descriptor properties if the MQMD:

- Has format MQFMT_MD_EXTENSION

- Is immediately followed by a valid MQMDE structure
- Is version one or contains the default version two fields only

It is possible for a selection string to resolve to either TRUE or FALSE before any matching against message properties takes place. , For example, it might be the case if the selection string is set to "TRUE <>FALSE". Such early evaluation is guaranteed to take place only when there are no message property references in the selection string.

If a selection string resolves to TRUE before any message properties are considered, all messages published to the topic subscribed to by the consumer are delivered. If a selection string resolves to FALSE before any message properties are considered, a reason code of MQRC_SELECTOR_ALWAYS_FALSE, and completion code MQCC_FAILED are returned on the function call that presented the selector.

Even if a message contains no message properties (other than header properties) then it can still be eligible for selection. If a selection string references a message property that does not exist, this property is assumed to have the value of NULL or 'Unknown'.

For example, a message might still satisfy a selection string like 'Color IS NULL', where 'Color' does not exist as a message property in the message.

Selection can be performed only on the properties that are associated with a message, not the message itself, unless an extended message selection provider is available. Selection can be performed on the message payload only if an extended message selection provider is available.

Each message property has a type associated with it. When you perform a selection, you must ensure that the values used in expressions to test message properties are of the correct type. If a type mismatch occurs, the expression in question resolves to FALSE.

It is your responsibility to ensure that the selection string and message properties use compatible types.

Selection criteria continue to be applied on behalf of inactive durable subscribers, so that only messages that match the selection string that was originally supplied are kept.

Selection strings are non-alterable when a durable subscription is resumed with alter (MQSO_ALTER). If a different selection string is presented when a durable subscriber resumes activity, then MQRC_SELECTOR_NOT_ALTERABLE is returned to the application.

Applications receive a return code of MQRC_NO_MSG_AVAILABLE if there is no message on a queue that meets the selection criteria.

If an application has specified a selection string containing property values then only those messages that contain matching properties are eligible for selection. For example, a subscriber specifies a selection string of "a = 3" and a message is published containing no properties, or properties where 'a' does not exist or is not equal to 3. The subscriber does not receive that message to its destination queue.

Messaging performance

Selecting messages from a queue requires IBM WebSphere MQ to sequentially inspect each message on the queue. Messages are inspected until a message is found that matches the selection criteria or there are no more messages to examine. Therefore, messaging performance suffers if message selection is used on deep queues.

To optimize message selection on deep queues when selection is based on JMSCorrelationID or JMSMessageID, use a selection string of the form JMSCorrelationID = ... or JMSMessageID = ... and reference only one property.

This method offers a significant improvement in performance for selection on JMSCorrelationID and offers a marginal performance improvement for JMSMessageID.

Using complex selectors

Selectors can contain many components, for example:

a and b or c and d or e and f or g and h or i and j ... or y and z

Use of such complex selectors can have serious performance implications and excessive resource requirements. As such, IBM WebSphere MQ will protect the system by failing to process overly complex selectors that could result in a system resource shortage. Protection can occur after approximately 100 tests on some platforms so selectors approaching that number of components may see failures. It is recommended that the use of selectors with many components are thoroughly tried and tested on the appropriate platforms to ensure that the protection limits are not reached.

The performance and complexity of selectors can be improved by simplifying them using additional parenthesis to combine components. For example:

(a and b or c and d) or (e and f or g and h) or (i and j) ...

Related concepts

Message selector syntax

A WebSphere MQ message selector is a string with syntax that is based on a subset of the SQL92 conditional expression syntax.

Selecting on the content of a message

It is possible to subscribe based on a selection of message payload content (also known as content filtering), but the decision about which messages should be delivered to such a subscription cannot be performed directly by WebSphere MQ; instead an extended message selection provider, for example IBM Integration Bus, is required to process the messages.

Message selector syntax

A WebSphere MQ message selector is a string with syntax that is based on a subset of the SQL92 conditional expression syntax.

The order in which a message selector is evaluated is from left to right within a precedence level. You can use parentheses to change this order. Predefined selector literals and operator names are written here in uppercase; however, they are not case-sensitive.

WebSphere MQ verifies the syntactic correctness of a message selector at the time it is presented. If the syntax of the selection string is incorrect or a property name is not valid, and an extended message selection provider is not available, MQRC_SELECTION_NOT_AVAILABLE is returned to the application. If the syntax of the selection string is incorrect or a property name is not valid when a subscription is resumed, a MQRC_SELECTOR_SYNTAX_ERROR is returned to the application. If property name validation was disabled when the property was set (by setting MQCMHO_NONE instead of MQCMHO_VALIDATE) and an application subsequently puts a message with an invalid property name, this message is never selected.

A selector can contain:

- Literals:
 - String literals are enclosed in single quotation marks. Two consecutive single quotation marks represent a single quotation mark. Examples are 'literal' and 'literal's'. Like Java string literals, these use the Unicode character encoding. You cannot use double quotation marks to enclose a string literal. Any sequence of bytes can be used between the single quotation marks.
 - A byte string is one or more pair of hexadecimal characters enclosed in double quotation marks and prefixed by 0x. Examples are "0x2F1C" or "0XD43A". The length of a byte string must be at least one byte. If a selector byte string is matched to a message property of type MQTYPE_BYTE_STRING, no special action is taken on leading or trailing zero. The bytes are treated as another character. Endianness is also not considered. The length of both selector and property byte strings must be equal, and the sequence of bytes must be the same.

Examples of byte string selections (assume *myBytes* = 0AFC23) which match are:

- "myBytes = "0x0AFC23"" = TRUE

The following string selections do not match:

- "myBytes = "0xAFC23"" = MQRC_SELECTOR_SYNTAX_ERROR (because number of bytes is not multiple of two)

- "myBytes = "0x0AFC2300" " = FALSE (because the trailing zero is significant in the comparison)
- "myBytes = "0x000AFC23" " = FALSE (because leading zero is significant in the comparison)
- "myBytes = "0x23FC0A" " = FALSE (because endianness is not considered)
- Hex numbers begin with a zero, followed by an uppercase or lowercase x. The remainder of the literal contains one or more valid hex characters. Examples are 0xA, 0xAF, 0X2020.
- A leading zero followed by one or more digits in the range 0-7 is always interpreted as being the start of an octal number. You cannot represent a zero-prefixed decimal number like this, for example, 09 returns a syntax error because 9 is not a valid octal digit. Examples of octal numbers are 0177, 0713.
- An exact numeric literal is a numeric value without a decimal point, such as 57, -957, and +62. An exact numeric literal can have a trailing uppercase or lowercase L; this does not affect how the number is stored or interpreted. WebSphere MQ supports exact numerals in the range -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.
- An approximate numeric literal is a numeric value in scientific notation, such as 7E3 or -57.9E2, or a numeric value with a decimal, such as 7., -95.7, or +6.2. WebSphere MQ supports numbers in the range -1.797693134862315E+308 to 1.797693134862315E+308.

The significand should follow an optional sign character (+ or -). The significand should be either an integer or a fraction. A fractional part of the significand need not have a leading digit.

An uppercase or lowercase E indicates the start of an optional exponent. The exponent has a decimal radix and the number part of the exponent can be prefixed by an optional sign character.

Approximate numeric literals can be terminated by an F or a D character (not case-sensitive). This syntax exists to support the cross-language method of tagging single or double precision numbers. These characters are optional and do not affect how an approximate numeric literal is stored or processed. These numbers are always stored and processed using double-precision.

- The boolean literals TRUE and FALSE.

Note: Non-finite IEEE-754 representations such as NaN, +Infinity, -Infinity are not supported in selection strings. It is therefore not possible to use these values as operands in an expression. Negative zero is treated the same as positive zero for mathematical operations.

- Identifiers:

An identifier is a variable-length character sequence that must begin with a valid identifier start character, followed by zero or more valid identifier part characters. The rules for identifier names are the same as those for message property names, see [“Property names” on page 19](#) and [“Property name restrictions” on page 20](#) for more information.

Note: Selection can be performed on the message payload only if an extended message selection provider is available.

Identifiers are either header field references or property references. The type of a property value in a message selector must correspond to the type used to set the property, although numeric promotion is performed where possible. If a type mismatch occurs then the result of the expression is FALSE. If a property that does not exist in a message is referenced, its value is NULL.

Type conversions that apply to the get methods for properties do not apply when a property is used in a message selector expression. For example, if you set a property as a string value and then use a selector to query it as a numeric value, the expression returns FALSE.

JMS field and property names that map to property names or MQMD field names are also valid identifiers in a selection string. WebSphere MQ maps the recognized JMS field and property names to the message property values. See [“Message selectors in JMS” on page 773](#) for more information. As an example, the selection string "JMSPriority >=" selects on the Pri property found in the jms folder of the current message.

- Overflow/underflow:

For both decimal and approximate numeric numbers, the following are undefined:

- Specifying a number that is out of the defined range
- Specifying an arithmetic expression which would cause overflow or underflow

No checks are performed for these conditions.

- White space:

Defined as a space, form-feed, newline, carriage return, horizontal tab, or vertical tab. The following Unicode characters are recognized as white space:

- \u0009 to \u000D
- \u0020
- \u001C
- \u001D
- \u001E
- \u001F
- \u1680
- \u180E
- \u2000 to \u200A
- \u2028
- \u2029
- \u202F
- \u205F
- \u3000

- Expressions:

- A selector is a conditional expression. A selector that evaluates to true matches; a selector that evaluates to false or unknown does not match.
- Arithmetic expressions are composed of themselves, arithmetic operations, identifiers (identifier value is treated as a numeric literal), and numeric literals.
- Conditional expressions are composed of themselves, comparison operations, and logical operations.

- Standard bracketing (), to set the order in which expressions are evaluated, is supported.

- Logical operators in precedence order: NOT, AND, OR.

- Comparison operators: =, >, >=, <, <=, <> (not equal).

- Two byte strings are equal only if the strings are of the same length and the sequence of bytes is equal.
- Only values of the same type can be compared. One exception is that it is valid to compare exact numeric values and approximate numeric values, (the type conversion required is defined by the rules of Java numeric promotion). If there is an attempt to compare different types, the selector is always false.
- String and boolean comparison is restricted to = and <>. Two strings are equal only if they contain the same sequence of characters.

- Arithmetic operators in precedence order:

- +, - unary.
- * multiplication, and / division.
- + addition, and - subtraction.
- Arithmetic operations on a NULL value are not supported. If they are attempted, the complete selector is always false.
- Arithmetic operations must use Java numeric promotion.

- arithmetic-expr1 [NOT] BETWEEN arithmetic-expr2 and arithmetic-expr3 comparison operator:

- Age BETWEEN 15 and 19 is equivalent to age >= 15 AND age <= 19.
- Age NOT BETWEEN 15 and 19 is equivalent to age < 15 OR age > 19.
- If any of the expressions of a BETWEEN operation are NULL, the value of the operation is false. If any of the expressions of a NOT BETWEEN operation are NULL, the value of the operation is true.
- identifier [NOT] IN (string-literal1, string-literal2,...) comparison operator where identifier has a String or NULL value.
 - Country IN ('UK', 'US', 'France') is true for 'UK' and false for 'Peru'. It is equivalent to the expression (Country = 'UK') OR (Country = 'US') OR (Country = 'France').
 - Country NOT IN ('UK', 'US', 'France') is false for 'UK' and true for 'Peru'. It is equivalent to the expression NOT ((Country = 'UK') OR (Country = 'US') OR (Country = 'France')).
 - If the identifier of an IN or NOT IN operation is NULL, the value of the operation is unknown.
- identifier [NOT] LIKE *pattern-value* [ESCAPE *escape-character*] comparison operator, where identifier has a string value. *pattern-value* is a string literal, where _ stands for any single character and % stands for any sequence of characters (including the empty sequence). All other characters stand for themselves. The optional *escape-character* is a single character string literal that is used to escape the special meaning of the _ and % in *pattern-value*. The LIKE operator must be used only to compare two string values.
 - phone LIKE '12%3' is true for 123 and 12993 and false for 1234.
 - word LIKE 'l_se' is true for lose and false for loose.
 - underscored LIKE '_%' ESCAPE '\' is true for _foo and false for bar.
 - phone NOT LIKE '12%3' is false for 123 and 12993 and true for 1234.
 - If the identifier of a LIKE or NOT LIKE operation is NULL, the value of the operation is unknown.

Note: The LIKE operator must be used to compare two string values. The value of Root.MQMD.CorrelId is a 24-byte byte array, not a character string. The selector string Root.MQMD.CorrelId LIKE 'ABC%' is accepted by the parser as syntactically valid, but it is evaluated to false. When you are comparing a byte array with a character string, LIKE therefore cannot be used.

- identifier IS NULL comparison operator tests for a NULL header field value, or a missing property value.
- identifier IS NOT NULL comparison operator tests for the existence of a non-null header field value or a property value.
- Null values

The evaluation of selector expressions that contain NULL values is defined by SQL 92 NULL semantics, in summary:

- SQL treats a NULL value as unknown.
- Comparison or arithmetic with an unknown value always yields an unknown value.
- The IS NULL and IS NOT NULL operators convert an unknown value into TRUE and FALSE values.

The boolean operators use three-valued logic (T=TRUE, F=FALSE, U=UNKNOWN)

Table 1. Boolean operator outcome when logic is A AND B		
Operator A	Operator B	Outcome (A AND B)
T	F	F
T	U	U
T	T	T
F	T	F

Table 1. Boolean operator outcome when logic is A AND B (continued)		
Operator A	Operator B	Outcome (A AND B)
F	U	F
F	F	F
U	T	U
U	U	U
U	F	F

Table 2. Boolean operator outcome when logic is A OR B		
Operator A	Operator B	Outcome (A OR B)
T	F	T
T	U	T
T	T	T
F	T	T
F	U	U
F	F	F
U	T	T
U	U	U
U	F	U

Table 3. Boolean operator outcome when logic is NOT A	
Operator A	Outcome (NOT A)
T	F
F	T
U	U

The following message selector selects messages with a message type of car, color of blue, and weight greater than 2500 lbs:

```
"JMSType = 'car' AND color = 'blue' AND weight > 2500"
```

Although SQL supports fixed decimal comparison and arithmetic, message selectors do not. This is why exact numeric literals are restricted to those without a decimal. It is also why there are numerics with a decimal as an alternative representation for an approximate numeric value.

SQL comments are not supported.

Related concepts

[Selection behavior](#)

Overview of IBM WebSphere MQ selection behavior.

[Selecting on the content of a message](#)

It is possible to subscribe based on a selection of message payload content (also known as content filtering), but the decision about which messages should be delivered to such a subscription cannot be

performed directly by WebSphere MQ; instead an extended message selection provider, for example IBM Integration Bus, is required to process the messages.

“Message properties” on page 18

Use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing MQMD or MQRFH2 headers. They also facilitate communication between WebSphere MQ and JMS applications.

Related reference

MsgHandle

MQBUFMH - Convert buffer into message handle

Selection string rules and restrictions

Familiarize yourself with these rules about how selection strings are interpreted and character restrictions to avoid potential problems when using selectors.

- Equivalence is tested using a single equals character; for example, `a = b` is correct, whereas `a == b` is incorrect.
- An operator used by many programming languages to represent 'not equal to' is `!=`. This representation is not a valid synonym for `<>`; for example, `a <> b` is valid, whereas `a != b` is not valid.
- Single quotation marks are recognized only if the `'` (U+0027) character is used. Similarly, double quotation marks, valid only when used to enclose byte strings, must use the `"` (U+0022) character.
- The symbols `&`, `&&`, `|` and `||` are not synonyms for logical conjunction/disjunction; for example, `a && b` must be specified as `a AND b`.
- The wildcard characters `*` and `?` are not synonyms for `%` and `_`.
- Selectors containing compound expressions such as `20 < b < 30` are not valid. The parser evaluates operators that have the same precedence from left to right. The example would therefore become `(20 < b) < 30`, which does not make sense. Instead the expression must be written as `(b > 20) AND (b < 30)`.
- Byte strings must be enclosed in double quotation marks; if single quotation marks are used, the byte string is taken to be a string literal. The number of characters (not the number that the characters represent) following the `0x` must be a multiple of two.
- The keyword `IS` is not a synonym for the equals character. Thus the selection strings `a IS 3` and `b IS 'red'` are not valid. The `IS` keyword exists only to support `IS NULL` and `IS NOT NULL` cases.

Related concepts

UTF-8 and Unicode considerations when using message selectors

UTF-8 and Unicode considerations when using message selectors

Characters, not enclosed in single quotation marks, that make up the reserved keywords of a selection string must be entered in Basic Latin Unicode (ranging from character U+0000 to U+0007F). It is not valid to use other code point representations of alphanumeric characters. For example, the number 1 must be expressed as U+0031 in Unicode, it is not valid to use the Fullwidth Digit equivalent U+FF11 or the Arabic equivalent U+0661.

Message property names can be specified using any valid sequence of Unicode characters. Message property names contained within selection strings that are encoded in UTF-8 will be validated even if they contain multi-byte characters. Validation of multi-byte UTF-8 is strict and you must ensure that valid UTF-8 sequences are used for message property names.

No extra processing is performed on property names or values when comparing for equality. This means for example that no pre/de-composition takes place and ligatures are not given any special meaning. For example, the pre-composed umlaut character U+00FC is not considered to be equivalent to U+0075 + U+0308 and the character sequence `ff` is not considered to be equivalent to the Unicode U+FB00 (LATIN SMALL LIGATURE FF)

Property data enclosed in single quotation marks can be represented by any sequence of bytes and is not validated.

Related concepts

Selection string rules and restrictions

Familiarize yourself with these rules about how selection strings are interpreted and character restrictions to avoid potential problems when using selectors.

Selecting on the content of a message

It is possible to subscribe based on a selection of message payload content (also known as content filtering), but the decision about which messages should be delivered to such a subscription cannot be performed directly by WebSphere MQ; instead an extended message selection provider, for example IBM Integration Bus, is required to process the messages.

When an application publishes on a topic string, where one or more subscribers have a selection string selecting on the content of the message, WebSphere MQ will request that the extended message selection provider parse the publication and inform WebSphere MQ whether the publication matches the selection criteria specified by each subscriber with a content filter.

If the extended message selection provider determines that the publication matches the subscriber's selection string, the message will continue to be delivered to the subscriber.

If the extended message selection provider determines that the publication does not match, the message is not delivered to the subscriber. This might cause the MQPUT or MQPUT1 call to fail with reason code MQRC_PUBLICATION_FAILURE. If the extended message selection provider is unable to parse the publication, reason code MQRC_CONTENT_ERROR is returned and the MQPUT or MQPUT1 call fails.

If the extended message selection provider is unavailable or is unable to determine whether the subscriber should receive the publication, reason code MQRC_SELECTION_NOT_AVAILABLE is returned and the MQPUT or MQPUT1 call fails.

When a subscription is being created with a content filter and the extended message selection provider is not available, the MQSUB call fails with reason code MQRC_SELECTION_NOT_AVAILABLE. If a subscription with a content filter is being resumed and the extended message selection provider is not available, the MQSUB call returns a warning of MQRC_SELECTION_NOT_AVAILABLE, but the subscription is allowed to be resumed.

Related concepts

Selection behavior

Overview of IBM WebSphere MQ selection behavior.

Message selector syntax

A WebSphere MQ message selector is a string with syntax that is based on a subset of the SQL92 conditional expression syntax.

Asynchronous consumption of IBM WebSphere MQ messages

Asynchronous consumption uses a set of Message Queue Interface (MQI) extensions, the MQI calls MQCB and MQCTL, which allow an MQI application to be written to consume messages from a set of queues. Messages are delivered to the application by invoking a 'unit of code', identified by the application passing either the message, or a token representing the message.

In the most straightforward of application environments, the 'unit of code' is defined by a function pointer, however in other environments the 'unit of code' can be defined by a program or module name.

In asynchronous consumption of messages, the following terms are used:

Message consumer

A programming construct that allows you to define a program, or function, to be invoked with a message when one which matches the applications requirement becomes available.

Event handler

A programming construct that allows you to define a program or function to invoke when an asynchronous event, such as queue manager quiescing, occurs.

Callback

A generic term used to refer to either a Message Consumer or an Event Handler routine.

Asynchronous consumption can simplify the design and implementation of new applications, especially those that process multiple input queues or subscriptions. However, if you are using more than one input queue and you are processing messages in priority sequence, priority sequence is observed independently within each queue: you might get low-priority messages from one queue ahead of high-priority messages from another. Message order across multiple queues is not guaranteed. Also note that if you use API exits, you might need to change them to include the MQCB and MQCTL calls.

The following illustrations give an example of how you can use this function.

Figure 5 on page 33 shows a multithreaded application consuming messages from two queues. The example shows all of the messages being delivered to a single function.

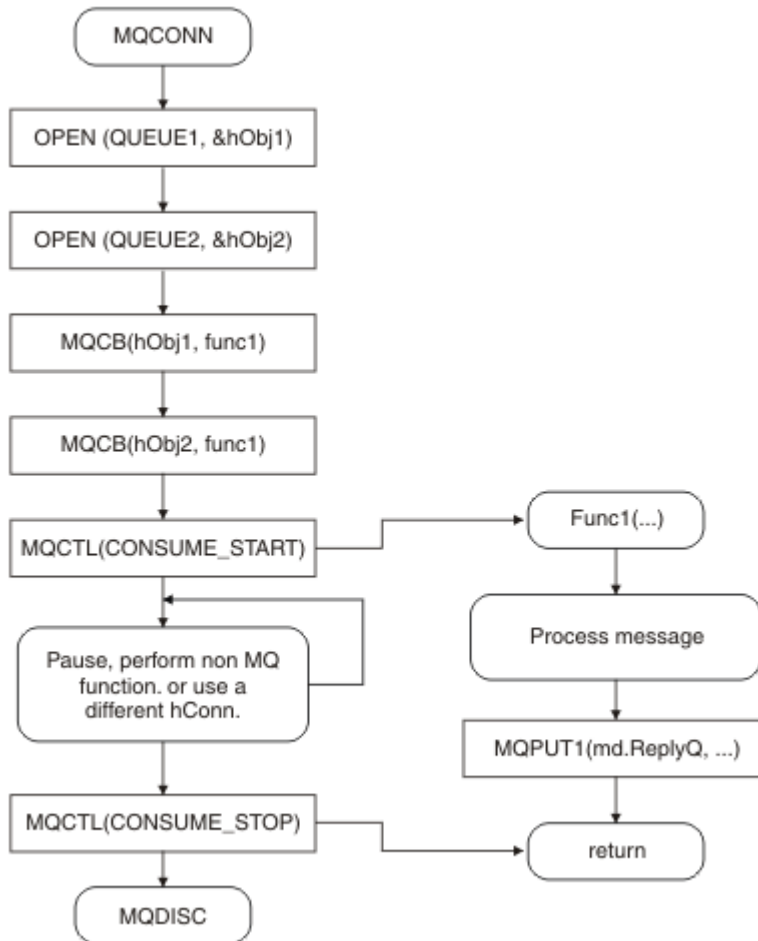


Figure 5. Standard Message Driven application consuming from two queues

Figure 6 on page 34 This sample flow shows a single threaded application consuming messages from two queues. The example shows all of the messages being delivered to a single function.

The difference from the asynchronous case is that control does not return to the issuer of MQCTL until all of the consumers have deactivated themselves; that is one consumer has issued an MQCTL STOP request or the queue manager quiesces.

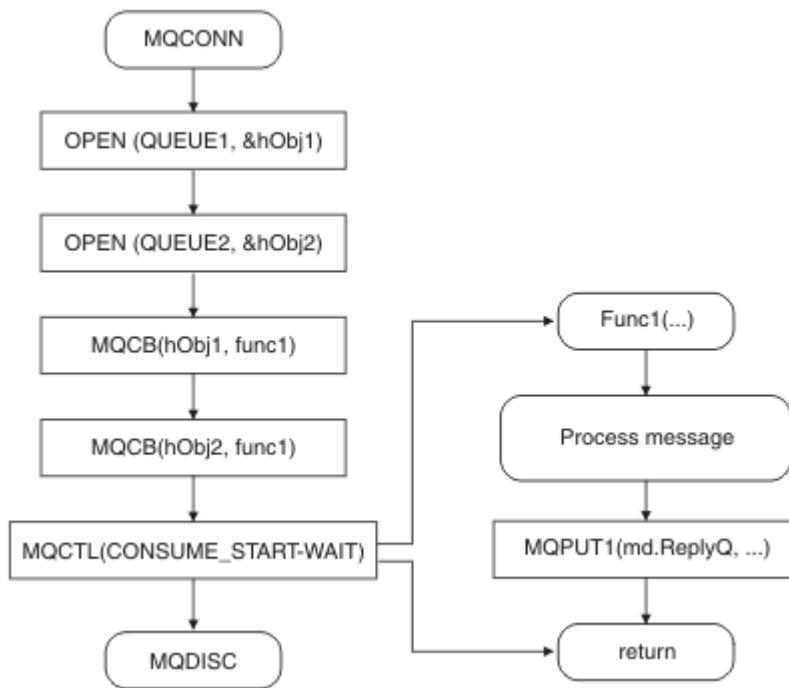


Figure 6. Single Threaded Message Driven application consuming from two queues

Message groups

Messages can occur within groups to allow ordering of messages.

Message groups allow multiple messages to be marked as related to one another, and a logical order to be applied to the group (see “[Logical and physical ordering](#)” on page 236). On platforms other than z/OS, a related concept, “[Message segmentation](#)” on page 252 enables large messages to be broken up into smaller segments. You cannot use grouped or segmented messages when putting to a topic.

The hierarchy within a group is as follows:

Group

This is the highest level in the hierarchy and is identified by a *GroupId*. It consists of one or more messages that contain the same *GroupId*. These messages can be stored anywhere on the queue.

Note: The term *message* is used here to denote one item on a queue, such as would be returned by a single MQGET that does not specify MQGMO_COMPLETE_MSG.

Figure 7 on page 34 shows a group of logical messages:




Figure 7. Group of logical messages

By opening a queue and specifying MQOO_BIND_ON_GROUP, you force all messages in a group that are sent to this queue to be sent to the same instance of the queue. For more information on the BIND_ON_GROUP option, see [Handling message affinities](#).

Logical message

Logical messages within a group are identified by the *GroupId* and *MsgSeqNumber* fields. The *MsgSeqNumber* starts at 1 for the first message within a group, and if a message is not in a group, the value of the field is 1.

Use logical messages within a group to:

- Ensure ordering (if this is not guaranteed under the circumstances in which the message is transmitted).
- Allow applications to group similar messages (for example, those that must all be processed by the same server instance).

Each message within a group consists of one physical message, unless it is split into segments. Each message is logically a separate message, and only the *GroupId* and *MsgSeqNumber* fields in the MQMD need bear any relationship to other messages in the group. Other fields in the MQMD are independent; some might be identical for all messages in the group whereas others might be different. For example, messages in a group can have different format names, CCSIDs, and encodings.

Segment

Segments are used to handle messages that are too large for either the putting or getting application or the queue manager (including intervening queue managers through which the message passes). For more information, see [“Message segmentation” on page 252](#).

An individual message is broken down into smaller messages called *segments*. A segment of a message is identified by the *GroupId*, *MsgSeqNumber*, and *Offset* fields. The *Offset* field starts at zero for the first segment within a message.

Each segment consists of one physical message that might belong to a group (Figure 8 on page 35 shows an example of messages within a group). A segment is logically part of a single message, so only the *MsgId*, *Offset*, and *SegmentFlag* fields in the MQMD should differ between separate segments of the same message. If a segment fails to arrive, reason code [MQRC_INCOMPLETE_GROUP](#) or [MQRC_INCOMPLETE_MSG](#) is returned as appropriate.

Figure 8 on page 35 shows a group of logical messages, some of which are segmented:




Figure 8. Segmented messages

You cannot use segmented or grouped messages with Publish/Subscribe.

For a description of logical and physical messages, see [“Logical and physical ordering” on page 236](#). For further information about segmenting messages, see [“Message segmentation” on page 252](#).

Message persistence

Persistent messages are written to logs and queue data files.

If a queue manager is restarted after a failure, it recovers these persistent messages as necessary from the logged data. Messages that are not persistent are discarded if a queue manager stops, whether the stoppage is as a result of an operator command or because of the failure of some part of your system.

When you create a message, if you initialize the message descriptor (MQMD) using the defaults, the persistence for the message is taken from the *DefPersistence* attribute of the queue specified in the MQOPEN command. Alternatively, you can set the persistence of the message using the *Persistence* field of the MQMD structure to define the message as persistent or nonpersistent.

The performance of your application is affected when you use persistent messages; the extent of the effect depends on the performance characteristics of the machine's I/O subsystem and how you use the sync point options on each platform:

- A persistent message, outside the current unit of work, is written to disk on every put and get operation. See [“Committing and backing out units of work”](#) on page 309.
- In IBM WebSphere MQ on UNIX systems, IBM WebSphere MQ on Linux systems, and IBM WebSphere MQ for Windows, a persistent message within the current unit of work is logged only when the unit of work is committed (and the unit of work could contain many queue operations).

Nonpersistent messages can be used for fast messaging. See [Safety of messages](#) for further information about fast messages.

Note: A combination of writing persistent messages within a unit of work, and writing persistent messages outside a unit of work, can cause potentially severe performance problems for your applications. This is particularly true when the same target queue is used for both operations.

Messages that fail to be delivered

When a queue manager cannot put a message on a queue, you have various options.

You can:

- Attempt to put the message on the queue again.
- Request that the message is returned to the sender.
- Put the message on the dead-letter queue.

See [“Handling program errors”](#) on page 523 for more information.

Messages that are backed out

When processing messages from a queue under the control of a unit of work, the unit of work can consist of one or more messages. If a backout occurs, the messages that were retrieved from the queue are reinstated on the queue, and they can be processed again in another unit of work. If the processing of a particular message is causing the problem, the unit of work is backed out again. This can cause a processing loop. Messages that were put to a queue are removed from the queue.

An application can detect messages that are caught up in such a loop by testing the *BackoutCount* field of MQMD. The application can either correct the situation, or issue a warning to an operator.

On WebSphere MQ for WebSphere MQ for Windows, WebSphere MQ on UNIX systems, WebSphere MQ on Linux systems the backout count always survives restarts of the queue manager. Any change to the *HardenGetBackout* attribute is ignored.

For more information on committing and backing out messages, see [“Committing and backing out units of work”](#) on page 309.

Reply-to queue and queue manager

There are occasions when you might receive messages in response to a message you send:

- A reply message in response to a request message
- A report message about an unexpected event or expiry
- A report message about a COA (Confirmation Of Arrival) or a COD (Confirmation Of Delivery) event
- A report message about a PAN (Positive Action Notification) or a NAN (Negative Action Notification) event

Using the MQMD structure, specify the name of the queue to which you want reply and report messages sent, in the *ReplyToQ* field. Specify the name of the queue manager that owns the reply-to queue in the *ReplyToQMGr* field.

If you leave the *ReplyToQMGr* field blank, the queue manager sets the contents of the following fields in the message descriptor on the queue:

ReplyToQ

If *ReplyToQ* is a local definition of a remote queue, the *ReplyToQ* field is set to the name of the remote queue; otherwise this field is not changed.

ReplyToQMgr

If *ReplyToQ* is a local definition of a remote queue, the *ReplyToQMgr* field is set to the name of the queue manager that owns the remote queue; otherwise the *ReplyToQMgr* field is set to the name of the queue manager to which your application is connected.

Note: You can request that a queue manager makes more than one attempt to deliver a message, and you can request that the message is discarded if it fails. If the message, after failing to be delivered, is not to be discarded, the remote queue manager puts the message on its dead-letter (undelivered message) queue (see [“Using the dead-letter \(undelivered message\) queue”](#) on page 526).

Message context

Message context information allows the application that retrieves the message to find out about the originator of the message.

The retrieving application might want to:

- Check that the sending application has the correct level of authority
- Perform some accounting function so that it can charge the sending application for any work that it has to perform
- Keep an audit trail of all the messages that it has worked with

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. For more information about how to specify context information, see [“Controlling context information”](#) on page 222.

The user context is used by the queue manager when generating the following types of report message:

- Confirm on delivery
- Expiry

When these report messages are generated, the user context is checked for +put and +passid authority on the destination of the report. Where the user context has insufficient authority, the report message is placed on the dead-letter queue if one has been defined. Where there is no dead-letter queue, the report message is discarded.

All context information is stored in the context fields of the message descriptor. The type of information falls into identity, origin, and user context information.

Identity context

Identity context information identifies the user of the application that first put the message on a queue. Suitably authorized applications can set the following fields:

- The queue manager fills the *UserIdentifier* field with a name that identifies the user; the way that the queue manager can do this depends on the environment in which the application is running.
- The queue manager fills the *AccountingToken* field with a token or number that it determined from the application that put the message.
- Applications can use the *ApplIdentityData* field for any extra information that they want to include about the user (for example, an encrypted password).

A Windows systems security identifier (SID) is stored in the *AccountingToken* field when a message is created under WebSphere MQ for Windows. The SID can be used to supplement the *UserIdentifier* field and to establish the credentials of a user.

For information about how the queue manager fills the *UserIdentifier* and *AccountingToken* fields, see the descriptions of these fields in [UserIdentifier](#) and [AccountingToken](#).

Applications that pass messages from one queue manager to another should also pass on the identity context information so that other applications know the identity of the originator of the message.

Origin context

Origin context information describes the application that put the message on the queue on which the message is *currently* stored. The message descriptor contains the following fields for origin context information:

<i>PutApplType</i>	The type of application that put the message (for example, a CICS transaction).
<i>PutApplName</i>	The name of the application that put the message (for example, the name of a job or transaction).
<i>PutDate</i>	The date on which the message was put on the queue.
<i>PutTime</i>	The time at which the message was put on the queue.
<i>ApplOriginData</i>	Any extra information that an application wants to include about the origin of the message. For example, it could be set by suitably authorized applications to indicate whether the identity data is trusted.

Origin context information is typically supplied by the queue manager. Greenwich Mean Time (GMT) is used for the *PutDate* and *PutTime* fields. See the descriptions of these fields in [PutDate](#) and [PutTime](#).

An application with enough authority can provide its own context. This allows accounting information to be preserved when a single user has a different user ID on each of the systems that process a message that they have originated.

WebSphere MQ objects

This information provides details on WebSphere MQ objects which include: queue managers, queue-sharing groups, queues, administrative topic objects, namelists, process definitions, authentication information objects, channels, storage classes, listeners, and services.

Queue managers define the properties (known as attributes) of these objects. The values of these attributes affect the way in which WebSphere MQ processes these objects. From your applications, you use the Message Queue Interface (MQI) to control these objects. Objects are identified by an *object descriptor* (MQOD) when addressed from a program.

When you use WebSphere MQ commands to define, alter, or delete objects, for example, the queue manager checks that you have the required level of authority to perform these operations. Similarly, when an application uses the MQOPEN call to open an object, the queue manager checks that the application has the required level of authority before it allows access to that object. The checks are made on the name of the object being opened.

Related concepts

[“Controlling context information” on page 222](#)

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

Related reference

[“MQOPEN options relating to message context” on page 213](#)

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

Preparing and running Microsoft Transaction Server applications

To prepare an MTS application to run as a WebSphere MQ MQI client application, follow these instructions as appropriate for your environment.

For general information about how to develop Microsoft Transaction Server (MTS) applications that access WebSphere MQ resources, see the section on MTS in the WebSphere MQ Help Center.

To prepare an MTS application to run as a WebSphere MQ MQI client application, do one of the following for each component of the application:

- If the component uses the C language bindings for the MQI, follow the instructions in [“Preparing C programs in Windows” on page 439](#) but link the component with the library mqicxa.lib instead of mqic.lib.
- If the component uses the WebSphere MQ C++ classes, follow the instructions in [“Building C++ programs on Windows” on page 624](#) but link the component with the library imqcx23vn.lib instead of imqc23vn.lib.
- If the component uses the Visual Basic language bindings for the MQI, follow the instructions in the [“Preparing Visual Basic programs in Windows” on page 443](#) but when you define the Visual Basic project, type MqType=3 in the **Conditional Compilation Arguments** field.
- If the component uses the WebSphere MQ Automation Classes for ActiveX (MQAX), define an environment variable, GMQ_MQ_LIB, with the value mqic32xa.dll .

You can define the environment variable from within your application, or you can define it so that its scope is system wide. However, defining it as system wide can cause any existing MQAX application, that does not define the environment variable from within the application, to behave incorrectly.

Using IBM WebSphere MQ with WebSphere Application Server

Use this topic to understand the use of IBM WebSphere MQ with WebSphere Application Server.

Applications written in Java that are running under WebSphere Application Server can use the Java Messaging Service (JMS) specification to perform messaging. Point-to-point messaging in this environment may be provided by an IBM WebSphere MQ queue manager

A benefit of using an IBM WebSphere MQ queue manager to provide the point-to-point messaging is that connecting JMS applications can participate fully in the functionality of an IBM WebSphere MQ network, which allows the applications to exchange messages with queue managers running on a multitude of platforms.

Applications can use either the *client transport* or *bindings transport* for the queue connection factory object. For *bindings transport* the queue manager must exist locally to the application which requires a connection. If the queue manager is not local to the application, then the *Client Attachment* should be installed to allow the application to connect to a queue manager running on another machine or image.

By default, JMS messages held on IBM WebSphere MQ queues use an MQRFH2 header to hold some of the JMS message header information. Many legacy IBM WebSphere MQ applications cannot process messages with these headers, and require their own characteristic headers, for example the MQCIH for CICS Bridge, or MQWIH for IBM WebSphere MQ Workflow applications. For more details about these special considerations, see [“Mapping JMS messages onto WebSphere MQ messages” on page 775](#).

Transactional support scenarios

Using transactional support you can enable your applications to work reliably with databases.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

This section introduces transactional support. The work required to enable your applications to use IBM WebSphere MQ with a database product spans the areas of application programming and system administration. Use the information here together with [“Committing and backing out units of work” on page 309](#).

We start by introducing the units of work that form transactions, then describe the ways in which you enable IBM WebSphere MQ to coordinate transactions with databases.

Related concepts

[“Introducing units of work” on page 40](#)

This topic introduces and defines the general concepts of unit of work, commit, backout, and sync point. It also contains two scenarios that illustrate global units of work.

[IBM WebSphere MQ and HP NonStop TMF](#)

Introducing units of work

This topic introduces and defines the general concepts of unit of work, commit, backout, and sync point. It also contains two scenarios that illustrate global units of work.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

When a program puts messages on queues within a unit of work, those messages are made visible to other programs only when the program *commits* the unit of work. To commit a unit of work, all updates must be successful to preserve data integrity.

If the program detects an error and decides not to make the put operation permanent, it can *back out* the unit of work. When a program performs a backout, WebSphere MQ restores the queues by removing the messages that were put on the queues by that unit of work.

Similarly, when a program gets messages from one or more queues within a unit of work, those messages remain on the queues until the program commits the unit of work, but the messages are not available to be retrieved by other programs. The messages are permanently deleted from the queues when the program commits the unit of work. If the program backs out the unit of work, WebSphere MQ restores the queues by making the messages available to be retrieved by other programs.

The decision to commit or back out the changes is taken, in the simplest case, at the end of a task. However, it can be more useful for an application to synchronize data changes at other logical points within a task. These logical points are called sync points (or synchronization points) and the period of processing a set of updates between two sync points is called a *unit of work*. Several MQGET calls and MQPUT calls can be part of a single unit of work.

With WebSphere MQ, we need to distinguish between *local* and *global* units of work:

Local units of work

Are those in which the only actions are puts to, and gets from, WebSphere MQ queues, and the coordination of each unit of work is provided within the queue manager using a *single-phase commit* process.

Use local units of work when the only resources to be updated are the queues that are managed by a single WebSphere MQ queue manager. Updates are committed by using the MQCMIT verb or backed out using MQBACK.

There are no system administration tasks, other than log management, which is involved in using local units of work. In your applications, where you use the MQPUT and MQGET calls with MQCMIT and MQBACK, try using the MQPMO_SYNCPOINT and MQGMO_SYNCPOINT options. (For information about log management, see [Managing log files](#).)

Global units of work

Are those in which other resources, such as tables in a relational database, are also updated. When more than one *resource manager* is involved, there is a need for *transaction manager* software that uses a *two-phase commit* process to coordinate the global unit of work.

Use global units of work when you also need to include updates to relational database manager software, such as Db2, Oracle, Sybase, and Informix.

There are several possible scenarios for using global units of work. Documented here are two scenarios:

1. In the first, the queue manager itself acts as the transaction manager. In this scenario, MQI verbs control the global units of work; they are started in applications using the MQBEGIN verb and then committed using MQCMIT or backed out using MQBACK.
2. In the second, the transaction manager role is performed by other software, such as TXSeries, Encina, or Tuxedo. In this scenario, an API provided by the transaction manager software is used to control the unit of work (for example, EXEC CICS SYNCPOINT for TXSeries).

The following sections describe all the steps necessary to use global units of work, organized by the two scenarios:

- [“Scenario 1: Queue manager performs the coordination” on page 41](#)
- [“Scenario 2: Other software provides the coordination” on page 66](#)

Scenario 1: Queue manager performs the coordination

In scenario 1, the queue manager acts as the transaction manager. In this scenario, MQI verbs control the global units of work; they are started in applications using the MQBEGIN verb and then committed using MQCMIT or backed out using MQBACK.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Isolation level

In IBM WebSphere MQ, a message on a queue might be visible before a database update, depending on the transaction isolation design implemented within the database.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

When an IBM WebSphere MQ queue manager is working as an XA transaction manager, to coordinate updates to XA resource managers, the following commit protocol is followed:

1. Prepare all XA resource managers.
2. Commit the IBM WebSphere MQ queue manager resource manager.
3. Commit other resource managers.

Between step 2 and 3, an application might see a message that is committed to the queue but the corresponding row in the database does not reflect this message.

This is not a problem if the database is configured such that the application's database API calls wait for pending updates to be completed.

You can resolve this by configuring the database differently. The type of configuration needed is referred to as the "isolation level". For more information on isolation levels, refer to the database documentation. You can, alternatively, configure the queue manager to commit the resource managers in the following reverse order:

1. Prepare all XA resource managers.
2. Commit other resource managers.
3. Commit the IBM WebSphere MQ queue manager resource manager.

When you change the protocol the IBM WebSphere MQ queue manager is committed last, so applications that read messages from the queues see a message only after the corresponding database update has been completed.

To configure the queue manager to use this changed protocol, set the **AMQ_REVERSE_COMMIT_ORDER** environment variable.

Set this environment variable in the environment from which the **strmqm** is run to start the queue manager. For example, run the following in the shell just before starting the queue manager:

```
export AMQ_REVERSE_COMMIT_ORDER=1
```

Note: Setting this environment variable might cause an extra log entry per transaction, so this will have a small impact on the performance of each transaction.

Database coordination

When the queue manager coordinates global units of work itself, it becomes possible to integrate database updates within the units of work. That is, a mixed MQI and SQL application can be written, and the MQCMIT and MQBACK verbs can be used to commit or roll back the changes to the queues and databases together.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The queue manager achieves this using the two-phase commit protocol described in *X/Open Distributed Transaction Processing: The XA Specification*. When a unit of work is to be committed, the queue manager first asks each participating database manager whether it is prepared to commit its updates. Only if all the participants, including the queue manager itself, are prepared to commit, are all the queue and database updates committed. If any participant cannot prepare its updates, the unit of work is backed out instead.

In general, a global unit of work is implemented in an application by the following method (in pseudocode):

```
MQBEGIN
MQGET (include the flag MQGMO_SYNCPOINT in the message options)
MQPUT (include the flag MQPMO_SYNCPOINT in the message options)
SQL INSERT
MQCMIT
```

The purpose of MQBEGIN is to denote the beginning of a global unit of work. The purpose of MQCMIT is to denote the end of the global unit of work, and to complete it with all participating resource managers, using the two-phase commit protocol.

When the unit of work (also known as a *transaction*) is completed successfully using MQCMIT, all actions taken within that unit of work are made permanent or irreversible. If, for any reason, the unit of work fails, all actions are instead backed out. It is not possible for one action in a unit of work to be made permanent while another is backed out. This is the principle of a unit of work: either all actions within the unit of work are made permanent or none of them are.

Note:

1. The application programmer can force a unit of work to be backed out by calling MQBACK. The unit of work is also backed out by the queue manager if the application or database *fails* before MQCMIT is called.
2. If an application calls MQDISC without calling MQCMIT, the queue manager behaves as if MQCMIT had been called, and commits the unit of work.

In between MQBEGIN and MQCMIT, the queue manager does not make any calls to the database to update its resources. That is, the only way a database's tables are changed is by your code (for example, the SQL INSERT in the pseudocode).

Full recovery support is provided if the queue manager loses contact with any of the database managers during the commit protocol. If a database manager becomes unavailable while it is in doubt, that is, it has successfully prepared to commit, but has yet to receive a commit or backout decision, the queue manager remembers the outcome of the unit of work until that outcome has been successfully delivered to the database. Similarly, if the queue manager terminates with incomplete commit operations outstanding,

these are remembered over queue manager restart. If an application terminates unexpectedly, the integrity of the unit of work is not compromised, but the outcome depends on where in the process the application terminated, as described in [Table 5](#) on page 43.

What happens when the database or application program fails is summarized in the following tables:

<i>Table 4. What happens when a database server fails</i>	
Failure occurrence	Outcome
Before the application call to MQCMIT.	The unit of work is backed out.
During the application call to MQCMIT, before all databases have indicated that they have successfully prepared.	The unit of work is backed out with a reason code of MQRC_BACKED_OUT.
During the application call to MQCMIT, after all databases have indicated that they have successfully prepared, but before all have indicated that they have successfully committed.	The unit of work is held in recoverable state by the queue manager, with a reason code of MQRC_OUTCOME_PENDING.
During the application call to MQCMIT, after all databases have indicated that they have successfully committed.	The unit of work is committed with a reason code of MQRC_NONE.
After the application call to MQCMIT.	The unit of work is committed with a reason code of MQRC_NONE.

<i>Table 5. What happens when an application program fails</i>	
Failure occurrence	Outcome
Before the application call to MQCMIT.	The unit of work is backed out.
During the application call to MQCMIT, before the queue manager has received the application's MQCMIT request.	The unit of work is backed out.
During the application call to MQCMIT, after the queue manager has received the application's MQCMIT request.	The queue manager tries to commit using two-phase commit (subject to the database products successfully executing and committing their parts of the unit of work).

In the case where the reason code on return from MQCMIT is MQRC_OUTCOME_PENDING, the unit of work is remembered by the queue manager until it has been able to reestablish contact with the database server, and tell it to commit its part of the unit of work. Refer to [“Considerations when contact is lost with the XA resource manager”](#) on page 59 for information on how and when recovery is done.

The queue manager communicates with database managers using the XA interface as described in *X/Open Distributed Transaction Processing: The XA Specification*. Examples of these function calls are xa_open, xa_start, xa_end, xa_prepare, and xa_commit. We use the terms *transaction manager* and *resource manager* in the same sense as they are used in the XA specification.

Restrictions

There are restrictions to the database coordination support.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The following restrictions apply:

- The ability to coordinate database updates within WebSphere MQ units of work is **not** supported in an MQI client application. The use of MQBEGIN in a client application fails. A program that calls MQBEGIN must run as a *server* application on the same machine as the queue manager.

Note: A *server* application is a program that has been linked with the necessary WebSphere MQ server libraries; a *client* application is a program that has been linked with the necessary WebSphere MQ client libraries. See “Building applications for WebSphere MQ MQI clients” on page 341 and “Building an IBM WebSphere MQ application” on page 409 for details on compiling and linking your programs.

- The database server can reside on a different machine from the queue manager server, as long as the database client is installed on the same machine as the queue manager, and it supports this function. Consult the database product's documentation to determine whether their client software can be used for two-phase commit systems.
- Although the queue manager behaves as a resource manager (for the purposes of being involved in Scenario 2 global units of work), it is not possible to make one queue manager coordinate another queue manager within its Scenario 1 global units of work.

Switch load files

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The switch load file is a shared library (a DLL on Windows systems) that is loaded by the code in your IBM WebSphere MQ application and the queue manager. Its purpose is to simplify the loading of the database's client shared library, and to return the pointers to the XA functions.

The details of the switch load file must be specified before the queue manager is started. The details are placed in the qm.ini file on Windows, UNIX and Linux systems.

- On Windows and Linux (x86 and x86-64 platforms) systems, use the IBM WebSphere MQ Explorer to update the qm.ini file.
- On all other systems edit the file, qm.ini, directly.

The C source for the switch load file is supplied with the IBM WebSphere MQ installation if it supports Scenario 1 global units of work. The source contains a function called MQStart. When the switch load file is loaded, the queue manager calls this function, which returns the address of a structure called an XA *switch*.

The XA switch structure exists in the database client shared library, and contains a number of function pointers, as described in Table 6 on page 44:

Table 6. XA switch function pointers		
Function pointer name	XA function	Purpose
xa_open_entry	xa_open	Connect to database
xa_close_entry	xa_close	Disconnect from database
xa_start_entry	xa_start	Start a branch of a global unit of work
xa_end_entry	xa_end	Suspend a branch of a global unit of work
xa_rollback_entry	xa_rollback	Roll back a branch of a global unit of work
xa_prepare_entry	xa_prepare	Prepare to commit a branch of a global unit of work
xa_commit_entry	xa_commit	Commit a branch of a global unit of work

Table 6. XA switch function pointers (continued)		
Function pointer name	XA function	Purpose
xa_recover_entry	xa_recover	Discover from the database whether it has an in-doubt unit of work
xa_forget_entry	xa_forget	Allow a database to forget a branch of a global unit of work
xa_complete_entry	xa_complete	Complete a branch of a global unit of work

During the first MQBEGIN call in your application, the IBM WebSphere MQ code that executes as part of MQBEGIN loads the switch load file, and calls the xa_open function in the database shared library. Similarly, during queue manager startup, and on other subsequent occasions, some queue manager processes load the switch load file and call xa_open.

You can reduce the number of xa_* calls by using *dynamic registration*. For a complete description of this optimization technique, see [“XA dynamic registration” on page 63](#).

Configuring your system for database coordination

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

There are several tasks that you must perform before a database manager can participate in global units of works coordinated by the queue manager. These are described here as follows:

- [“Installing and configuring the database product” on page 45](#)
- [“Creating switch load files” on page 46](#)
- [“Adding configuration information to the queue manager” on page 47](#)
- [“Writing and modifying your applications” on page 48](#)
- [“Testing the system” on page 49](#)

Installing and configuring the database product

To install and configure your database product, see the product's own documentation. This topics in this section describe general configuration issues and how they relate to interoperation between WebSphere MQ and the database.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Database connections

An application that establishes a standard connection to the queue manager is associated with a thread in a separate local queue manager agent process. (A connection that is not a *fastpath* connection is a *standard* connection in this context. For more information, see [“Connecting to a queue manager using the MQCONN call” on page 200](#).)

When the application issues MQBEGIN, both it and the agent process call the xa_open function in the database client library. In response to this, the database client library code *connects* to the database that is to be involved in the unit of work *from both the application and queue manager processes*. These database connections are maintained as long as the application remains connected to the queue manager.

This is an important consideration if the database supports only a limited number of users or connections, because two connections are being made to the database to support the one application program.

Client/server configuration

The database client library that is loaded into the WebSphere MQ queue manager and application processes **must** be able to send to and receive from its server. Ensure that:

- The database's client/server configuration files have the correct details
- The relevant environment variables are set in the environment of the queue manager **and** the application processes

Creating switch load files

WebSphere MQ comes with a sample makefile, used to build switch load files for the supported database managers.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

The sample makefile, together with all the associated C source files required to build the switch load files, is installed in the following directories:

- For WebSphere MQ for Windows, in the `MQ_INSTALLATION_PATH\tools\c\samples\xatm\` directory
- For WebSphere MQ for UNIX and Linux systems, in the `MQ_INSTALLATION_PATH/samp/xatm/` directory

The sample source modules used to build the switch load files are:

- For DB2, `db2swit.c`
- For Oracle, `oraswit.c`
- For Informix, `infswit.c`
- For Sybase, `sybswit.c`

When you generate switch load files, install 32-bit switch load files in `/var/mqm/exits` and install 64-bit switch load files in `/var/mqm/exits64`.

If you have 32-bit queue managers then the sample make file, `xaswit.mak`, installs a 32-bit switch load file in `/var/mqm/exits`.

If you have 64-bit queue managers then the sample make file, `xaswit.mak`, installs a 32-bit switch load file in `/var/mqm/exits`, and a 64-bit switch load file in `/var/mqm/exits64`.

File security

It is possible that your operating system might fail the loading of the switch load file by WebSphere MQ, for reasons outside the control of WebSphere MQ. If this occurs, error messages are written to the WebSphere MQ error logs, and potentially the `MQBEGIN` call can fail. To help to ensure that your operating system does not fail the loading of the switch load file, you must fulfil the following requirements:

1. The switch load file must be available in the location that is given in the `qm.ini` file.
2. The switch load file must be accessible to all processes that need to load it, including the queue manager processes and application processes.
3. All of the libraries upon which the switch load file depends, including the libraries that are provided by the database product, must be present and accessible.

Adding configuration information to the queue manager

When you have created a switch load file for your database manager, and placed it in a safe location, you must specify that location to your queue manager.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

To specify the location, perform the following steps:

- On Windows and Linux (x86 and x86-64 platforms) systems use the WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager.
- On all other systems specify the details of the switch load file in the XAResourceManager stanza in the queue manager's qm.ini file.

Add an XAResourceManager stanza for the database that your queue manager is going to coordinate. The most common case is for there to be only one database, and therefore only one XAResourceManager stanza. For details of more complicated configurations involving multiple databases, see [“Multiple database configurations”](#) on page 58. The attributes of the XAResourceManager stanza are as follows:

Name=name

User-chosen string that identifies the resource manager. In effect, it gives a name to the XAResourceManager stanza. The name is mandatory and can be up to 31 characters in length.

The name you choose must be unique; there must be only one XAResourceManager stanza with this name in this qm.ini file. The name should also be meaningful, because the queue manager uses it to refer to this resource manager both in queue manager error log messages and in output when the `dspmqrtn` command is used. (See [“Displaying outstanding units of work with the dspmqrtn command”](#) on page 60 for more information.)

Once you have chosen a name, and have started the queue manager, do not change the Name attribute. For more details about changing configuration information, see [“Changing configuration information”](#) on page 62.

SwitchFile=name

This is the name of the XA switch load file you built earlier. This is a mandatory attribute. The code in the queue manager and WebSphere MQ application processes tries to load the switch load file on two occasions:

1. At queue manager startup
2. When you make the first call to MQBEGIN in your WebSphere MQ application process

The security and permissions attributes of your switch load file must allow these processes to perform this action.

XAOpenString=string

This is a string of data that WebSphere MQ code passes in its calls to the database manager's `xa_open` function. This is an optional attribute; if it is omitted a zero-length string is assumed.

The code in the queue manager and WebSphere MQ application processes call the `xa_open` function on two occasions:

1. At queue manager startup
2. When you make the first call to MQBEGIN in your WebSphere MQ application process

The format for this string is particular to each database product, and will be described in the documentation for that product. In general, the `xa_open` string contains authentication information (user name and password) to allow a connection to the database in both the queue manager and the application processes.

XACloseString=string

This is a string of data that WebSphere MQ code passes in its calls to the database manager's `xa_close` function. This is an optional attribute; if it is omitted a zero-length string is assumed.

The code in the queue manager and WebSphere MQ application processes call the `xa_close` function on two occasions:

1. At queue manager startup
2. When you make a call to `MQDISC` in your WebSphere MQ application process, having earlier made a call to `MQBEGIN`

The format for this string is particular to each database product, and will be described in the documentation for that product. In general, the string is empty, and it is common to omit the `XACloseString` attribute from the `XAResourceManager` stanza.

ThreadOfControl=THREAD|PROCESS

The `ThreadOfControl` value can be `THREAD` or `PROCESS`. The queue manager uses it for serialization purposes. This is an optional attribute; if it is omitted, the value `PROCESS` is assumed.

If the database client code allows threads to call the XA functions without serialization, the value for `ThreadOfControl` can be `THREAD`. The queue manager assumes that it can call the XA functions in the database client shared library from multiple threads at the same time, if necessary.

If the database client code does not allow threads to call its XA functions in this way, the value for `ThreadOfControl` must be `PROCESS`. In this case, the queue manager serializes all calls to the database client shared library so that only one call at a time is made from within a particular process. You probably also need to ensure that your application performs similar serialization if it runs with multiple threads.

Note that this issue, of the database product's ability to cope with multi-threaded processes in this way, is an issue for that product's vendor. Consult the database product's documentation for details on whether you can set the `ThreadOfControl` attribute to `THREAD` or `PROCESS`. We recommend that, if you can, you set `ThreadOfControl` to `THREAD`. If in doubt, the *safer* option is to set it to `PROCESS`, although you will lose the potential performance benefits of using `THREAD`.

Writing and modifying your applications
How to implement a global unit of work.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The sample application programs for Scenario 1 global units of work that are supplied with a WebSphere MQ installation are described in the [“Introducing units of work”](#) on page 40.

In general, a global unit of work is implemented in an application by the following method (in pseudocode):

```
MQBEGIN
MQGET
MQPUT
SQL INSERT
MQCMIT
```

The purpose of `MQBEGIN` is to denote the beginning of a global unit of work. The purpose of `MQCMIT` is to denote the end of the global unit of work, and to complete it with all participating resource managers, using the two-phase commit protocol.

In between `MQBEGIN` and `MQCMIT`, the queue manager does not make any calls to the database to update its resources. That is, the only way a database's tables are changed is by your code (for example, the `SQL INSERT` in the pseudocode).

The role of the queue manager, as far as the database is concerned, is to tell it when a global unit of work has started, when it has ended, and whether the global unit of work should be committed or rolled-back.

As far as your application is concerned, the queue manager performs two roles: a resource manager (where the resources are messages on queues) and the transaction manager for the global unit of work.

Start with the supplied sample programs, and work through the various WebSphere MQ and database API calls that are being made in those programs. The API calls concerned are fully documented in [“Sample WebSphere MQ programs”](#) on page 91, [Data types used in the MQI](#), and (in the case of the database's own API) the database's own documentation.

Testing the system

You know whether your application and system are correctly configured only by running them during testing. You can test the system's configuration (the successful communication between queue manager and database) by building and running one of the supplied sample programs.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Configuring Db2

DB2 support and configuration information.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The supported levels of Db2 are defined at the [IBM WebSphere MQ detailed system requirements](#) page.

Note: 32-bit instances of Db2 are not supported on platforms where the queue manager is 64-bit.

Do the following:

1. Check the environment variable settings.
2. Create the Db2 switch load file.
3. Add resource manager configuration information.
4. Change Db2 configuration parameters if necessary.

Read this information in conjunction with the general information provided in [“Configuring your system for database coordination”](#) on page 45.

Warning: If you run `db2profile` on UNIX and Linux platforms, the environment variable `LIBPATH` and `LD_LIBRARY_PATH` are set. It is advisable to unset these environment variables, see appropriate *Quick Beginnings* Guide.

Checking the Db2 environment variable settings

Ensure that your Db2 environment variables are set for queue manager processes **as well as in** your application processes. In particular, you must always set the `DB2INSTANCE` environment variable **before** you start the queue manager. The `DB2INSTANCE` environment variable identifies the Db2 instance containing the Db2 databases that are being updated. For example:

- On UNIX and Linux systems, use:

```
export DB2INSTANCE=db2inst1
```

- On Windows systems, use:

```
set DB2INSTANCE=DB2
```

On Windows with a Db2 database, you must add the user `MUSR_MQADMIN` to the `DB2USERS` group, to enable the queue manager to start.

Creating the Db2 switch load file

The easiest way to create the Db2 switch load file is to use the sample file `xaswit.mak`, which WebSphere MQ provides to build the switch load files for a variety of database products.

On Windows systems, you can find xaswit.mak in the directory `MQ_INSTALLATION_PATH\tools\c\samples\xatm`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed. To create the Db2 switch load file with Microsoft Visual C++, use:

```
nmake /f xaswit.mak db2swit.dll
```

The generated switch file is placed in `c:\Program Files\IBM\WebSphere MQ\exits`.

You can find xaswit.mak in the directory `MQ_INSTALLATION_PATH/samp/xatm`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Edit xaswit.mak to *uncomment* the lines appropriate to the version of Db2 you are using. Then execute the makefile using the command:

```
make -f xaswit.mak db2swit
```

The generated 32-bit switch load file is placed in `/var/mqm/exits`.

The generated 64-bit switch load file is placed in `/var/mqm/exits64`.

Adding resource manager configuration information for Db2

You must modify the configuration information for the queue manager to declare Db2 as a participant in global units of work. Modifying configuration information in this way is described in more details in [“Adding configuration information to the queue manager” on page 47](#).

- On Windows and Linux (x86 and x86-64 platforms) systems, use the WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager.
- On all other systems specify the details of the switch load file in the XAResourceManager stanza in the queue manager's `qm.ini` file.

Figure 9 on page 50 is a UNIX sample, showing an XAResourceManager entry where the database to be coordinated is called mydbname, this name being specified in the `XAOpenString`:

```
XAResourceManager:
  Name=mydb2
  SwitchFile=db2swit
  XAOpenString=mydbname,myuser,mypasswd,toc=t
  ThreadOfControl=THREAD
```

Figure 9. Sample XAResourceManager entry for Db2 on UNIX platforms

Note:

1. `ThreadOfControl=THREAD` cannot be used with Db2 versions earlier than version 8. Set `ThreadOfControl` and the `XAOpenString` parameter `toc` to one of the following combinations:
 - `ThreadOfControl=THREAD` and `toc=t`
 - `ThreadOfControl=PROCESS` and `toc=p`

If you are using the `jdbcdb2` XA switch load file to enable JDBC/JTA coordination, you must use `ThreadOfControl=PROCESS` and `toc=p`.

Changing Db2 configuration parameters

For each Db2 database that the queue manager is coordinating, you must set database privileges, change the `tp_mon_name` parameter, and reset the `maxappls` parameter. To do this, perform the following steps:

Set database privileges

The queue manager processes run with effective user and group mqm on UNIX and Linux systems. On Windows systems, they run as the user that started the queue manager. This can be one of:

1. The user who issued the `strmqm` command, or
2. The user under which the IBM MQSeries® Service COM server runs

By default, this user is called `MUSR_MQADMIN`.

If you have not specified a user name and password on the `xa_open` string, **the user under which the queue manager is running** is used by Db2 to authenticate the `xa_open` call. If this user (for example, user `mqm` on UNIX and Linux systems) does not have minimal privileges in the database, the database refuses to authenticate the `xa_open` call.

The same considerations apply to your application process. If you have not specified a user name and password on the `xa_open` string, the user under which your application is running is used by Db2 to authenticate the `xa_open` call that is made during the first `MQBEGIN`. Again, this user must have minimal privileges in the database for this to work.

For example, give the `mqm` user connect authority in the `mydbname` database by issuing the following Db2 commands:

```
db2 connect to mydbname
db2 grant connect on database to user mqm
```

See [“Security considerations” on page 59](#) for more information about security.

Windows Change the TP_MON_NAME parameter

For Db2 for Windows systems only, change the `TP_MON_NAME` configuration parameter to name the DLL that Db2 uses to call the queue manager for dynamic registration.

Use the command `db2 update dbm cfg using TP_MON_NAME mqmax` to name `MQMAX.DLL` as the library that Db2 uses to call the queue manager. This must be present in a directory within `PATH`.

Reset the maxappls parameter

You might need to review your setting for the `maxappls` parameter, which limits the maximum number of applications that can be connected to a database. Refer to [“Installing and configuring the database product” on page 45](#).

Configuring Oracle

Oracle support and configuration information.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Complete the following steps:

1. Check environment variable settings.
2. Create the Oracle switch load file.
3. Add resource manager configuration information.
4. Change the Oracle configuration parameters, if necessary.

A current list of levels of Oracle supported by IBM WebSphere MQ is provided at the [IBM WebSphere MQ detailed system requirements page](#).

Checking the Oracle environment variable settings

Ensure that your Oracle environment variables are set for queue manager processes as well as in your application processes. In particular, always set the following environment variables before starting the queue manager:

ORACLE_HOME

The Oracle home directory. For example, on UNIX and Linux systems, use:

```
export ORACLE_HOME=/opt/oracle/product/8.1.6
```

On Windows systems, use:

```
set ORACLE_HOME=c:\oracle\ora81
```

ORACLE_SID

The Oracle SID being used. If you are using Net8 for client/server connectivity, you might not have to set this environment variable. Consult your Oracle documentation.

The subsequent example is an example of setting this environment variable, on UNIX and Linux systems:

```
export ORACLE_SID=sid1
```

The equivalent on Windows systems is:

```
set ORACLE_SID=sid1
```

Note: The PATH environment variable must be set to include the binaries directory (for example ORACLE_INSTALL_DIR/VERSION/32BIT_NAME/bin or ORACLE_INSTALL_DIR/VERSION/64BIT_NAME/bin), otherwise you might see a message stating that oraclient libraries are missing from the machine.

If you run queue managers on Windows 64 bit systems, then both 64 bit and 32 bit Oracle clients must be installed. You must install both clients because the queue manager runs as 32 bit processes that use a 32 bit switch load file, which must in turn start a 32 bit Oracle client dll.

The switch load file, loaded by 64 bit queue managers, must access the Oracle 64 bit client libraries. 32 bit queue managers must access the 32 bit Oracle client when IBM WebSphere MQ is running on a Windows 64 bit system.

Creating the Oracle switch load file

To create the Oracle switch load file, use the sample file `xaswit.mak`, which IBM WebSphere MQ provides to build the switch load files for various database products. On Windows systems, you can find `xaswit.mak` in the directory `C:\Program Files\IBM\WebSphere MQ\tools\c\samples\xatm`. To create the Oracle switch load file with Microsoft Visual C++, use: `nmake /f xaswit.mak oraswit.dll`

The generated switch file is placed in `MQ_INSTALLATION_PATH\exits.MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

You can find `xaswit.mak` in the directory `MQ_INSTALLATION_PATH\samp\xatm`.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

Edit `xaswit.mak` to uncomment the lines appropriate to the version of Oracle you are using. Then execute the makefile using the command:

```
make -f xaswit.mak oraswit
```

The generated 32 bit switch load file is placed in `/var/mqm/exits`.

The generated 64 bit switch load file is placed in `/var/mqm/exits64`.

Adding resource manager configuration information for Oracle

You must modify the configuration information for the queue manager to declare Oracle as a participant in global units of work. Modifying the configuration information for the queue manager in this way is described in more detail in [“Adding configuration information to the queue manager”](#) on page 47.

- On Windows and Linux (x86 and x86-64 platforms) systems, use the IBM WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager.
- On all other systems specify the details of the switch load file in the XAResourceManager stanza in the `qm.ini` file of the queue manager.

Figure 10 on page 53 is a UNIX and Linux systems sample showing an XAResourceManager entry. You must add a LogDir to the XA open string so that all error and tracing information is logged to the same place.

```
XAResourceManager:
  Name=myoracle
  SwitchFile=oraswit
  XAOpenString=Oracle_XA+Acc=P/myuser/mypasswd+SesTm=35+LogDir=/tmp+threads=true
  ThreadOfControl=THREAD
```

Figure 10. Sample XAResourceManager entry for Oracle on UNIX and Linux platforms

Note:

1. In Figure 10 on page 53, the `xa_open` string has been used with four parameters. Additional parameters can be included as described in Oracle's documentation.
2. When using the IBM WebSphere MQ parameter `ThreadOfControl=THREAD` you must use the Oracle parameter `+threads=true` in the XAResourceManager stanza.

See the *Oracle8 Server Application Developer's Guide* for more information about the `xa_open` string.

Changing Oracle configuration parameters

For each Oracle database that the queue manager is coordinating, you must review your maximum sessions and set database privileges. To do so, complete these steps:

Review your maximum sessions

You might have to review your `LICENSE_MAX_SESSIONS` and `PROCESSES` settings to take into account the additional connections required by processes belonging to the queue manager. See [“Installing and configuring the database product”](#) on page 45 for more details.

Set database privileges

The Oracle user name specified in the `xa_open` string must have privileges to access the `DBA_PENDING_TRANSACTIONS` view, as described in the Oracle documentation.

The necessary privilege can be given using the following example command:

```
grant select on DBA_PENDING_TRANSACTIONS to myuser;
```

Configuring Informix

Informix support and configuration information.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Complete the following steps:

1. Ensure that you have installed the appropriate Informix client SDK:

- 32 bit queue managers and applications require a 32 bit Informix client SDK.
 - 64 bit queue managers and applications require a 64 bit Informix client SDK.
2. Ensure that Informix databases are created correctly.
 3. Check environment variable settings.
 4. Build the Informix switch load file.
 5. Add resource manager configuration information.

A current list of levels of Informix supported by WebSphere MQ is provided at the [IBM WebSphere MQ detailed system requirements](#) page.

Ensuring that Informix databases are created correctly

Every Informix database that is to be coordinated by a WebSphere MQ queue manager must be created specifying the `log` parameter. For example:

```
create database mydbname with log;
```

WebSphere MQ queue managers are unable to coordinate Informix databases that do not have the `log` parameter specified on creation. If a queue manager attempts to coordinate an Informix database that does not have the `log` parameter specified on creation, the `xa_open` call to Informix fails, and a number of FFST errors are generated.

Checking the Informix environment variable settings

Ensure that your Informix environment variables are set for queue manager processes **as well as in** your application processes. In particular, always set the following environment variables **before** starting the queue manager:

INFORMIXDIR

The directory of the Informix product installation.

- For 32 bit UNIX and Linux applications, use the following command:

```
export INFORMIXDIR=/opt/informix/32-bit
```

- For 64 bit UNIX and Linux applications, use the following command:

```
export INFORMIXDIR=/opt/informix/64-bit
```

- For Windows applications, use the following command:

```
set INFORMIXDIR=c:\informix
```

For systems that have 64 bit queue managers that must support both 32 bit and 64 bit applications, you need both the Informix 32 bit and 64 bit client SDKs installed. The sample makefile `xaswit.mak`, used for creating a switch load file also sets both product installation directories.

INFORMIXSERVER

The name of the Informix server. For example, on UNIX and Linux systems, use:

```
export INFORMIXSERVER=hostname_1
```

On Windows systems, use:

```
set INFORMIXSERVER=hostname_1
```

ONCONFIG

The name of the Informix server configuration file. For example, on UNIX and Linux systems, use:

```
export ONCONFIG=onconfig.hostname_1
```

On Windows systems, use:

```
set ONCONFIG=onconfig.hostname_1
```

Creating the Informix switch load file

To create the Informix switch load file, use the sample file `xaswit.mak`, which WebSphere MQ provides to build the switch load files for various database products. On Windows systems, you can find `xaswit.mak` in the directory `MQ_INSTALLATION_PATH\tools\c\samples\xatm`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed. To create the Informix switch load file with Microsoft Visual C++, use:

```
nmake /f xaswit.mak infswit.dll
```

The generated switch file is placed in `c:\Program Files\IBM\WebSphere MQ\exits`.

You can find `xaswit.mak` in the directory `MQ_INSTALLATION_PATH\samp\xatm`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Edit `xaswit.mak` to *uncomment* the lines appropriate to the version of Informix you are using. Then execute the makefile using the command:

```
make -f xaswit.mak infswit
```

The generated 32 bit switch load file is placed in `/var/mqm/exits`.

The generated 64 bit switch load file is placed in `/var/mqm/exits64`.

Adding resource manager configuration information for Informix

You must modify the configuration information for the queue manager to declare Informix as a participant in global units of work. Modifying the configuration information for the queue manager in this way is described in more detail in [“Adding configuration information to the queue manager” on page 47](#).

- On Windows and Linux (x86 and x86-64 platforms) systems, use the WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager.
- On all other systems specify the details of the switch load file in the XAResourceManager stanza in the `qm.ini` file of the queue manager.

Figure 11 on page 55 is a UNIX sample, showing a `qm.ini` XAResourceManager entry where the database to be coordinated is called `mydbname`, this name being specified in the `XAOpenString`:

```
XAResourceManager:
  Name=myinformix
  SwitchFile=infswit
  XAOpenString=DB=mydbname@myinformixserver\;USER=myuser\;PASSWD=mypasswd
  ThreadOfControl=THREAD
```

Figure 11. Sample XAResourceManager entry for Informix on UNIX platforms

Note: By default the sample `xaswit.mak` on UNIX platforms creates a switch load file that uses threaded Informix libraries. You must ensure that `ThreadOfControl` is set to `THREAD` when using these Informix libraries. In [Figure 11 on page 55](#), the `qm.ini` file XAResourceManager stanza attribute `ThreadOfControl`

is set to THREAD. When THREAD is specified, applications must be built using the threaded Informix libraries and the WebSphere MQ threaded API libraries.

The XAOpenString attribute must contain the database name, followed by the @ symbol, and then followed by the Informix server name.

To use the nonthreaded Informix libraries, you must ensure that the qm.ini file XAResourceManager stanza attribute ThreadOfControl is set to PROCESS. You must also make the following changes to the sample xaswit.mak:

1. Uncomment the generation of a nonthreaded switch load file.
2. Comment out the generation of the threaded switch load file.

Sybase configuration

Sybase support and configuration information.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Complete the following steps:

1. Ensure you have installed the Sybase XA libraries, for example by installing the XA DTM option.
2. Check environment variable settings.
3. Enable Sybase XA support.
4. Create the Sybase switch load file.
5. Add resource manager configuration information.

A current list of levels of Sybase supported by WebSphere MQ is provided at the [IBM WebSphere MQ detailed system requirements](#) page.

Checking the Sybase environment variable settings

Ensure that your Sybase environment variables are set for queue manager processes **as well as in** your application processes. In particular, always set the following environment variables **before** starting the queue manager:

SYBASE

The location of the Sybase product installation. For example, on UNIX and Linux systems, use:

```
export SYBASE=/sybase
```

On Windows systems, use:

```
set SYBASE=c:\sybase
```

SYBASE_OCS

The directory under SYBASE where you have installed the Sybase client files. For example, on UNIX and Linux systems, use:

```
export SYBASE_OCS=OCS-12_0
```

On Windows systems, use:

```
set SYBASE_OCS=OCS-12_0
```


Enabling Sybase XA support

Within the Sybase XA configuration file `$SYBASE/$SYBASE_OCS/xa_config`, define a Logical Resource Manager (LRM) for each connection to the Sybase server that is being updated. An example of the contents of `$SYBASE/$SYBASE_OCS/xa_config` is shown in Figure 12 on page 57.

```
# The first line must always be a comment

[xa]

    LRM=lrname
    server=servername
```

Figure 12. Example contents of `$SYBASE/$SYBASE_OCS/xa_config`

Creating the Sybase switch load file

To create the Sybase switch load file, use the sample files supplied with WebSphere MQ. On Windows systems, you can find `xaswit.mak` in the directory `C:\Program Files\IBM\WebSphere MQ\tools\c\samples\xatm`. To create the Sybase switch load file with Microsoft Visual C++, use:

```
nmake /f xaswit.mak sybswit.dll
```

The generated switch file is placed in `c:\Program Files\IBM\WebSphere MQ\exits`.

You can find `xaswit.mak` in the directory `MQ_INSTALLATION_PATH/samp/xatm`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Edit `xaswit.mak` to *uncomment* the lines appropriate to the version of Sybase you are using. Then execute the makefile using the command:

```
make -f xaswit.mak sybswit
```

The generated 32-bit switch load file is placed in `/var/mqm/exits`.

The generated 64-bit switch load file is placed in `/var/mqm/exits64`.

Adding resource manager configuration information for Sybase

You must modify the configuration information for the queue manager to declare Sybase as a participant in global units of work. Modifying the configuration information is described in more detail in [“Adding configuration information to the queue manager”](#) on page 47.

- On Windows and Linux (x86 and x86-64 platforms) systems, use the WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager.
- On all other systems specify the details of the switch load file in the `XAResourceManager` stanza in the queue manager's `qm.ini` file.

Figure 13 on page 58 shows a UNIX and Linux sample, which uses the database associated with the `lrname` LRM definition in the Sybase XA configuration file, `$SYBASE/$SYBASE_OCS/xa_config`. Include a log file name if you want XA function calls to be logged:

```
XAResourceManager:  
  Name=mysybase  
  SwitchFile=sybswit  
  XAOpenString=-User -Ppassword -Nlrmname -L/tmp/sybase.log -Txa  
  ThreadOfControl=THREAD
```

Figure 13. Sample XAResourceManager entry for Sybase on UNIX and Linux platforms

Using multi-threaded programs with Sybase

If you are using multi-threaded programs with WebSphere MQ global units of work incorporating updates to Sybase, you **must** use the value THREAD for the ThreadOfControl parameter. Also ensure that you link your program (and the switch load file) with the threadsafe Sybase libraries (the _r versions). Using the value THREAD for the ThreadOfControl parameter is shown in [Figure 13 on page 58](#).

Multiple database configurations

If you want to configure the queue manager so that updates to multiple databases can be included within global units of work, add an XAResourceManager stanza for each database.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

If the databases are all managed by the same database manager, each stanza defines a separate database. Each stanza specifies the same *SwitchFile*, but the contents of the *XAOpenString* are different because it specifies the name of the database being updated. For example, the stanzas shown in [Figure 14 on page 58](#) configure the queue manager with the Db2 databases *MQBankDB* and *MQFeeDB* on UNIX and Linux systems.

Important: You cannot have multiple stanzas pointing to the same database. This configuration does not work under any circumstances, and if you try this configuration it fails.

You will receive errors of the form when the MQ code makes its second *xa_open* call in any process in this environment, the database software fails the second *xa_open* with a -5 error, XAER_INVAL.

```
XAResourceManager:  
  Name=DB2 MQBankDB  
  SwitchFile=db2swit  
  XAOpenString=MQBankDB  
  
XAResourceManager:  
  Name=DB2 MQFeeDB  
  SwitchFile=db2swit  
  XAOpenString=MQFeeDB
```

Figure 14. Sample XAResourceManager entries for multiple Db2 databases

If the databases to be updated are managed by different database managers, add an XAResourceManager stanza for each. In this case, each stanza specifies a different *SwitchFile*. For example, if *MQFeeDB* is managed by Oracle instead of Db2, use the following stanzas on UNIX and Linux systems:

```
XAResourceManager:
  Name=DB2 MQBankDB
  SwitchFile=db2swit
  XAOpenString=MQBankDB

XAResourceManager:
  Name=Oracle MQFeeDB
  SwitchFile=oraswit
  XAOpenString=Oracle_XA+Acc=P/myuser/mypassword+SesTm=35+LogDir=/tmp/ora.log+DB=MQFeeDB
```

Figure 15. Sample XAResourceManager entries for a DB2 and Oracle database

In principle, there is no limit to the number of database instances that can be configured with a single queue manager.

Note: For information on support for including Informix databases in multiple database updates within global units of work, check the product readme file.

Security considerations

Considerations for running your database under the XA model.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The following information is provided for guidance only. In all cases, refer to the documentation provided with the database manager to determine the security implications of running your database under the XA model.

An application process denotes the start of a global unit of work using the MQBEGIN verb. The first MQBEGIN call that an application issues connects to all participating databases by calling their client library code at the xa_open entry point. All the database managers provide a mechanism for supplying a user ID and password in their XAOpenString. This is the only time that authentication information flows.

Note that, on UNIX and Linux platforms, fastpath applications must run with an effective user ID of mqm while making MQI calls.

Considerations when contact is lost with the XA resource manager

The queue manager tolerates database managers not being available. This means you can start and stop the queue manager independently from the database server. When contact is restored, the queue manager and database resynchronize. You can also use the rsvmqtrn command to manually resolve in-doubt units of work.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

In normal operations, only a minimal amount of administration is necessary after you have completed the configuration steps. The administration job is made easier because the queue manager tolerates database managers not being available. In particular this means that:

- The queue manager can start at any time without first starting each of the database managers.
- The queue manager does not need to stop and restart if one of the database managers becomes unavailable.

This allows you to start and stop the queue manager independently from the database server.

Whenever contact is lost between the queue manager and a database, they need to resynchronize when both become available again. Resynchronization is the process by which any in-doubt units of work involving that database are completed. In general, this occurs automatically without the need for user intervention. The queue manager asks the database for a list of units of work that are in doubt. It then instructs the database to either commit or roll back each of these in-doubt units of work.

When a queue manager starts, it resynchronizes with each database. When an individual database becomes unavailable, only that database needs to be resynchronized the next time that the queue manager notices it is available again.

The queue manager regains contact with a previously unavailable database automatically as new global units of work are started with MQBEGIN. It does this by calling the xa_open function in the database client library. If this xa_open call fails, MQBEGIN returns with a completion code of MQCC_WARNING and a reason code of MQRC_PARTICIPANT_NOT_AVAILABLE. You can retry the MQBEGIN call later.

Do not continue to attempt a global unit of work that involves updates to a database that has indicated failure during MQBEGIN. There will not be a connection to that database through which updates can be made. Your only options are to end the program, or to retry MQBEGIN periodically in the hope that the database might become available again.

Alternatively, you can use the `rsvmqtrn` command to resolve explicitly all in-doubt units of work.

In-doubt units of work

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

A database might be left with in-doubt units of work if contact with the queue manager is lost after the database manager has been instructed to prepare. Until the database server receives the outcome from the queue manager (commit or roll back), it needs to retain the database locks associated with the updates.

Because these locks prevent other applications from updating or reading database records, resynchronization needs to take place as soon as possible.

If, for some reason, you cannot wait for the queue manager to resynchronize with the database automatically, you can use facilities provided by the database manager to commit or roll back the database updates manually. In the *X/Open Distributed Transaction Processing: The XA Specification*, this is called making a *heuristic* decision. Use it only as a last resort because of the possibility of compromising data integrity; you might, for example, mistakenly roll back the database updates when all other participants have committed their updates.

It is far better to restart the queue manager, or use the `rsvmqtrn` command when the database has been restarted, to initiate automatic resynchronization.

Displaying outstanding units of work with the dspmqtrn command

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

While a database manager is unavailable, you can use the **dspmqtrn** command to check the state of outstanding global units of work involving that database.

The **dspmqtrn** command displays only those units of work in which one or more participants are in doubt. The participants are awaiting the decision from the queue manager to commit or roll back the prepared updates.

For each of these global units of work, the state of each participant is displayed in the output from **dspmqtrn**. If the unit of work did not update the resources of a particular resource manager, it is not displayed.

With respect to an in-doubt unit of work, a resource manager is said to have done one of the following things:

Prepared

The resource manager is prepared to commit its updates.

Committed

The resource manager has committed its updates.

Rolled-back

The resource manager has rolled back its updates.

Participated

The resource manager is a participant, but has not prepared, committed, or rolled back its updates.

When the queue manager is restarted, it asks each database having an XAResourceManager stanza for a list of its in-doubt global units of work. If the database has not been restarted, or is otherwise unavailable, the queue manager cannot yet deliver to the database the final outcomes for those units of work. The outcome of the in-doubt units of work is delivered to the database at the first opportunity when the database is again available.

In this case, the database manager is reported as being in *prepared* state until resynchronization has occurred.

Whenever the dspmqtrn command displays an in-doubt unit of work, it first lists all the possible resource managers that might be participating. These are allocated a unique identifier, *RMId*, which is used instead of the *Name* of the resource managers when reporting their state with respect to an in-doubt unit of work.

Sample dspmqtrn output shows the result of issuing the following command:

```
dspmqtrn -m MY_QMGR
```

```
AMQ7107: Resource manager 0 is MQSeries.  
AMQ7107: Resource manager 1 is DB2 MQBankDB.  
AMQ7107: Resource manager 2 is DB2 MQFeedB.  
  
AMQ7056: Transaction number 0,1.  
      XID: formatID 5067085, gtrid_length 12, bqual_length 4  
          gtrid [3291A5060000201374657374]  
          bqual [00000001]  
AMQ7105: Resource manager 0 has committed.  
AMQ7104: Resource manager 1 has prepared.  
AMQ7104: Resource manager 2 has prepared.
```

where *Transaction number* is the ID of the transaction which can be used with the *rsvmqtrn* command. See AMQ7000-7999: WebSphere MQ product for further information about AMQ7056 message. The *XID* variables are part of the *X/Open XA Specification*; for the most up-to-date information about this specification see: <https://publications.opengroup.org/c193>.

Figure 16. Sample dspmqtrn output

The output in Sample dspmqtrn output shows that there are three resource managers associated with the queue manager. The first is resource manager 0, which is the queue manager itself. The other two resource manager instances are the MQBankDB and MQFeedB Db2 databases.

The example shows only a single in-doubt unit of work. A message is issued for all three resource managers, which means that updates were made to the queue manager and both Db2 databases within the unit of work.

The updates made to the queue manager, resource manager 0, have been *committed*. The updates to the Db2 databases are in *prepared* state, which means that Db2 must have become unavailable before it was called to commit the updates to the MQBankDB and MQFeedB databases.

The in-doubt unit of work has an external identifier called an XID (*transaction id*). This is a piece of data given to Db2 by the queue manager to identify its portion of the global unit of work.

Resolving outstanding units of work with the rsvmqtrn command

Outstanding units of work complete when the queue manager and DB2 resynchronize.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The output shown in [Figure 16 on page 61](#) shows a single in-doubt unit of work in which the commit decision has yet to be delivered to both DB2 databases.

To complete this unit of work, the queue manager and DB2 need to resynchronize when DB2 next becomes available. The queue manager uses the start of new units of work as an opportunity to regain contact with DB2. Alternatively, you can instruct the queue manager to resynchronize explicitly using the **rsvmqtrn** command.

Do this soon after DB2 has been restarted, so that any database locks associated with the in-doubt unit of work are released as quickly as possible. Use the -a option, which tells the queue manager to resolve all in-doubt units of work. In the following example, DB2 has restarted, so the queue manager can resolve the in-doubt unit of work:

```
> rsvmqtrn -m MY_QMGR -a
Any in-doubt transactions have been resolved.
```

Mixed outcomes and errors

Although the queue manager uses a two-phase commit protocol, this does not completely remove the possibility of some units of work completing with mixed outcomes. This is where some participants commit their updates and some back out their updates.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Units of work that complete with a mixed outcome have serious implications because shared resources that should have been updated as a single unit of work are no longer in a consistent state.

Mixed outcomes are mainly caused when heuristic decisions are made about units of work instead of allowing the queue manager to resolve in-doubt units of work itself. Such decisions are outside the queue manager's control.

Whenever the queue manager detects a mixed outcome, it produces FFST information and documents the failure in its error logs, with one of two messages:

- If a database manager rolls back instead of committing:

```
AMQ7606 A transaction has been committed but one or more resource
managers have rolled back.
```

- If a database manager commits instead of rolling back:

```
AMQ7607 A transaction has been rolled back but one or more resource
managers have committed.
```

Further messages identify the databases that are heuristically damaged. It is then your responsibility to locally restore consistency to the affected databases. This is a complicated procedure in which you need first to isolate the update that has been wrongly committed or rolled back, then to undo or redo the database change manually.

Changing configuration information

After the queue manager has successfully started to coordinate global units of work, do not change any of the resource manager configuration information.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

If you need to change the configuration information you can do so at any time, but the changes do not take effect until after the queue manager has been restarted.

If you remove the resource manager configuration information for a database, you are effectively removing the ability for the queue manager to contact that database manager.

Never change the *Name* attribute in any of your resource manager configuration information. This attribute uniquely identifies that database manager instance to the queue manager. If you change this unique identifier, the queue manager assumes that the database has been removed and a completely new instance has been added. The queue manager still associates outstanding units of work with the old *Name*, possibly leaving the database in an in-doubt state.

Removing database manager instances

If you need to remove a database from your configuration permanently, ensure that the database is not in doubt before you restart the queue manager.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Database products provide commands for listing in-doubt transactions. If there are any in-doubt transactions, first allow the queue manager to resynchronize with the database. Do this by starting the queue manager. You can verify that resynchronization has taken place by using the **rsvmqtrn** command or the database's own command for viewing in-doubt units of work. Once you are satisfied that resynchronization has taken place, end the queue manager and remove the database's configuration information.

If you fail to observe this procedure the queue manager still remembers all in-doubt units of work involving that database. A warning message, AMQ7623, is issued every time the queue manager is restarted. If you are never going to configure this database with the queue manager again, use the -r option of the **rsvmqtrn** command to instruct the queue manager to forget about the database's participation in its in-doubt transactions. The queue manager forgets about such transactions only when in-doubt transactions have been completed with all participants.

There are times when you might need to remove some resource manager configuration information temporarily. On UNIX and Linux systems this is best achieved by commenting out the stanza so that it can be easily reinstated at a later time. You might decide to do this if there are errors every time the queue manager contacts a particular database or database manager. Temporarily removing the resource manager configuration information concerned allows the queue manager to start global units of work involving all the other participants. Here is an example of a commented-out XAResourceManager stanza follows:

```
# This database has been temporarily removed
#XAResourceManager:
#  Name=mydb2
#  SwitchFile=db2swit
#  XAOpenString=mydbname,myuser,mypassword,toc=t
#  ThreadOfControl=THREAD
```

Figure 17. Commented-out XAResourceManager stanza on UNIX and Linux systems

On Windows systems, use the WebSphere MQ Explorer to delete the information about the database manager instance. Take great care to type in the correct name in the *Name* field when reinstating it. If you mistype the name, you may face in-doubt problems, as described in ["Changing configuration information"](#) on page 62.

XA dynamic registration

The XA specification provides a way of reducing the number of xa_* calls that a transaction manager makes to a resource manager. This optimization is known as *dynamic registration*.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Dynamic registration is supported by DB2. Other databases might support it; consult the documentation for your database product for details.

Why is the dynamic registration optimization useful? In your application, some global units of work might contain updates to database tables; others might not contain such updates. When no persistent update has been made to a database's tables, there is no need to include that database in the commit protocol that occurs during MQCMIT.

Whether or not your database supports dynamic registration, your application calls `xa_open` during the first MQBEGIN call on a WebSphere MQ connection. It also calls `xa_close` on the subsequent MQDISC call. The pattern of subsequent XA calls depends on whether the database supports dynamic registration:

If your database does not support dynamic registration...

Every global unit of work involves several XA function calls made by WebSphere MQ code into the database client library, regardless of whether you made a persistent update to the tables of that database within your unit of work. These include:

- `xa_start` and `xa_end` from the application process. These are used to declare the beginning and end of a global unit of work.
- `xa_prepare`, `xa_commit`, and `xa_rollback` from the queue manager agent process, `amqzlaa0`. These are used to deliver the outcome of the global unit of work: the commit or rollback decision.

In addition, the queue manager agent process also calls `xa_open` during the first MQBEGIN.

If your database supports dynamic registration...

The WebSphere MQ code makes only those XA function calls that are necessary. For a global unit of work that has **not** involved persistent updates to database resources, there are **no** XA calls to the database. For a global unit of work that **has** involved such persistent updates, the calls are to:

- `xa_end` from the application process to declare the end of the global unit of work.
- `xa_prepare`, `xa_commit`, and `xa_rollback` from the queue manager agent process, `amqzlaa0`. These are used to deliver the outcome of the global unit of work: the commit or rollback decision.

For dynamic registration to work, it is vital that the database has a way of telling WebSphere MQ when it has performed a persistent update that it wants to be included in the current global unit of work. WebSphere MQ provides the `ax_reg` function for this purpose.

The database's client code that runs in your application process finds the `ax_reg` function and calls it, to *dynamically register* the fact it has done persistent work within the current global unit of work. In response to this `ax_reg` call, WebSphere MQ records that the database has participated. If this is the first `ax_reg` call on this WebSphere MQ connection, the queue manager agent process calls `xa_open`.

The database client code make this `ax_reg` call when it is running in your process, for example, during an SQL UPDATE call or whatever call in the database's client API is responsible

Error conditions

In XA dynamic registration there is a possibility of a confusing failure in the queue manager.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

A common example is if you forget to set your database environment variables properly before starting your queue manager, the queue manager's calls to `xa_open` fail. No global units of work can be used.

To avoid this, ensure that you have set the relevant environment variables before starting the queue manager. Review your database product's documentation, and the advice given in [“Configuring Db2” on page 49](#), [“Configuring Oracle” on page 51](#), and [“Sybase configuration” on page 56](#).

With all database products, the queue manager calls `xa_open` once at queue manager startup, as part of the recovery session (as explained in [“Considerations when contact is lost with the XA resource manager” on page 59](#)). This `xa_open` call fails if you set your database environment variables incorrectly, but it does not cause the queue manager to fail to start. This is because the same `xa_open` error code is used by the database client library to indicate that the database server is unavailable. WebSphere MQ does not treat

this as a serious error, as the queue manager must be able to start to continue processing data outside global units of work involving that database.

Subsequent calls to `xa_open` are made from the queue manager during the first `MQBEGIN` on a WebSphere MQ connection (if dynamic registration is not being used) or during a call by the database client code to the WebSphere MQ-provided `ax_reg` function (if dynamic registration is being used).

The **timing** of any error conditions (or, occasionally, FFST reports) depends on whether you are using dynamic registration:

- If you are using dynamic registration, your `MQBEGIN` call could succeed, but your SQL UPDATE (or similar) database call will fail.
- If you are not using dynamic registration, your `MQBEGIN` call will fail.

Ensure that your environment variables are set correctly in your application and queue manager processes.

Summarizing XA calls

Here is a list of the calls that are made to the XA functions in a database client library as a result of the various MQI calls that control global units of work. This is not a complete description of the protocol described in the XA specification; it is provided as a brief overview.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Note that `xa_start` and `xa_end` calls are always called by WebSphere MQ code in the application process, whereas `xa_prepare`, `xa_commit`, and `xa_rollback` are always called from the queue manager agent process, `amqzlaa0`.

The `xa_open` and `xa_close` calls shown in this table are all made from the application process. The queue manager agent process calls `xa_open` in the circumstances described in [“Error conditions” on page 64](#).

Table 7. Summary of XA function calls		
MQI call	XA calls made with dynamic registration	XA calls made without dynamic registration
First <code>MQBEGIN</code>	<code>xa_open</code>	<code>xa_open</code> <code>xa_start</code>
Subsequent <code>MQBEGIN</code>	No XA calls	<code>xa_start</code>
<code>MQCMIT</code> (without <code>ax_reg</code> being called during the current global unit of work)	No XA calls	<code>xa_end</code> <code>xa_prepare</code> <code>xa_commit</code> <code>xa_rollback</code>
<code>MQCMIT</code> (with <code>ax_reg</code> being called during the current global unit of work)	<code>xa_end</code> <code>xa_prepare</code> <code>xa_commit</code> <code>xa_rollback</code>	Not applicable. No calls are made to <code>ax_reg</code> in non-dynamic mode.
<code>MQBACK</code> (without <code>ax_reg</code> being called during the current global unit of work)	No XA calls	<code>xa_end</code> <code>xa_rollback</code>

Table 7. Summary of XA function calls (continued)		
MQI call	XA calls made with dynamic registration	XA calls made without dynamic registration
MQBACK (with ax_reg being called during the current global unit of work)	xa_end xa_rollback	Not applicable. No calls are made to ax_reg in non-dynamic mode.
MQDISC, where MQCMIT or MQBACK was called first. If they were not, MQCMIT processing is first done during MQDISC.	xa_close	xa_close
Notes: 1. For MQCMIT, xa_commit is called if xa_prepare is successful. Otherwise, xa_rollback is called.		

Scenario 2: Other software provides the coordination

In scenario 2, an external transaction manager coordinates global units of work, starting and committing them under control of the transaction manager's API. The MQBEGIN, MQCMIT, and MQBACK verbs are unavailable.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

This section describes this scenario, including:

- [“External sync point coordination” on page 66](#)
- [“Using CICS” on page 69](#)
- [“Using the Microsoft Transaction Server \(COM+\)” on page 73](#)

The IBM WebSphere MQ client for HP Integrity NonStop Server can use the HP NonStop Transaction Management Facility (TMF) to coordinate global units of work. For more information, see [Using HP NonStop TMF](#).

External sync point coordination

A global unit of work can also be coordinated by an external X/Open XA-compliant transaction manager. Here the WebSphere MQ queue manager participates in, but does not coordinate, the unit of work.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The flow of control in a global unit of work coordinated by an external transaction manager is as follows:

1. An application tells the external sync point coordinator (for example, TXSeries) that it wants to start a transaction.
2. The sync point coordinator tells known resource managers, such as WebSphere MQ, about the current transaction.
3. The application issues calls to resource managers associated with the current transaction. For example, the application could issue MQGET calls to WebSphere MQ.
4. The application issues a commit or backout request to the external sync point coordinator.
5. The sync point coordinator completes the transaction by issuing the appropriate calls to each resource manager, typically using two-phase commit protocols.

The supported levels of external sync point coordinators that can provide a two-phase commit process for transactions in which WebSphere MQ participates are defined at [IBM WebSphere MQ detailed system requirements](#).

The rest of this section describes how to enable external units of work.

The IBM WebSphere MQ XA switch structure

Each resource manager participating in an externally coordinated unit of work must provide an XA switch structure. This structure defines both the capabilities of the resource manager and the functions that are to be called by the sync point coordinator.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

IBM WebSphere MQ provides two versions of this structure:

- *MQRMIXASwitch* for static XA resource management
- *MQRMIXASwitchDynamic* for dynamic XA resource management

Consult your transaction manager documentation to determine whether to use the static or dynamic resource management interface. Wherever a transaction manager supports it, we recommend that you use dynamic XA resource management.

Some 64-bit transaction managers treat the *long* type in the XA specification as 64-bit, and some treat it as 32-bit. WebSphere MQ supports both models:

- If your transaction manager is 32-bit, or your transaction manager is 64-bit but treats the *long* type as 32-bit, use the switch load file listed in [Table 8 on page 67](#).
- If your transaction manager is 64-bit and treats the *long* type as 64-bit, use the switch load file listed in [Table 9 on page 68](#).

A list of known 64-bit transaction managers that treat the *long* type as 64-bit is provided in [Table 10 on page 68](#). Consult your transaction manager documentation if you are unsure which model your transaction manager uses.

<i>Table 8. XA switch load file names</i>		
Platform	Switch load file name (server)	Switch load file name (extended transactional client)
Windows	<i>mqmxa.dll</i>	<i>mqcxa.dll</i>
AIX (nonthreaded)	<i>libmqmxa.a</i>	<i>libmqcxa.a</i>
AIX (threaded)	<i>libmqmxa_r.a</i>	<i>libmqcxa_r.a</i>
HP-UX (nonthreaded)	<i>libmqmxa.so</i>	<i>libmqcxa.so</i>
HP-UX (threaded)	<i>libmqmxa_r.so</i>	<i>libmqcxa_r.so</i>
Linux (nonthreaded)	<i>libmqmxa.so</i>	<i>libmqcxa.so</i>
Linux (threaded)	<i>libmqmxa_r.so</i>	<i>libmqcxa_r.so</i>
Solaris	<i>libmqmxa.so</i>	<i>libmqcxa.so</i>

Table 9. Alternative 64-bit XA switch load file names

Platform	Switch load file name (server)	Switch load file name (extended transactional client)
AIX (nonthreaded)	<i>libmqmxa64.a</i>	<i>libmqcxa64.a</i>
AIX (threaded)	<i>libmqmxa64_r.a</i>	<i>libmqcxa64_r.a</i>
HP-UX (nonthreaded)	<i>libmqmxa64.so</i>	<i>libmqcxa64.so</i>
HP-UX (threaded)	<i>libmqmxa64_r.so</i>	<i>libmqcxa64_r.so</i>
Linux (nonthreaded)	<i>libmqmxa64.so</i>	<i>libmqcxa64.so</i>
Linux (threaded)	<i>libmqmxa64_r.so</i>	<i>libmqcxa64_r.so</i>
Solaris	<i>libmqmxa64.so</i>	<i>libmqcxa64.so</i>

Table 10. 64-bit transaction managers that require the alternative 64-bit switch load file

Transaction Manager
Tuxedo

Some external sync point coordinators (not CICS) require that each resource manager participating in a unit of work supplies its name in the name field of the XA switch structure. The WebSphere MQ resource manager name is MQSeries_XA_RMI.

The sync point coordinator defines how the WebSphere MQ XA switch structure links to it. Information about linking the WebSphere MQ XA switch structure with CICS is provided in “Using CICS” on page 69. For information about linking the WebSphere MQ XA switch structure with other XA-compliant sync point coordinators, consult the documentation supplied with those products.

The following considerations apply to using WebSphere MQ with all XA-compliant sync point coordinators:

- The `xa_info` structure passed on any `xa_open` call by the sync point coordinator includes the name of a WebSphere MQ queue manager. The name takes the same form as the queue-manager name passed to the `MQCONN` call. If the name passed on the `xa_open` call is blank, the default queue manager is used.
Alternatively, the `xa_info` structure can contain values for the `TPM` and `AXLIB` parameters. The `TPM` parameter specifies the transaction manager being used. The valid values are CICS, TUXEDO and ENCINA. The `AXLIB` parameter specifies the name of the library that contains the transaction manager's `ax_reg` and `ax_unreg` functions. For more information on these parameters, see [Configuring an extended transactional client](#). If the `xa_info` structure contains either of these parameters, the queue manager name is specified in the `QMNAME` parameter, unless the default queue manager is being used.
- Only one queue manager at a time can participate in a transaction coordinated by an instance of an external sync point coordinator. The sync point coordinator is effectively connected to the queue manager, and is subject to the rule that only one connection at a time is supported.
- All applications that include calls to an external sync point coordinator can connect only to the queue manager that is participating in the transaction managed by the external coordinator (because they are already effectively connected to that queue manager). However, such applications must issue an `MQCONN` call to obtain a connection handle, and an `MQDISC` call before they exit.
- A queue manager with resource updates coordinated by an external sync point coordinator must start before the external sync point coordinator. Similarly, the sync point coordinator must end before the queue manager.
- If your external sync point coordinator terminates abnormally, stop and restart your queue manager **before** restarting the sync point coordinator to ensure that any messaging operations uncommitted at the time of the failure are properly resolved.

Using CICS

CICS is one of the elements of TXSeries.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The versions of TXSeries that are XA-compliant (and use a two-phase commit process) are defined at: [IBM WebSphere MQ detailed system requirements](#)

WebSphere MQ also supports other transaction managers. See [IBM WebSphere MQ detailed system requirements](#) for the current lists of supported software.

Requirements of the two-phase commit process

Requirements of the two-phase commit process when you use the CICS two-phase commit process with WebSphere MQ. These requirements do not apply to z/OS.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Note the following requirements:

- WebSphere MQ and CICS must reside on the same physical machine.
- WebSphere MQ does not support CICS on a WebSphere MQ MQI client.
- You must start the queue manager, with its name specified in the XAD resource definition stanza, **before** you attempt to start CICS. Failure to do this will prevent you from starting CICS if you have added an XAD resource definition stanza for WebSphere MQ to the CICS region.
- Only one WebSphere MQ queue manager can be accessed at a time from a single CICS region.
- A CICS transaction must issue an MQCONN request before it can access WebSphere MQ resources. The MQCONN call must specify the name of the WebSphere MQ queue manager specified on the XAOpen entry of the XAD resource definition stanza for the CICS region. If this entry is blank, the MQCONN request must specify the default queue manager.
- A CICS transaction that accesses WebSphere MQ resources must issue an MQDISC call from the transaction before returning to CICS. Failure to do this might mean that the CICS application server is still connected, leaving queues open. Additionally, if you do not install a task termination exit (see [“Sample task termination exit” on page 72](#)), the CICS application server might later end abnormally, perhaps during a subsequent transaction.
- You must ensure that the CICS user ID (cics) is a member of the mqm group, so that the CICS code has the authority to call WebSphere MQ.

For transactions running in a CICS environment, the queue manager adapts its methods of authorization and determining context as follows:

- The queue manager queries the user ID under which CICS runs the transaction. This is the user ID checked by the Object Authority Manager, and is used for context information.
- In the message context, the application type is MQAT_CICS.
- The application name in the context is copied from the CICS transaction name.

General XA support

General XA is not supported on IBM i. An XA switch load module is provided to enable you to link CICS with WebSphere MQ on UNIX and Linux systems. Additionally, sample source code files are provided to enable you to develop the XA switches for other transaction messages.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The names of the switch load modules provided are:

Table 11. Essential code for CICS applications: XA initialization routine	
C (source)	C (exec) - add one of the following to your XAD.Stanza
amqzscix.c	amqzsc - TXSeries for AIX, Version 5.1, amqzsc - TXSeries for HP-UX, Version 5.1 amqzsc - TXSeries for Sun Solaris, Version 5.1
amqzscin.c	mqmc4swi - TXSeries for Windows, Version 5.1

Building libraries for use with TXSeries for Multiplatforms

Use this information when building libraries for use with TXSeries for Multiplatforms.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Pre-built switch load files are shared libraries (called DLLs on the Windows system) that you can use with CICS programs, which require a 2-phase commit transaction by using the XA protocol. The names of these pre-built libraries are in the table [Essential code for CICS applications: XA initialization routine](#). Sample source code is also supplied in the following directories:

Table 12. Installation directories on Windows, UNIX and Linux operating systems		
Platform	Directory	Source file
UNIX and Linux	MQ_INSTALLATION_PATH/ samp/	amqzscix.c
Windows	MQ_INSTALLATION_PATH\Tools \c\Samples	amqzscin.c

where `MQ_INSTALLATION_PATH` is the directory in which you installed IBM WebSphere MQ.

To build the switch load file from the sample source, follow the instructions appropriate for your operating system:

AIX

Issue the following command:

```
export MQM_HOME=/usr/mqm
echo "amqzscix" > tmp.exp
xlc_r $MQM_HOME/samp/amqzscix.c -I/usr/lpp/cics/include -I$MQM_HOME/inc -e amqzscix -bE:tmp.exp
-bM:SRE -o amqzsc /usr/lpp/cics/lib/regxa_swxa.o -L$MQM_HOME/lib -L/usr/lpp/cics/lib -lcicsrt -lEncina
-lEncServer -lpthreads -lsarpc -lmqmcics_r -lmqmx_r -lmqzi_r -lmqmcs_r
rm tmp.exp
```

Solaris

Issue the following command:

```
/opt/SUNWsp/ro/bin/cc -s -l/opt/encina/include amqzscix.c -G -o amqzscix -e
CICS_XA_Init -LMQ_INSTALLATION_PATH/lib -L/opt/encina/lib
-L/opt/dcelocal/lib /opt/cics/lib/reqxa_swxa.o
-lmqmcics -lmqmx -lmqzi -lmqmcs -lmqmzse -lcicsrt -lEncina -lEncSfs -ldce
```

HP-UX

Issue the following command:

```
cc -c -s -I/opt/encina/include MQ_INSTALLATION_PATH/samp/amqzscix.c -Aa +z -o amqzscix.o ld -b  
-o amqzscix amqzscix.o /opt/cics/lib/regxa_swxa.o +e CICS_XA_Init \  
-LMQ_INSTALLATION_PATH/lib -L/opt/encina/lib -L/opt/cics/lib  
-lmqmxr_r -lmqzi_r -lmqmcs_r -lmqmqse -ldbm -lc -lm
```

Linux platforms

Issue the following command:

```
gcc -m32 -shared -fPIC -o amqzscix amqzscix.c  
-IMQ_INSTALLATION_PATH/inc -I CICS_INSTALLATION_PATH/include  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib -Wl,-rpath-link,/usr/lib -Wl,--no-undefined  
-Wl,--allow-shlib-undefined -L CICS_LIB_PATH/regxa_swxa.o -lpthread -ldl -lc  
-shared -lmqzi_r -lmqmxr_r -lmqmcics_r -ldl -lc
```

Windows

Follow these steps:

1. Use the `cl` command to build `amqzscin.obj` by compiling at least the following variables:

```
cl.exe -c -IEncinaPath\include -IMQ_INSTALLATION_PATH\include -Gz -LD amqzscin.c
```

2. Create a module definition file named `mqmc1415.def`, which contains the following lines:

```
LIBRARY MQMC4SWI  
EXPORTS  
CICS_XA_Init
```

3. Use the **lib** command to build an export file and an import library by using at least the following option:

```
lib -def:mqmc4swi.def -out:mqmc4swi.lib
```

If the `lib` command is successful then an `mqmc4swi.exp` file is also built.

4. Use the `link` command to build `mqmc4swi.dll` by using at least the following option:

```
link.exe -dll -nod -out:mqmc4swi.dll  
amqzscin.obj CicsPath\lib\regxa_swxa.obj  
mqmc4swi.exp mqmcics4.lib  
CicsPath\lib\libcicsrt.lib  
DcePath\lib\libdce.lib DcePath\lib\pthreads.lib  
EncinaPath\lib\libEncina.lib  
EncinaPath\lib\libEncServer.lib  
msvcrt.lib kernel32.lib
```

IBM WebSphere MQ XA support and Tuxedo

IBM WebSphere MQ on Windows, UNIX and Linux systems can block Tuxedo-coordinated XA applications indefinitely in `xa_start`.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

This can occur only when two or more processes coordinated by Tuxedo in a single global transaction attempt to access IBM WebSphere MQ using the same transaction branch ID (XID). If Tuxedo gives each process in the global transaction a different XID to use with IBM WebSphere MQ, this cannot occur.

To avoid the problem, configure each application in Tuxedo that accesses IBM WebSphere MQ under a single global transaction ID (gtrid), within its own Tuxedo server group. Processes in the same server group use the same XID when accessing resource managers on behalf of a single gtrid, and are therefore vulnerable to blocking in `xa_start` in IBM WebSphere MQ. Processes in different server groups use separate XIDs when accessing resource managers and so do not have to serialize their transaction work in IBM WebSphere MQ.

Enabling the CICS two-phase commit process

To enable CICS to use a two-phase commit process to coordinate transactions that include MQI calls, add a CICS XAD resource definition stanza entry to the CICS region. Note, this topic is not applicable to z/OS.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Here is an example of adding an XAD stanza entry for WebSphere MQ for Windows, where <Drive> is the drive where WebSphere MQ is installed (for example, D:).

```
cicsadd -cxad -r<cics_region> \  
ResourceDescription="MQM XA Product Description" \  
SwitchLoadFile="<Drive>:\Program Files\IBM\WebSphere MQ\bin\mqmc4swi.dll" \  
XAOpen=<queue_manager_name>
```

For extended transactional clients, use the switch load file mqcc4swi.dll.

Here is an example of adding an XAD stanza entry for WebSphere MQ for UNIX and Linux systems, where `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed:

```
cicsadd -cxad -r<cics_region> \  
ResourceDescription="MQM XA Product Description" \  
SwitchLoadFile="MQ_INSTALLATION_PATH/lib/amqzsc" \  
XAOpen=<queue_manager_name>
```

For extended transactional clients, use the switch load file amqczsc.

For information about using the **cicsadd** command, see the *CICS Administration Reference*, or *CICS Administration Guide* for your platform.

Calls to WebSphere MQ can be included in a CICS transaction, and the WebSphere MQ resources will be committed or rolled back as directed by CICS. This support is not available to client applications.

You **must** issue an MQCONN from your CICS transaction in order to access WebSphere MQ resources, followed by a corresponding MQDISC on exit.

Enabling CICS user exits

A CICS user exit *point* (normally referred to as a *user exit*) is a place in a CICS module at which CICS can transfer control to a program that you have written (a user exit *program*), and at which CICS can resume control when your exit program has finished its work.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

Before using a CICS user exit, read the *CICS Administration Guide* for your platform.

Sample task termination exit

WebSphere MQ supplies sample source code for a CICS task termination exit.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

The sample source code is in the following directories:

Table 13. CICS task termination exits		
Platform	Directory	Source file
UNIX and Linux systems	<code>MQ_INSTALLATION_PATH/samp</code>	<code>amqzscgx.c</code>

Table 13. CICS task termination exits (continued)		
Platform	Directory	Source file
Windows	<code>MQ_INSTALLATION_PATH\Tools</code> <code>\c\Samples</code>	<code>amqzscgn.c</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

The build instructions for the sample task termination exit are contained in the comments near the top of each source file.

This exit is invoked by CICS at normal and abnormal task termination (after any sync point has been taken). No recoverable work is permitted in the exit program.

These functions are only used in a WebSphere MQ and CICS context in which the CICS version supports the XA interface. CICS refers to these libraries as "programs" or "user exits".

CICS has a number of user exits and `amqzscgx`, if used, is defined and enabled on CICS as the "Task termination user exit (UE014015)", that is, exit number 15.

When the task termination exit is called by CICS, CICS has already informed WebSphere MQ of the task's termination state and WebSphere MQ has taken the appropriate action (commit or rollback). All the exit does is to issue an MQDISC to clean up.

One purpose of installing and configuring your CICS system to use a task termination exit is to protect your system against some of the consequences of faulty application code. For example, if your CICS transaction ends abnormally without first calling MQDISC, and has no task termination exit installed, then you might see (within around 10 seconds) a subsequent unrecoverable failure of the CICS region. This is because WebSphere MQ's health thread, that runs in the `cicsas` process, will not have been posted and given time to clean up and return. The symptoms might be that the `cicsas` process ends immediately, having written FFST reports to `/var/mqm/errors` or the equivalent location on Windows.

Using the Microsoft Transaction Server (COM+)

COM+ (Microsoft Transaction Server) is designed to help users run business logic applications in a typical middle tier server.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

See [Features that can be used only with the primary installation on Windows](#) for important information.

COM+ divides work up into *activities*, which are typically short independent chunks of business logic, such as *transfer funds from account A to account B*. COM+ relies heavily on object orientation and in particular on COM; loosely a COM+ activity is represented by a COM (business) object.

COM+ is an integrated part of the operating system. To use COM+ on Windows 2000 and Windows XP, you need Hotfix Q313582 (also known as COM+ Rollup Package 19.1).

COM+ provides three services to the business object administrator, removing much of the worry from the business object programmer:

- Transaction management
- Security
- Resource pooling

You usually use COM+ with front-end code that is a COM client to the objects held within COM+, and back-end services such as a database, with WebSphere MQ bridging between the COM+ business object and the back-end.

The front-end code can be a stand-alone program, or an Active Server Page (ASP) hosted by the Microsoft Internet Information Server (IIS). The front-end code can be on the same computer as COM+ and its business objects, with connection through COM. Alternatively, the front-end code can be on a different

computer, with connection through DCOM. You can use different clients to access the same COM+ business object in different situations.

The back-end code can be on the same computer as COM+ and its business objects, or on a different computer with connection through any of the WebSphere MQ supported protocols.

Expiring global units of work

The queue manager can be configured to expire global units of work after a pre-configured interval of inactivity.

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

To enable this behavior, set the following environment variables:

- `AMQ_TRANSACTION_EXPIRY_RESCAN=<rescan interval in milliseconds>`
- `AMQ_XA_TRANSACTION_EXPIRY=<timeout interval in milliseconds>`



Attention: The environment variables only affect transactions that are in *Idle* state in table 6-4 of the XA Specification. That is, transactions that are not associated on any application thread, but for which the external Transaction Manager software has not yet called the **xa_prepare** function call.

External transaction managers keep only a log of transactions that are prepared, committed, or rolled back. If the external transaction manager goes down for any reason, on its return it drives prepared, committed, and rolled back transactions to completion but any active transactions that have yet to be prepared become orphaned. To avoid this, set the `AMQ_XA_TRANSACTION_EXPIRY` to allow for the expected interval between an application making MQI transactional API calls and completing the transaction, having carried out transactional work on other resource managers.

To ensure a timely cleanup after the `AMQ_XA_TRANSACTION_EXPIRY` expires, set the `AMQ_TRANSACTION_EXPIRY_RESCAN` value to a lower value than the `AMQ_XA_TRANSACTION_EXPIRY` interval, ideally so that the rescan occurs more than once within the `AMQ_XA_TRANSACTION_EXPIRY` interval.

Unit of recovery disposition

WebSphere MQ for z/OS provides unit of recovery dispositions. This feature allows you to configure whether the second phase of 2-phase commit transactions can be driven, for example, during recovery, when connected to another queue manager within the same queue-sharing group (QSG).

Note: This topic is also available in IBM MQ Version 8.0 and later versions. However, you cannot switch to a later version using the "Change version" list box. To go to the topic in a later version, edit the version number in the URL box in your browser.

WebSphere MQ for z/OS V7.0.1 and later supports the unit of recovery disposition.

Unit of recovery disposition

The unit of recovery disposition is related to an application's connection and subsequently any transactions that it starts. There are two possible unit of recovery dispositions.

- A GROUP unit of recovery disposition identifies that a transactional application is logically connected to the queue-sharing group and does not have an affinity to any specific queue manager. Any 2-phase commit transactions that it starts that have completed phase-1 of the commit process, that is, they are in doubt, can be inquired and resolved, when connected to any queue manager within the QSG. In a recovery scenario this means that the transaction coordinator does not have to reconnect to the same queue manager, which might be unavailable.
- A QMGR unit of recovery disposition identifies that an application has a direct affinity to the queue manager to which it is connected and any transactions that it starts also have this disposition.

In a recovery scenario the transaction coordinator must reconnect to the same queue manager to inquire, and resolve, any in-doubt transactions, irrespective of whether the queue manager belongs to a queue-sharing group.

Deciding which programming language to use

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

IBM WebSphere MQ provides support for the following programming procedural languages:

- C
- Visual Basic (Windows systems only)
- COBOL

These languages use the message queue interface (MQI) to access message queueing services. For more information about support for these languages see, [“Using procedural languages with WebSphere MQ” on page 75](#).

IBM WebSphere MQ provides support for:

- .NET
- ActiveX
- C++
- Java
- JMS

These languages use the IBM WebSphere MQ Object Model, which provides classes that provide the same functionality as WebSphere MQ calls and structures, but that are a more natural way of programming in an object-oriented environment. Some of the languages that use the IBM WebSphere MQ Object Model provide additional functions that are not available in the message queue interface (MQI). For more information about support for these languages see, [“Object-oriented programming with WebSphere MQ” on page 76](#).

Using procedural languages with WebSphere MQ

For detailed information about how to write your applications in your chosen language, see the following links:

- [“Coding in C” on page 79](#)
- [“Coding in Visual Basic” on page 83](#)
- [“Coding in COBOL” on page 81](#)

For an overview of the call interface for procedural languages, see [Call descriptions](#). This topic contains a list of the MQI calls, and each call shows you how to code the calls in each of these languages.

WebSphere MQ provides data definition files to help you to write your applications. For a full description, see [“IBM WebSphere MQ data definition files” on page 77](#).

If you can choose which language to code your programs in, consider the maximum length of the messages that your programs will process. If your programs will process only messages of a known maximum length, you can code them in any of the supported programming languages. But if you do not know the maximum length of the messages that the programs will have to process, the language you choose will depend on whether you are writing a CICS, IMS, or batch application:

IMS and batch

Code the programs in C, PL/I, or assembler language to use the facilities these languages offer for obtaining and releasing arbitrary amounts of memory. Alternatively, you could code your programs in COBOL, but use assembler language, PL/I, or C subroutines to get and release storage.

CICS

Code the programs in any language supported by CICS. The EXEC CICS interface provides the calls for managing memory, if necessary.

Object-oriented programming with WebSphere MQ

Some of the languages and programming frameworks that use the IBM WebSphere MQ Object Model provide additional functions that are not available in the message queue interface (MQI). For details of the classes, methods and properties provided by the IBM WebSphere MQ Object Model, see [“The IBM WebSphere MQ Object Model”](#) on page 83.

.NET

See [Using .NET](#) for information about coding .NET programs using the WebSphere MQ .NET classes. Message Service Clients for C/C++ and .NET provide an application programming Interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API.

C++

IBM WebSphere MQ provides C++ classes equivalent to WebSphere MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI. See [Using C++](#) for information about coding programs using the WebSphere MQ Object Model in C++. Message Service Clients for C/C++ and .NET provide an application programming Interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API.

Java

See [Using Java](#) for information about coding programs using the WebSphere MQ Object Model in Java. For information about the differences between IBM WebSphere MQ classes for Java and IBM WebSphere MQ classes to help you decide which to use, see [“Should I use IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS?”](#) on page 85.

JMS

WebSphere MQ also provides classes that implement the Java Message Service (JMS) specification. For details of the WebSphere MQ classes for JMS, see [Using Java](#). For information about the differences between IBM WebSphere MQ classes for Java and IBM WebSphere MQ classes to help you decide which to use, see [“Should I use IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS?”](#) on page 85.

Message Service Clients for C/C++ and .NET provide an application programming Interface (API) called XMS that has the same set of interfaces as the Java Message Service (JMS) API.

ActiveX

The WebSphere MQ ActiveX is commonly known as the MQAX. The MQAX is included as part of WebSphere MQ for Windows. Support for ActiveX has been stabilized at the WebSphere MQ Version 6.0 level. To exploit features introduced to WebSphere MQ later than Version 6.0, consider using .NET instead. Refer to [Using the Component Object Model Interface \(WebSphere MQ Automation Classes for ActiveX\)](#) for information about coding programs using the WebSphere MQ Object Model in ActiveX.

Related concepts

[Technical overview](#)

[“Developing applications”](#) on page 7

IBM WebSphere MQ provides several ways in which you can develop applications to send and receive messages that you need to support your business processes. You can also develop applications to manage your queue managers and related resources.

[“Application development concepts”](#) on page 8

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

Related reference

[Application development reference](#)

IBM WebSphere MQ data definition files

IBM WebSphere MQ provides data definition files to help you to write your applications.

Data definition files are also known as:

Language	Data definitions
C	Include files or header files
Visual Basic	Module files (32-bit versions only)
COBOL	Copy files
Assembler	Macros
PL/I	Include files

The data definition files to help you to write channel exits are described in [WebSphere MQ COPY, header, include, and module files](#).

The data definition files to help you to write installable services exits are described in [“User exits, API exits, and WebSphere MQ installable services”](#) on page 357.

For data definition files supported on C++, see [Using C++](#).

The names of the data definition files have the prefix CMQ, and a suffix that is determined by the programming language:

Suffix	Language
a	Assembler language
b	Visual Basic
c	C
l	COBOL (without initialized values)
p	PL/I
v	COBOL (with default values set)

Installation library

The name **thlqual** is the high-level qualifier of the installation library on z/OS.

This topic introduces WebSphere MQ data definition files, under these headings:

- [“C language include files”](#) on page 77
- [“Visual Basic module files”](#) on page 78
- [“COBOL copy files”](#) on page 78

C language include files

The WebSphere MQ C include files are listed in [C header files](#). They are installed in the following directories or libraries:

Platform	Installation directory or library
UNIX platforms	<i>MQ_INSTALLATION_PATH</i> /inc/
Windows systems	<i>MQ_INSTALLATION_PATH</i> \Tools\c\include

where *MQ_INSTALLATION_PATH* represents the high-level directory in which WebSphere MQ is installed.

Note: For UNIX platforms, the include files are symbolically linked into */usr/include*.

For more information about the structure of directories, see [Planning file system support](#).

Visual Basic module files

WebSphere MQ for Windows provides four Visual Basic module files.

They are listed in [Visual Basic module files](#) and installed in

```
MQ_INSTALLATION_PATH\Tools\Samples\VB\Include
```

COBOL copy files

For COBOL, WebSphere MQ provides separate copy files containing the named constants, and two copy files for each of the structures.

There are two copy files for each structure because each is provided both with and without initial values:

- In the WORKING-STORAGE SECTION of a COBOL program, use the files that initialize the structure fields to default values. These structures are defined in the copy files that have names suffixed with the letter V (values).
- In the LINKAGE SECTION of a COBOL program, use the structures without initial values. These structures are defined in copy files that have names suffixed with the letter L (linkage).

The WebSphere MQ COBOL copy files are listed in [COBOL COPY files](#). They are installed in the following directories:

Platform	Installation directory or library
Other UNIX platforms	<i>MQ_INSTALLATION_PATH</i> /inc/
Windows	<i>MQ_INSTALLATION_PATH</i> \Tools\cobol\copybook (for Micro Focus COBOL) <i>MQ_INSTALLATION_PATH</i> \Tools\cobol\copybook\VAcobol (for IBM VisualAge® COBOL)

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Include in your program only those files that you need. Do this with one or more COPY statements after a level-01 declaration. This means that you can include multiple versions of the structures in a program if necessary. Note that CMQV is a large file.

Here is an example of COBOL code to include the CMQMDV copy file:

```
01 MQM-MESSAGE-DESCRIPTOR.  
  COPY CMQMDV.
```

Each structure declaration begins with a level-01 item; you can declare several instances of the structure by coding the level-01 declaration followed by a COPY statement to copy in the remainder of the structure declaration. To refer to the appropriate instance, use the IN keyword.

Here is an example of COBOL code to include two instances of CMQMDV:

```
* Declare two instances of MQMD  
01 MY-CMQMD.  
  COPY CMQMDV.  
01 MY-OTHER-CMQMD.  
  COPY CMQMDV.  
*  
* Set MSGTYPE field in MY-OTHER-CMQMD  
  MOVE MQMT-REQUEST TO MQMD-MSGTYPE IN MY-OTHER-CMQMD.
```

Align the structures on 4-byte boundaries. If you use the COPY statement to include a structure following an item that is not the level-01 item, ensure that the structure is a multiple of 4-bytes from the start of the level-01 item. If you do not do this, you might reduce the performance of your application.

The structures are described in [Data types used in the MQI](#). The descriptions of the fields in the structures show the names of fields without a prefix. In COBOL programs, prefix the field names with the name of the structure followed by a hyphen, as shown in the COBOL declarations. The fields in the structure copy files are prefixed in this way.

The field names in the declarations in the structure copy files are in uppercase. You can use mixed case or lowercase instead. For example, the field *StrucId* of the MQGMO structure is shown as MQGMO-STRUCID in the COBOL declaration and in the copy file.

The V-suffix structures are declared with initial values for all the fields, so you need to set only those fields where the value required is different from the initial value.

Coding in C

Note the information in the following sections when coding WebSphere MQ programs in C.

- [“Parameters of the MQI calls” on page 79](#)
- [“Parameters with undefined data type” on page 79](#)
- [“Data types” on page 79](#)
- [“Manipulating binary strings” on page 80](#)
- [“Manipulating character strings” on page 80](#)
- [“Initial values for structures” on page 80](#)
- [“Initial values for dynamic structures” on page 81](#)
- [“Use from C++” on page 81](#)

Parameters of the MQI calls

Parameters that are *input-only* and of type MQHCONN, MQHOBJ, MQHMSG, or MQLONG are passed by value; for all other parameters, the *address* of the parameter is passed by value.

Not all parameters that are passed by address need to be specified every time a function is invoked. Where a particular parameter is not required, a null pointer can be specified as the parameter on the function invocation, in place of the address of the parameter data. Parameters for which this is possible are identified in the call descriptions.

No parameter is returned as the value of the function; in C terminology, this means that all functions return void.

The attributes of the function are defined by the MQENTRY macro variable; the value of this macro variable depends on the environment.

Parameters with undefined data type

The MQGET, MQPUT, and MQPUT1 functions each have a *Buffer* parameter that has an undefined data type. This parameter is used to send and receive the application's message data.

Parameters of this sort are shown in the C examples as arrays of MQBYTE. You can declare the parameters in this way, but it is typically more convenient to declare them as the structure that describes the layout of the data in the message. The function parameter is declared as a pointer-to-void, and so the address of any data can be specified as the parameter on the function invocation.

Data types

All data types are defined with the typedef statement.

For each data type, the corresponding pointer data type is also defined. The name of the pointer data type is the name of the elementary or structure data type prefixed with the letter P to denote a pointer. The attributes of the pointer are defined by the MQPOINTER macro variable; the value of this macro variable depends on the environment. The following code illustrates how to declare pointer data types:

```
#define MQPOINTER          /* depends on environment */
...
typedef MQLONG  MQPOINTER PMQLONG; /* pointer to MQLONG */
typedef MQMD    MQPOINTER PMQMD;   /* pointer to MQMD */
```

Manipulating binary strings

Strings of binary data are declared as one of the MQBYTEn data types.

Whenever you copy, compare, or set fields of this type, use the C functions `memcpy`, `memcmp`, or `memset`:

```
#include <string.h>
#include "cmqc.h"

MQMD MyMsgDesc;

memcpy(MyMsgDesc.MsgId,          /* set "MsgId" field to nulls */
       MQMI_NONE,               /* ...using named constant */
       sizeof(MyMsgDesc.MsgId));

memset(MyMsgDesc.CorrelId,       /* set "CorrelId" field to nulls */
       0x00,                    /* ...using a different method */
       sizeof(MQBYTE24));
```

Do not use the string functions `strcpy`, `strcmp`, `strncpy`, or `strncmp` because these do not work correctly with data declared as MQBYTE24.

Manipulating character strings

When the queue manager returns character data to the application, the queue manager always pads the character data with blanks to the defined length of the field. The queue manager does not return null-terminated strings, but you can use them in your input. Therefore, when copying, comparing, or concatenating such strings, use the string functions `strncpy`, `strncmp`, or `strncat`.

Do not use the string functions that require the string to be terminated by a null (`strcpy`, `strcmp`, and `strcat`). Also, do not use the function `strlen` to determine the length of the string; use instead the `sizeof` function to determine the length of the field.

Initial values for structures

The include file `<cmqc.h>` defines various macro variables that you can use to provide initial values for the structures when declaring instances of those structures. These macro variables have names of the form `MQxxx_DEFAULT`, where `MQxxx` represents the name of the structure. Use them like this:

```
MQMD  MyMsgDesc = {MQMD_DEFAULT};
MQPMO MyPutOpts = {MQPMO_DEFAULT};
```

For some character fields, the MQI defines particular values that are valid (for example, for the *StrucId* fields or for the *Format* field in MQMD). For each of the valid values, two macro variables are provided:

- One macro variable defines the value as a string with a length, excluding the implied null, that exactly matches the defined length of the field. For example, the symbol `~` represents a blank character:

```
#define MQMD_STRUC_ID "MD~~"
#define MQFMT_STRING "MQSTR~~~"
```

Use this form with the `memcpy` and `memcmp` functions.

- The other macro variable defines the value as an array of `char`; the name of this macro variable is the name of the string form suffixed with `_ARRAY`. For example:


```
#define MQMD_STRUC_ID_ARRAY 'M','D',' ',' '
#define MQFMT_STRING_ARRAY 'M','Q','S','T','R',' ',' ',' '

```

Use this form to initialize the field when an instance of the structure is declared with values different from those provided by the MQMD_DEFAULT macro variable.

Initial values for dynamic structures

When a variable number of instances of a structure are required, the instances are typically created in main storage obtained dynamically using the calloc or malloc functions.

To initialize the fields in such structures, the following technique is recommended:

1. Declare an instance of the structure using the appropriate MQxxx_DEFAULT macro variable to initialize the structure. This instance becomes the *model* for other instances:

```
MQMD ModelMsgDesc = {MQMD_DEFAULT};
/* declare model instance */

```

Code the static or auto keywords on the declaration to give the model instance static or dynamic lifetime, as required.

2. Use the calloc or malloc functions to obtain storage for a dynamic instance of the structure:

```
PMQMD InstancePtr;
InstancePtr = malloc(sizeof(MQMD));
/* get storage for dynamic instance */

```

3. Use the memcpy function to copy the model instance to the dynamic instance:

```
memcpy(InstancePtr,&ModelMsgDesc,sizeof(MQMD));
/* initialize dynamic instance */

```

Use from C++

For the C++ programming language, the header files contain the following additional statements that are included only when a C++ compiler is used:

```
#ifndef __cplusplus
extern "C" {
#endif

/* rest of header file */

#ifdef __cplusplus
}
#endif

```

Coding in COBOL

Note the information in the following section when coding WebSphere MQ programs in COBOL.

Named constants

The names of constants are shown containing the underscore character (_) as part of the name. In COBOL, you must use the hyphen character (-) in place of the underscore. Constants that have character-string values use the single quotation mark character (') as the string delimiter. To make the compiler accept this character, use the compiler option APOST.

The copy file CMQV contains declarations of the named constants as level-10 items. To use the constants, declare the level-01 item explicitly, then use the COPY statement to copy in the declarations of the constants:

```
WORKING-STORAGE SECTION.  
01 MQM-CONSTANTS.  
   COPY CMQV.
```

However, this method causes the constants to occupy storage in the program even if they are not referred to. If the constants are included in many separate programs within the same run unit, multiple copies of the constants will exist; this might result in a significant amount of main storage being used. You can avoid this by adding the GLOBAL clause to the level-01 declaration:

```
* Declare a global structure to hold the constants  
01 MQM-CONSTANTS GLOBAL.  
   COPY CMQV.
```

This allocates storage for only *one* set of constants within the run unit; the constants, however, can be referred to by *any* program within the run unit, not just the program that contains the level-01 declaration.

Ensuring structure alignment

Care should be taken to ensure IBM WebSphere MQ structures that are passed on to start on the MQ call's must be aligned on word boundaries. A word boundary is 4 bytes for 32-bit processes, 8 bytes for 64 bit processes and 16 bytes for 128 bit processes (IBM i).

Where possible, place all IBM WebSphere MQ structures together so they are all boundary aligned.

Coding in pTAL

Note the information in the following section when coding IBM WebSphere MQ programs in pTAL.

HP Integrity NonStop Server

Defining and initializing IBM WebSphere MQ structures

pTAL structure definitions for IBM WebSphere MQ structures are provided with names that end with ^DEF. For example, the following pTAL declarations would be coded to create a IBM WebSphere MQ Message Descriptor (MQMD) structure and a IBM WebSphere MQ Put Message Options (MQPMO) structure.

```
STRUCT MYMD(MQMD^DEF);      ! Declare an MQMD structure  
STRUCT MYPMO(MQPMO^DEF);    ! Declare an MQPMO structure
```

IBM WebSphere MQ provides pTAL DEFINE with names that end with ^DEFAULT to initialize IBM WebSphere MQ structures with default values. The following pTAL statements are coded to assign default values to the declared MQMD and the MQPMO structures:

```
MQMD^DEFAULT(MYMD);          ! Assign default values to an MQMD structure  
MQPMO^DEFAULT(MYPMO);        ! Assign default values to an MQPMO structure
```

You can declare and initialize other IBM WebSphere MQ structures by using similar code.

pTAL and the CRE

pTAL programs cannot initialize the Common Runtime Environment and therefore they must be used with a C-language or COBOL main routine.

The pTAL samples that are provided with IBM WebSphere MQ use a C-language mainline routine that is called AMQSPTM0.C

Parameters with MQCHAR data type

The MQGET, MQPUT, and MQPUT1 procedures each have a **Buffer** parameter that has an MQCHAR .EXT data type. This parameter is used to send and receive the application's message data.

Parameters of this sort are shown in the pTAL samples as arrays of string. You can declare the parameters in this way, but it is typically more convenient to declare them as the structure that describes the layout of the data in the message. The procedure parameter is declared as an MQCHAR .EXT, but the address of any data can be specified as the parameter on the procedure invocation.

Manipulating character strings

When the queue manager returns character data to the application, the queue manager always pads the character data with blanks to the defined length of the field. The queue manager does not return null-terminated strings, but you can use them in your input.

Coding in Visual Basic

Note the information in the following section when coding WebSphere MQ programs in Visual Basic.

Note: Outside the .NET environment, support for Visual Basic (VB) in WebSphere MQ has been stabilized at the V6.0 level. Most new function added to WebSphere MQ 7.0 or later is not available to VB applications. If you are programming in VB.NET, use the WebSphere MQ .NET classes. For more information, see [Using .NET](#).

Visual Basic is supported only on Windows.

To avoid unintended translation of binary data passing between Visual Basic and WebSphere MQ, use an MQBYTE definition instead of MQSTRING. CMQB.BAS defines several new MQBYTE types that are equivalent to a C byte definition and uses these within WebSphere MQ structures. For example, for the MQMD (message descriptor) structure, MsgId (message identifier) is defined as MQBYTE24.

Visual Basic does not have a pointer data type, so references to other WebSphere MQ data structures are by offset rather than pointer. Declare a compound structure consisting of the two component structures, and specify the compound structure on the call. WebSphere MQ support for Visual Basic provides an MQCONNAny call to make this possible and allow client applications to specify the channel properties on a client connection. It accepts an untyped structure (MQCNOCD) in place of the typical MQCNO structure.

The MQCNOCD structure is a compound structure consisting of an MQCNO followed by an MQCD. This structure is declared in the exits header file CMQXB. Use the routine MQCNOCD_DEFAULTS to initialize an MQCNOCD structure. A sample making MQCONN calls is provided (amqscnxb.vbp).

MQCONNAny has the same parameters as MQCONN, except that the *ConnectOpts* parameter is declared as being of Any data type rather than of MQCNO data type. This allows the function to accept either the MQCNO or the MQCNOCD structure. This function is declared in the main header file CMQB.

The IBM WebSphere MQ Object Model

The IBM WebSphere MQ Object Model consists of classes, methods and properties. Use this information to learn about each of these concepts.

The IBM WebSphere MQ Object Model consists of the following:

- *Classes* representing familiar WebSphere MQ concepts such as queue managers, queues, and messages.
- *Methods* on each class corresponding to MQI calls.
- *Properties* on each class corresponding to attributes of WebSphere MQ objects.

When creating a WebSphere MQ application using the WebSphere MQ Object Model, you create instances of these classes in the program. An instance of a class in object-oriented programming is called an *object*. When an object has been created, you interact with the object by examining or setting the values of the

object's properties (the equivalent of issuing an MQINQ or MQSET call), and by making method calls against the object (the equivalent of issuing the other MQI calls).

These topics describe each of the WebSphere MQ Object Models in detail:

- [“Classes” on page 84](#)
- [“Object references” on page 85](#)
- [“Return codes” on page 85](#)

Classes

The WebSphere MQ Object Model provides the following base set of classes.

The actual implementation of the model varies slightly between the different supported object-oriented environments.

MQQueueManager

An object of the MQQueueManager class represents a connection to a queue manager. It has methods to Connect(), Disconnect(), Commit(), and Backout() (the equivalent of MQCONN or MQCONNEX, MQDISC, MQCMIT, and MQBACK). It has properties corresponding to the attributes of a queue manager. Accessing a queue manager attribute property implicitly connects to the queue manager if not already connected. Destroying an MQQueueManager object implicitly disconnects from the queue manager.

MQQueue

An object of the MQQueue class represents a queue. It has methods to Put() and Get() messages to and from the queue (the equivalent of MQPUT and MQGET). It has properties corresponding to the attributes of a queue. Accessing a queue attribute property, or issuing a Put() or Get() method call, implicitly opens the queue (the equivalent of MQOPEN). Destroying an MQQueue object implicitly closes the queue (the equivalent of MQCLOSE).

MQTopic

An object of the MQTopic class represents a topic. It has methods to Put() (publish) and Get() (receive or subscribe) messages to and from the topic (the equivalent of MQPUT and MQGET). It has properties corresponding to the attributes of a topic. An MQTopic object can only be accessed for publication or subscription, not both simultaneously. When used for receiving messages the MQTopic object can be created with an unmanaged or managed subscription and as a durable or non-durable subscriber - multiple overloaded constructors are provided for these differing scenarios.

MQMessage

An object of the MQMessage class represents a message to be put on a queue or got from a queue. It contains a buffer, and encapsulates both application data and MQMD. It has properties corresponding to MQMD fields, and methods that allow you to write and read user data of different types (for example, strings, long integers, short integers, single bytes) to and from the buffer.

MQPutMessageOptions

An object of the MQPutMessageOptions class represents the MQPMO structure. It has properties corresponding to MQPMO fields.

MQGetMessageOptions

An object of the MQGetMessageOptions class represents the MQGMO structure. It has properties corresponding to MQGMO fields.

MQProcess

An object of the MQProcess class represents a process definition (used with triggering). It has properties that represent the attributes of a process definition.

MQDistributionList

An object of the MQDistributionList class represents a distribution list (used to send multiple messages with a single MQPUT). It contains a list of MQDistributionListItem objects.

MQDistributionListItem

An object of the MQDistributionListItem class represents a single distribution list destination. It encapsulates the MQOR, MQRR, and MQPMR structures, and has properties corresponding to the fields of these structures.

Object references

In a WebSphere MQ program that uses the MQI, WebSphere MQ returns connection handles and object handles to the program.

These handles must be passed as parameters on subsequent WebSphere MQ calls. With the WebSphere MQ Object Model, these handles are hidden from the application program. Instead, the creation of an object from a class results in an object reference being returned to the application program. It is this object reference that is used when making method calls and property accesses against the object.

Return codes

Issuing a method call or setting a property value results in return codes being set.

These return codes are a completion code and a reason code, and are themselves properties of the object. The values of completion code and reason code are the same as those defined for the MQI, with some extra values specific to the object-oriented environment.

Should I use IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS?

A Java application can use either IBM WebSphere MQ classes for Java or IBM WebSphere MQ classes for JMS to access IBM WebSphere MQ resources. Each approach has its advantages.

IBM WebSphere MQ classes for Java encapsulates the Message Queue Interface (MQI), the native IBM WebSphere MQ API, and uses the same object model as other object-oriented interfaces, whereas IBM WebSphere MQ classes for Java Message Service implements Sun's Java Message Service (JMS) interfaces.

If you are familiar with IBM WebSphere MQ in environments other than Java, using either procedural or object-oriented languages, you can transfer your existing knowledge to the Java environment by using IBM WebSphere MQ classes for Java. You can also exploit the full range of features of IBM WebSphere MQ, not all of which are available in IBM WebSphere MQ classes for JMS.

If you are not familiar with IBM WebSphere MQ, or already have JMS experience, you might find it easier to use the familiar JMS API to access IBM WebSphere MQ resources, by using IBM WebSphere MQ classes for JMS. JMS is also an integral part of the Java Platform, Enterprise Edition (Java EE) platform. Java EE applications can use message-driven beans (MDBs) to process messages asynchronously, and MDBs can process only JMS messages. JMS is also the standard mechanism for Java EE to interact with asynchronous messaging systems such as IBM WebSphere MQ. Every application server that is Java EE compliant must include a JMS provider, therefore you can use JMS to communicate between different application servers or you can port an application from one JMS provider to another without any change to the application.

Designing IBM WebSphere MQ applications

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

When designing an IBM WebSphere MQ application consider the following questions and options:

Type of application

What is the purpose of your application? See the following links for information about that different types of application you can develop:

- Server

- Client
- Publish/subscribe
- Web services
- User exits, API exits, and installable services

Additionally, you can also write your own applications to automate administration of IBM WebSphere MQ. For more information, see [Introduction to the WebSphere MQ Administration Interface \(MQAI\)](#) and [Automating administration tasks](#).

Programming language

IBM WebSphere MQ supports a number of procedural and object-oriented programming languages for writing applications. For more information see, [“Deciding which programming language to use”](#) on page 75.

Applications for more than one platform

Will your application run on more than one platform? Do you have a strategy to move to a different platform from the one that you use today? If the answer to either of these questions is yes, ensure that you code your programs for platform independence.

If you are using C, code in ANSI standard C. Use a standard C library function rather than an equivalent platform-specific function even if the platform-specific function is faster or more efficient. The exception is when efficiency in the code is paramount, when you should code for both situations using `#ifdef`. For example:

```
#ifdef _AIX
    AIX specific code
#else
    generic code
#endif
```

Types of queues

Do you want to create a queue each time that you need one, or do you want to use queues that have already been set up? Do you want to delete a queue when you have finished using it, or is it going to be used again? Do you want to use alias queues for application independence? To see what types of queues are supported, refer to [Queues](#).

Using queue manager clusters

You might want to take advantage of the simplified system administration, and increased availability, scalability, and workload balancing that are possible when you use clusters. See [Queue manager clusters](#) for more information.

Types of messages

You might want to use datagrams for simple messages, but request messages (for which you expect replies) for other situations. You might want to assign different priorities to some of your messages. For more information about designing messages, see [“Designing your messages”](#) on page 88.

Using publish/subscribe or point-to-point messaging

Using publish/subscribe messaging, a sending application sends the information that it wants to share in an IBM WebSphere MQ message to a standard destination managed by IBM WebSphere MQ publish/subscribe, and lets IBM WebSphere MQ handle the distribution of that information. The target application does not have to know anything about the source of the information it receives, it just registers an interest in one or more topics and receives that information when it is available. For more information about publish/subscribe messaging, see [Introduction to IBM WebSphere MQ publish/subscribe messaging](#).

Using point-to-point messaging, a sending application sends a message to a specific queue, from where it knows a receiving application will retrieve it. A receiving application gets messages from a specific queue and acts on their contents. An application will often function both as a sender and a receiver, sending a query to another application and receiving a response.

Controlling your IBM WebSphere MQ programs

You might want to start some programs automatically or make programs wait until a particular message arrives on a queue (using the IBM WebSphere MQ *triggering* feature, see [“Starting IBM WebSphere MQ applications using triggers”](#) on page 315). Alternatively, you might want to start another instance of an application when the messages on a queue are not getting processed fast enough (using the IBM WebSphere MQ *instrumentation events* feature as described in [Instrumentation events](#)).

Running your application on an IBM WebSphere MQ client

The full MQI is supported in the client environment and this enables almost any IBM WebSphere MQ application to be relinked to run on an IBM WebSphere MQ MQI client. Link the application on the IBM WebSphere MQ MQI client to the MQIC library, rather than to the MQI library.

Note: An application running on an IBM WebSphere MQ client can connect to more than one queue manager concurrently, or use a queue manager name with an asterisk (*) on an MQCONN or MQCONNEX call. Change the application if you want to link to the queue manager libraries instead of the client libraries, as this function will not be available.

See [“Running applications in the IBM WebSphere MQ MQI client environment”](#) on page 343 for more information.

Application performance

Design decisions can impact your application performance, for suggestions for enhancing performance of IBM WebSphere MQ applications, see [“Application design and performance”](#) on page 89.

Advanced IBM WebSphere MQ techniques

For more advanced applications you might want to use some advanced IBM WebSphere MQ techniques such as correlating replies, and generating and sending IBM WebSphere MQ context information. For more information, see [“Advanced IBM WebSphere MQ techniques”](#) on page 90.

Securing your data and maintaining its integrity

You can use the context information that is passed with a message to test that the message has been sent from an acceptable source. You can use the syncpointing facilities provided by IBM WebSphere MQ or your operating system to ensure that your data remains consistent with other resources (see [“Committing and backing out units of work”](#) on page 309 for further details). You can use the *persistence* feature of IBM WebSphere MQ messages to assure the delivery of important messages.

Testing IBM WebSphere MQ applications

The application development environment for IBM WebSphere MQ programs is no different from that for any other application, so you can use the same development tools as well as the IBM WebSphere MQ trace facilities.

Handling exceptions and errors

You need to consider how to process messages that cannot be delivered, and how to resolve error situations that are reported to you by the queue manager. For some reports, you must set report options on MQPUT.

Related concepts

[IBM WebSphere MQ technical overview](#)

[“Application development concepts”](#) on page 8

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Writing a queuing application”](#) on page 186

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Writing client applications”](#) on page 335

What you need to know to write client applications on WebSphere MQ.

[“Using .NET”](#) on page 534

WebSphere MQ classes for .NET allow a program written in the .NET programming framework to connect to WebSphere MQ as a WebSphere MQ MQI client or to connect directly to a WebSphere MQ server.

[“Using C++” on page 601](#)

WebSphere MQ provides C++ classes equivalent to WebSphere MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

[“Using WebSphere MQ classes for JMS” on page 684](#)

WebSphere MQ classes for Java Message Service (WebSphere MQ classes for JMS) is the JMS provider that is supplied with WebSphere MQ. As well as implementing the interfaces defined in the javax.jms package, WebSphere MQ classes for JMS provides two sets of extensions to the JMS API.

[“Using WebSphere MQ classes for Java” on page 625](#)

WebSphere MQ classes for Java enable you to use WebSphere MQ in a Java environment. A Java application can use either WebSphere MQ classes for Java or WebSphere MQ classes for JMS to access WebSphere MQ resources.

[“Using the Component Object Model Interface \(WebSphere MQ Automation Classes for ActiveX\)” on page 987](#)

The WebSphere MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access WebSphere MQ.

Designing your messages

Consider the aspects given in this information to help you design messages.

You create a message when you use an MQI call to put the message on a queue. As input to the call, you supply some control information in a *message descriptor* (MQMD) and the data that you want to send to another program. But at the design stage, you need to consider the following, because they affect the way that you create your messages:

Type of message to use

Are you designing a simple application in which you can send a message, then take no further action? Or are you asking for a reply to a question? If you are asking a question, you might include in the message descriptor the name of the queue on which you want to receive the reply.

Do you want your request and reply messages to be synchronous? This implies that you set a timeout period for the reply to answer your request, and if you do not receive the reply within that period, it is treated as an error.

Or would you prefer to work asynchronously, so that your processes do not have to depend upon the occurrence of specific events, such as common timing signals?

Another consideration is whether you have all your messages inside a unit of work.

Assigning different priorities to messages

You can assign a priority value to each message, and define the queue so that it maintains its messages in order of their priority. If you do this, when another program retrieves a message from the queue, it always gets the message with the highest priority. If the queue does not maintain its messages in priority order, a program that retrieves messages from the queue will retrieve them in the order in which they were added to the queue.

Programs can also select a message using the identifier that the queue manager assigned when the message was put on the queue. Alternatively, you can generate your own identifiers for each of your messages.

Effect of restarting queue manager on messages

The queue manager preserves all persistent messages, recovering them when necessary from the WebSphere MQ log files, when it is restarted. Nonpersistent messages and temporary dynamic queues are not preserved. Any messages that you do not want discarded must be defined as persistent when they are created. When writing an application for WebSphere MQ for Windows or WebSphere MQ on UNIX and Linux systems, make sure that you know how your system has been set up in respect of log file allocation to reduce the risk of designing an application that will run to the log file limits.

Giving information about yourself to the recipient of messages

Usually, the queue manager sets the user ID, but suitably authorized applications can also set this field, so that you can include your own user ID and other information that the receiving program can use for accounting or security purposes.

Amount of receiving queues

If a message might need to be put on several queues, you can use a distribution list, or publish to a topic.

Application design and performance

There are a number of ways in which poor program design can affect performance. These can be difficult to detect because the program can appear to perform well itself, but affect the performance of other tasks. Several problems specific to programs making WebSphere MQ calls are explained in this topic.

Here are a few ideas to help you to design efficient applications:

- Design your application so that processing goes on in parallel with a user's thinking time:
 - Display a panel and allow the user to start typing while the application is still initializing.
 - Get the data that you need in parallel from different servers.
- Keep connections and queues open if you are going to reuse them instead of repeatedly opening and closing, connecting and disconnecting.
- However, a server application that is putting only one message should use MQPUT1.
- Queue managers are optimized for messages between 4 KB and 100 KB in size. Very large messages are inefficient; it is probably better to send 100 messages of 1 MB each than a single 100 MB message. Very small messages are also inefficient. The queue manager does the same amount of work for a single-byte message as for a 4 KB message.
- Keep your messages within a unit of work, so that they can be committed or backed out simultaneously.
- Use the nonpersistent option for messages that do not need to be recoverable.
- If you need to send a message to a number of target queues, consider using a distribution list.

Effect of message length

The amount of data in a message can affect the performance of the application that processes the message. To achieve the best performance from your application, send only the essential data in a message. For example, in a request to debit a bank account, the only information that might need to be passed from the client to the server application is the account number and the amount of the debit.

Effect of message persistence

Persistent messages are usually logged. Logging messages reduces the performance of your application, so use persistent messages for essential data only. If the data in a message can be discarded if the queue manager stops or fails, use a nonpersistent message.

Searching for a particular message

The MQGET call usually retrieves the first message from a queue. If you use the message and correlation identifiers (*MsgId* and *CorrelId*) in the message descriptor to specify a particular message, the queue manager has to search the queue until it finds that message. Using the MQGET call in this way affects the performance of your application.

Queues that contain messages of different lengths

If your application cannot use messages of a fixed length, grow and shrink the buffers dynamically to suit the typical message size. If the application issues an MQGET call that fails because the buffer is too small, the size of the message data is returned. Add code to your application so that the buffer is resized accordingly and the MQGET call is re-issued.

Note: if you do not set the *MaxMsgLength* attribute explicitly, it defaults to 4 MB, which might be very inefficient if this is used to influence the application buffer size.

Frequency of sync points

Programs that issue very large numbers of MQPUT or MQGET calls within sync point, without committing them, can cause performance problems. Affected queues can fill up with messages that are currently inaccessible, while other tasks might be waiting to get these messages. This has implications in terms of storage, and in terms of threads tied up with tasks that are attempting to get messages.

Use of the MQPUT1 call

Use the MQPUT1 call only if you have a single message to put on a queue. If you want to put more than one message, use the MQOPEN call, followed by a series of MQPUT calls and a single MQCLOSE call.

Number of threads in use

For WebSphere MQ for Windows, an application might require a large number of threads. Each queue manager process is allocated a maximum allowable number of application threads.

Applications might use too many threads. Consider whether the application takes into account this possibility and that it takes actions either to stop or to report this type of occurrence.

Advanced IBM WebSphere MQ techniques

For a simple IBM WebSphere MQ application, you need to decide which WebSphere MQ objects to use in your application, and which types of message you want to use. For a more advanced application, you might want to use some of the techniques introduced in the following sections.

Waiting for messages

A program that is serving a queue can await messages by:

- Waiting until either a message arrives, or a specified time interval expires (see [“Waiting for messages” on page 256](#)).
- Establishing a callback exit to be driven when a message arrives; see [“Asynchronous consumption of IBM WebSphere MQ messages” on page 32](#).
- Making periodic calls on the queue to see whether a message has arrived (*polling*). This is not typically advisable because it can have performance implications.

Correlating replies

In WebSphere MQ applications, when a program receives a message that requests it to do some work, the program typically sends one or more reply messages to the requester.

To help the requester to associate these replies with its original request, an application can set a *correlation identifier* field in the descriptor of each message. Programs then copy the message identifier of the request message into the correlation identifier field of their reply messages.

Setting and using context information

Context information is used for associating messages with the user who generated them, and for identifying the application that generated the message. Such information is useful for security, accounting, auditing, and problem determination.

When you create a message, you can specify an option that requests that the queue manager associates default context information with your message.

For more information about using and setting context information, see [“Message context” on page 37](#).

Starting WebSphere MQ programs automatically

Use WebSphere MQ *triggering* to start a program automatically when messages arrive on a queue.

You can set trigger conditions on a queue so that a program starts to process that queue:

- Every time that a message arrives on the queue
- When the first message arrives on the queue
- When the number of messages on the queue reaches a predefined number

For more information about triggering, see [“Starting IBM WebSphere MQ applications using triggers” on page 315](#). Triggering is just one way of starting a program automatically. For example, you can start a program automatically on a timer using non-WebSphere MQ facilities.

WebSphere MQ can define service objects to start WebSphere MQ programs when the queue manager starts; see [Service objects](#).

Generating WebSphere MQ reports

You can request the following reports within an application:

- Exception reports
- Expiry reports
- Confirm-on-arrival (COA) reports
- Confirm-on-delivery (COD) reports
- Positive action notification (PAN) reports
- Negative action notification (NAN) reports

These are described in [“Report messages” on page 11](#).

Clusters and message affinities

Before starting to use clusters with multiple definitions for the same queue, examine your applications to see whether there are any that require an exchange of related messages.

Within a cluster, a message can be routed to any queue manager that hosts an instance of the appropriate queue. Therefore, the logic of applications with message affinities can be upset.

For example, you might have two applications that rely on a series of messages flowing between them in the form of questions and answers. It might be important that all the questions are sent to the same queue manager and that all the answers are sent back to the other queue manager. In this situation, it is important that the workload management routine does not send the messages to any queue manager that just happens to host an instance of the appropriate queue.

Where possible, remove the affinities. Removing message affinities improves the availability and scalability of applications.

For more information, see [Handling message affinities](#).

Sample WebSphere MQ programs

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

- [“Sample programs for distributed platforms” on page 92](#)

Related concepts

[“Application development concepts” on page 8](#)

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Writing a queuing application” on page 186](#)

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Writing client applications” on page 335](#)

What you need to know to write client applications on WebSphere MQ.

[“Using Web services in WebSphere MQ” on page 904](#)

You can develop IBM WebSphere MQ applications for Web services using the IBM WebSphere MQ transport for SOAP or the IBM WebSphere MQ bridge for HTTP.

[“Writing publish/subscribe applications” on page 266](#)

Start writing publish/subscribe WebSphere MQ applications.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

[“Handling program errors” on page 523](#)

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Sample programs for distributed platforms

This topic describes the sample programs delivered with IBM WebSphere MQ, written in C and COBOL. The samples demonstrate typical uses of the Message Queue Interface (MQI).

The samples are not intended to demonstrate general programming techniques, so some error checking that you might want to include in a production program is omitted. However, these samples are suitable for use as a base for your own message queuing programs.

The source code for all the samples is provided with the product; this source includes comments that explain the message queuing techniques demonstrated in the programs.

C++ sample programs: See [Using C++](#) for a description of the sample programs available in C++.

The names of the samples start with the prefix amq. The fourth character indicates the programming language, and the compiler where necessary.

s	C language
0	COBOL language on both IBM and Micro Focus compilers
i	COBOL language on IBM compilers only
m	COBOL language on Micro Focus compilers only

The eighth character of the executable indicates whether the sample runs in local binding mode or client mode. If there is no eighth character, then the sample runs in local bindings mode. If the eighth character is 'c' then the sample runs in client mode. To set up the queue manager to accept client connections, see [“Preparing and running the sample programs” on page 104](#) for details.

Use the following links to find out more about the sample programs:

- [“Features demonstrated in the sample programs” on page 93](#)
- [“The Publish/Subscribe sample programs” on page 129](#)
- [“The Put sample programs” on page 134](#)
- [“The Distribution List sample program” on page 122](#)
- [“The Browse sample programs” on page 110](#)

- [“The Browser sample program” on page 111](#)
- [“The Get sample programs” on page 123](#)
- [“The Reference Message sample programs” on page 134](#)
- [“The Request sample programs” on page 140](#)
- [“The Inquire sample programs” on page 128](#)
- [“The Inquire Properties of a Message Handle sample program” on page 129](#)
- [“The Set sample programs” on page 144](#)
- [“The Echo sample programs” on page 122](#)
- [“The Data-Conversion sample program” on page 114](#)
- [“The Triggering sample programs” on page 148](#)
- [“The Asynchronous Put sample program” on page 109](#)
- [“Database coordination samples” on page 114](#)
- [“The CICS transaction sample” on page 112](#)
- [“TUXEDO samples” on page 149](#)
- [“Dead-letter queue handler sample” on page 121](#)
- [“The Connect sample program” on page 112](#)
- [“The API exit sample program” on page 108](#)
- [“Using the SSPI security exit on Windows systems” on page 162](#)
- [“Running the samples using remote queues” on page 163](#)
- [“The Cluster Queue Monitoring sample program \(AMQSCLM\)” on page 163](#)
- [“Sample program for Connection Endpoint Lookup \(CEPL\)” on page 172](#)

Features demonstrated in the sample programs

A collection of tables that show the techniques demonstrated by the WebSphere MQ sample programs.

All the samples open and close queues using the MQOPEN and MQCLOSE calls, so these techniques are not listed separately in the tables. See the heading that includes the platform that you are interested in.

Samples for UNIX and Linux systems

This topic shows the techniques demonstrated by the sample programs for WebSphere MQ on UNIX and Linux systems.

See [“Preparing and running sample programs on UNIX systems” on page 106](#) to find out where the sample programs for WebSphere MQ on UNIX and Linux systems are stored.

[Table 14 on page 93](#) The table lists which C and COBOL source files are provided, and whether a server or client executable is included.

<i>Table 14. WebSphere MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL)</i>				
Technique	C (source) (“1” on page 96)	COBOL (source) (“2” on page 96)	Server (C executable)	Client (C executable) (“3” on page 96)
Using the publish/subscribe interface	amqspuba amqssuba amqssbxa	no sample	amqspub amqssub amqssbx	no sample
Putting messages using the MQPUT call	amqsput0	amq0put0	amqsput	amqsputc

Table 14. WebSphere MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source) (“1” on page 96)	COBOL (source) (“2” on page 96)	Server (C executable)	Client (C executable) (“3” on page 96)
Putting a single message using the MQPUT1 call	amqsinqa amqsecha	amqminqx amqmechx amqiinqx amqiechx	amqsinq amqsech	amqsechc
Putting messages to a distribution list (“4” on page 96)	amqsptl0	amq0ptl0.cbl	amqsptl	amqsptlc
Replying to a request message	amqsinqa	amqminqx amqiinqx	amqsinq	no sample
Getting messages (no wait)	amqsgbr0	amq0gbr0	amqsgbr	no sample
Getting messages (wait with a time limit)	amqsget0	amq0get0	amqsget	amqsgetc
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc
Getting messages (with data conversion)	amqsecha	no sample	amqsech	no sample
Putting Reference Messages to a queue (“4” on page 96)	amqsprma	no sample	amqsprm	amqsprmc
Getting Reference Messages from a queue (“4” on page 96)	amqsgrma	no sample	amqsgrm	amqsgrmc
Reference Message channel exit (“4” on page 96)	amqsqrma amqsxrma	no sample	amqsxrm	no sample
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Browsing complete messages	amqsbcg0	no sample	amqsbcg	amqsbcg
Using a shared input queue	amqsinqa	amqminqx amqiinqx	amqsinq	amqsinqc
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminqx amqiinqx	amqsinq	no sample
Using the MQSET call	amqsseta	amqmsetx amqisetx	amqsset	amqssetc
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	amqsreqc
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	no sample
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	no sample
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	no sample
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	(“5” on page 96)	no sample	no sample	no sample

Table 14. WebSphere MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source) (“1” on page 96)	COBOL (source) (“2” on page 96)	Server (C executable)	Client (C executable) (“3” on page 96)
WebSphere MQ (coordinating XA-compliant database managers) accessing a single database using SQL	amqsxas0.sqc DB2 amqsxas0.ec Informix	amq0xas0.sqb	no sample	no sample
WebSphere MQ (coordinating XA-compliant database managers) accessing two databases using SQL	amqsxag0.c amqsxab0.sqc amqsxaf0.sqc	amq0xag0.cbl amq0xab0.sqb amq0xaf0.sqb	no sample	no sample
CICS transaction (“6” on page 96)	amqscic0.ccs	no sample	amqscic0	no sample
Encina transaction (“4” on page 96)	amqsxae0	no sample	amqsxae0	no sample
TUXEDO transaction to put messages (“7” on page 96)	amqstxpx	no sample	no sample	no sample
TUXEDO transaction to get messages (“7” on page 96)	amqstxgx	no sample	no sample	no sample
Server for TUXEDO (“7” on page 96)	amqstxsx	no sample	no sample	no sample
Dead-letter queue handler	Directory ./tools/c/Samples/dlq (“8” on page 96)	no sample	amqsdldq	no sample
From an MQI client, putting a message	no sample	no sample	no sample	amqsputc
From an MQI client, getting a message	no sample	no sample	no sample	amqsgetc
Connecting to the queue manager using MQCONN	amqscnxc	no sample	no sample	amqscnxc
Using API exits	amqsaxe0	no sample	amqsaxe	no sample
Cluster workload balancing exit	amqswlm0	no sample	amqswlm	no sample
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0	no sample	amqsapt	amqsaptc
Reconnectable clients	amqsphac amqsghac amqsmhac	no sample	not applicable	amqsphac amqsghac amqsmhac
Using message consumers to asynchronously consume messages from multiple queues	amqscbf0	no sample	amqscbf	amqscbfc
Specifying SSL/TLS connection information on MQCONN	amqssslc	no sample	not applicable	amqssslc

Table 14. WebSphere MQ on UNIX and Linux sample programs demonstrating use of the MQI (C and COBOL) (continued)

Technique	C (source) ("1" on page 96)	COBOL (source) ("2" on page 96)	Server (C executable)	Client (C executable) ("3" on page 96)
-----------	-----------------------------	---------------------------------	-----------------------	----------------------------------------

Notes:

1. The executable version of the WebSphere MQ MQI client samples share the same source as the samples that run in a server environment.
2. Compile programs beginning 'amqm' with the Micro Focus COBOL compiler, those beginning 'amqi' with the IBM COBOL compiler, and those beginning 'amq0' with either.
3. The executable versions of the WebSphere MQ MQI client samples are not available on WebSphere MQ for HP-UX.
4. Supported on WebSphere MQ for AIX, WebSphere MQ for HP-UX, and WebSphere MQ for Solaris only.
5. On WebSphere MQ for AIX, WebSphere MQ for HP-UX, and WebSphere MQ for Solaris this program is called amqsvfc0.c
6. CICS is supported by WebSphere MQ for AIX and WebSphere MQ for HP-UX only.
7. TUXEDO is not supported by WebSphere MQ for Linux on System p.
8. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.

Detailed information about support for UNIX and Linux systems is available at the WebSphere MQ systems requirements page at [System Requirements for IBM WebSphere MQ](#).

Samples for IBM WebSphere MQ client for HP Integrity NonStop Server

This topic shows the techniques demonstrated by the sample programs for the IBM WebSphere MQ client on HP Integrity NonStop Server systems.

Table 15 on page 96The table lists the C, COBOL, and pTAL source sample programs that are provided.

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL								
Technique	C				COBOL		pTAL	
	OSS (Source)	OSS (Executable)	Guardian (Source)	Guardian (Executable)	OSS (Source)	Guardian (Source)	OSS (Source)	Guardian (Source)
Using the publish/subscribe interface	amqspub a.c amqssbxa .c amqssuba .c amqspse 0.c	amqspub c amqssbxc amqssubc	MQSPUBC MQSSBXC MQSSUBC	AMQSPU BC AMQSSBX C AMQSSUB C	amq0pu b0.cbl amq0su b0.cbl	MQSPUBL MQSSUBL	amqtpub0 .tal amqtsub0 .tal	MQSPUBT MQSSUBT
Putting messages using the MQPUT call	amqsput0 .c	amqsputc	MQSPUTC	AMQSPUT C	amq0put 0.cbl	MQSPUTL	amqtput0 .tal	MQSPUTT

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Putting a single message using the MQPUT1 call	amqsecha.c	amqsechc	MQSECHC	AMQSECHC			amqtech0.tal	MQSECHT
Putting messages to a distribution list	amqsptl0.c	amqsptlc	MQSPTLC	AMQSPTLC	amq0ptl0.cbl	MQSPTLL		
Replying to a request message	amqsinqa.c	amqsinqc	MQSINQC	AMQSINQC				
Getting messages (no wait)	amqsgbr0.c	amqsgbrc	MQSGBRC	AMQSGBRC	amq0gbr0.cbl	MQSGBRL		
Getting messages (wait with a time limit)	amqsget0.c	amqsgetc	MQSGETC	AMQSGETC	amq0get0.cbl	MQSGETL	amqtget0.tal	MQSGETT
Getting messages (unlimited wait)	amqstrg0.c	amqstrgc	MQSTRGC	AMQSTRGC				
Getting messages (with data conversion)	amqsecha.c	amqsechc	MQSECHC	AMQSECHC				
Putting Reference Messages to a queue	amqsprm.a.c	amqsprm.c	MQSPRMC	AMQSPRMC				

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Getting Reference Messages from a queue	amqsgm a.c	amqsgm c	MQSGRM C	AMQSGR MC				
Reference Message channel exit	amqsqrm a.c amqsxrm a.c		MQSQRM C MQSXRM C					
Browsing first 20 characters of a message	amqsgbr0 .c	amqsgbrc	MQSGBR C	AMQSGB RC	amq0gbr 0.cbl	MQSGBRL		
Browsing complete messages	amqsbcg 0.c	amqsbcgc	MQSBCG C	AMQSBC GC				
Using a shared input queue	amqsinqa .c	amqsinqc	MQSINQC	MQSINQC				
Using an exclusive input queue	amqstrg0. c	amqstrgc	MQSTRGC	AMQSTRG C				
Using the MQINQ call	amqsinqa .c	amqsinqc	MQSINQC	AMQSINQ C				
Using the MQSET call	amqsseta. c	amqssetc	MQSSETC	AMQSSET C				
Using a reply-to queue	amqsreq0 .c	amqsreqc	MQSREQC	AMQSREQ C	amq0req 0.cbl	MQSREQL		
Requesting message exceptions	amqsreq0 .c	amqsreqc	MQSREQC	AMQSREQ C	amq0req 0.cbl	MQSREQL		

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Accepting a truncated message	amqsgbr0.c	amqsgbrC	MQSGBRC	AMQSGBR	amq0gbr0.cbl	MQSGBRL		
Using a resolved queue name	amqsgbr0.c	amqsgbrC	MQSGBRC	AMQSGBR	amq0gbr0.cbl	MQSGBRL		
Triggering a process	amqstrg0.c	amqstrgC	MQSTRGC	AMQSTRGC				
Using data conversion	amqsvfc0.c							
Dead letter queue handler (1)	Directory . /samp/dlq							
Connecting to a queue manager using MQCONN	amqscnxc.c	amqscnxc	MQSCNXC					
Using API exits	amqsaxe0.c amqsaem0.c							
Cluster workload balancing exit	amqswlm0.c		MQSWLMC					
Cluster queue monitor	amqsclma.c							

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0.c	amqsaptc	MQSAPTC	MQSAPTC				
Reconnectable clients	amqsg hac .c amqsm hac .c amqsphac .c	amqsg hac amqsm hac amqsphac	MQSGHAC MQSMHAC MQSPHAC MQSFHAC	AMQSGHAC AMQSMHAC AMQSPHAC AMQSFHAC				
Using the message consumers to asynchronously consume messages from multiple queues	amqscbf0.c	amqscbfc						
Specifying SSL/TLS connection information on MQCONNX	amqssslc.c	amqssslc	MQSSSLC	AMQSSSLC				
Activity trace	amqsact0.c	amqsactc	MQSACTC	AMQSACTC				
Message properties	amqsiqm a.c amqsstm a.c	amqsiqm c amqsstm c	MQSIQMC MQSSTMC	AMQSIQMC AMQSSTMC				
Command server	amqsstop.c		MQSSTOC					

Table 15. IBM WebSphere MQ on HP Integrity NonStop Server sample programs demonstrating use of C, COBOL, and pTAL (continued)

Technique	C				COBOL		pTAL	
Log events	amqslog0.c	amqslogc	MQSLOGC	AMQSLOGC				
Accounting	amqsmon0.c	amqsmonc	MQSMONC	AMQSMONC				
Administration interface	amqsaicq.c amqsaie m.c amqsailq.c							
An example of a C language main function for invoking pTAL			MQSPTMC					

Notes:

1. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.
2. For information about developing applications for your IBM WebSphere MQ client on the HP Integrity NonStop Server platform, see:
 - [“Building your application on HP Integrity NonStop Server” on page 416](#)
 - [“Preparing C programs in HP Integrity NonStop Server” on page 418](#)
 - [“Preparing COBOL programs” on page 419](#)
 - [“Preparing pTAL programs” on page 420](#)

Samples for IBM WebSphere MQ for Windows

This topic shows the techniques demonstrated by the sample programs for IBM WebSphere MQ for Windows.

Table 16 on page 101The table lists which C and COBOL source files are provided, and whether a server or client executable is included.

<i>Table 16. IBM WebSphere MQ for Windows sample programs demonstrating use of the MQI (C and COBOL)</i>				
Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
Using the publish/subscribe interface	amqspuba amqssuba amqssbxa	no sample	amqspub amqssub amqssbx	no sample
Putting messages using the MQPUT call	amqsput0	amq0put0	amqsput	amqsputc

Table 16. IBM WebSphere MQ for Windows sample programs demonstrating use of the MQI (C and COBOL)
(continued)

Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
Putting a single message using the MQPUT1 call	amqsinqa amqsecha	amqminq2 amqmech2 amqiinq2 amqiech2	amqsinq amqsech	amqsinqc amqsechc
Putting messages to a distribution list	amqsptl0	amq0ptl0.cbl	amqsptl	amqsptlc
Replying to a request message	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Getting messages (no wait)	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Getting messages (wait with a time limit)	amqsget0	amq0get0	amqsget	amqsgetc
Getting messages (unlimited wait)	amqstrg0	no sample	amqstrg	amqstrgc
Getting messages (with data conversion)	amqsecha	no sample	amqsech	amqsechc
Putting Reference Messages to a queue	amqsprma	no sample	amqsprm	amqsprmc
Getting Reference Messages from a queue	amqsgrma	no sample	amqsgrm	amqsgrmc
Reference Message channel exit	amqsqrma amqsxrma	no sample	amqsxrm	no sample
Browsing first 20 characters of a message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Browsing complete messages	amqsbcg0	no sample	amqsbcg	amqsbcg
Using a shared input queue	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Using an exclusive input queue	amqstrg0	amq0req0	amqstrg	amqstrgc
Using the MQINQ call	amqsinqa	amqminq2 amqiinq2	amqsinq	amqsinqc
Using the MQSET call	amqsseta	amqmset2 amqiset2	amqsset	amqssetc
Using the MQINQMP call	amqsiqma	no sample	no sample	no sample
Using a reply-to queue	amqsreq0	amq0req0	amqsreq	amqsreqc
Requesting message exceptions	amqsreq0	amq0req0	amqsreq	amqsreqc
Accepting a truncated message	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Using a resolved queue name	amqsgbr0	amq0gbr0	amqsgbr	amqsgbrc
Triggering a process	amqstrg0	no sample	amqstrg	amqstrgc
Using data conversion	amqsvfc0	no sample	no sample	no sample
WebSphere MQ (coordinating XA-compliant database managers) accessing a single database using SQL	amqsxas0.sqc DB2 amqsxas0.ec Informix	amq0xas0.sq b	no sample	no sample

Table 16. IBM WebSphere MQ for Windows sample programs demonstrating use of the MQI (C and COBOL)
(continued)

Technique	C (source)	COBOL (source)	Server (C executable)	Client (C executable)
WebSphere MQ (coordinating XA-compliant database managers) accessing two databases using SQL	amqsxag0.c amqsxab0.sq c DB2 amqsxaf0.sqc DB2	amq0xag0.cbl amq0xab0.sq b amq0xaf0.sqb	no sample	no sample
TUXEDO transaction to put messages	amqstxpx	no sample	no sample	no sample
TUXEDO transaction to get messages	amqstxgx	no sample	no sample	no sample
Server for TUXEDO	amqstxsx	no sample	no sample	no sample
Dead-letter queue handler	Directory ./tools/c/Samples/dlq ("1" on page 103)	no sample	amqsdlq	no sample
From a WebSphere MQ MQI client, putting a message	no sample	no sample	no sample	amqsputc
From a WebSphere MQ MQI client, getting a message	no sample	no sample	no sample	amqsgetc
Connecting to the queue manager using MQCONN	amqscnxc	no sample	no sample	amqscnxc
Using API exits	amqsaxe0	no sample	amqsaxe	no sample
Cluster workload balancing	amqswlm0	no sample	amqswlm	no sample
SSPI security routines	amqsspin	no sample	amqrsplin.dll	amqrsplin.dll
Putting messages asynchronously and getting status using the MQSTAT call	amqsapt0	no sample	amqsapt	amqsaptc
Reconnectable clients	amqsphac amqsghac amqsmhac	no sample	Not applicable	amqsphac amqsghac amqsmhac
Using message consumers to asynchronously consume messages from multiple queues	amqscbf0	no sample	amqscbf	amqscbfc
Specifying SSL/TLS connection information on MQCONN	amqssslc	no sample	not applicable	amqssslc
Notes: 1. The source for the dead-letter queue handler consists of several files and is provided in a separate directory.				

Visual Basic samples for IBM WebSphere MQ for Windows

This topic shows the techniques demonstrated by the Visual Basic sample programs for IBM WebSphere MQ for Windows.

Table 17 on page 104 shows the techniques demonstrated by the IBM WebSphere MQ for Windows sample programs.

A project can contain several files. When you open a project within Visual Basic, the other files are loaded automatically. No executable programs are provided.

All the sample projects, except mqtrivc.vbp, are set up to work with the IBM WebSphere MQ server. To find out how to change the sample projects to work with the IBM WebSphere MQ clients see [“Preparing Visual Basic programs in Windows”](#) on page 443.

Table 17. IBM WebSphere MQ for Windows sample programs demonstrating use of the MQI (Visual Basic)	
Technique	Project file name
Putting messages using the MQPUT call	amqsputb.vbp
Getting messages using the MQGET call	amqsgetb.vbp
Browsing a queue using the MQGET call	amqsbcgb.vbp
Simple MQGET and MQPUT sample (client)	mqtrivc.vbp
Simple MQGET and MQPUT sample (server)	mqtrivs.vbp
Putting and getting strings and user-defined structures using MQPUT and MQGET	strings.vbp
Using PCF structures to start and stop a channel	pcfsamp.vbp
Creating a queue using the MQAI	amqsaicq.vbp
Listing a queue manager's queues using the MQAI	amqsailq.vbp
Monitoring events using the MQAI	amqsaiem.vbp

Preparing and running the sample programs

Configure your queue manager to securely accept incoming connection requests from applications running in client mode.

Before you begin

Ensure the queue manager already exists and has been started. Determine whether channel authentication records are already enabled as follows:

```
DISPLAY QMGR CHLAUTH
```

This task expects that channel authentication records are enabled. If this is a queue manager used by other users and applications, changing this setting will affect all other users and applications. If your queue manager does not make use of channel authentication records then step [“4” on page 105](#) can be replaced with an alternate authentication method (for example a security exit) which sets the MCAUSER to the *non-privileged-user-id* you will obtain in step [“1” on page 105](#).

You must know which channel name your application expects to use so that the application can be permitted to use the channel. You must also know which objects, for example queues or topics, your application expects to use so that your application can be permitted to use them.

About this task

This task creates a non-privileged user ID to be used for a client application which connects to the queue manager. Access is granted for the client application only to be able to use the channel it needs and the queue it needs by use of this user ID.

Procedure

1. Obtain a user ID on the system your queue manager is running on. For this task this user ID must not be a privileged administrative user. This user ID will be the authority under which the client connection will run on the queue manager.

2. Start a listener program with the following commands where:

qmgr is the name of your queue manager

nnnn is your chosen port number

- a) For UNIX and Windows systems:

```
runmqclsr -t tcp -m qmgr -p nnnn
```

3. If your application uses the SYSTEM.DEF.SVRCONN then this channel is already defined. If your application uses another channel, create it by issuing the MQSC command:

```
DEFINE CHANNEL('channel-name') CHLTYPE(SVRCONN) TRPTYPE(TCP) +  
DESCR('Channel for use by sample programs')
```

channel-name is the name of your channel.

4. Create a channel authentication rule allowing only the IP address of your client system to use the channel by issuing the MQSC command:

```
SET CHLAUTH('channel-name') TYPE(ADDRESSMAP) ADDRESS('client-machine-IP-address') +  
MCAUSER('non-privileged-user-id')
```

channel-name is the name of your channel.

client-machine-IP-address is the IP address of your client system.

If your sample client application is running on the same machine as the queue manager then use an IP address of '127.0.0.1' if your application is going to connect using 'localhost'. If several different client machines are going to connect in, you can use a pattern or a range instead of a single IP address. See [Generic IP addresses](#) for details.

non-privileged-user-id is the user ID you obtained in step “1” on [page 105](#)

5. If your application uses the SYSTEM.DEFAULT.LOCAL.QUEUE then this queue is already defined. If your application uses another queue, create it by issuing the MQSC command:

```
DEFINE QLOCAL('queue-name') DESCR('Queue for use by sample programs')
```

queue-name is the name of your queue.

6. Grant access to connect to and inquire the queue manager:

- a) For UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC OBJTYPE(QMGR) PRINCIPAL('non-privileged-user-id') +  
AUTHADD(CONNECT, INQ)
```

non-privileged-user-id is the user ID you obtained in step “1” on [page 105](#)

7. If your application is a point-to-point application, that is it makes use of queues, grant access to allow inquiring and the putting and getting messages using your queue by the user ID to be used, by issuing the MQSC commands:

- a) For UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC PROFILE('queue-name') OBJTYPE(QUEUE) +  
PRINCIPAL('non-privileged-user-id') AUTHADD(PUT, GET, INQ, BROWSE)
```

queue-name is the name of your queue.

non-privileged-user-id is the user ID you obtained in step “1” on [page 105](#)

8. If your application is a publish/subscribe application, that is it makes use of topics, grant access to allow publishing and subscribing using your topic by the user ID to be used, by issuing the MQSC commands:

a) For UNIX and Windows systems issue the MQSC commands:

```
SET AUTHREC PROFILE('SYSTEM.BASE.TOPIC') OBJTYPE(TOPIC) +  
PRINCIPAL('non-privileged-user-id') AUTHADD(PUB, SUB)
```

non-privileged-user-id is the user ID you obtained in step “1” on page 105

This will give *non-privileged-user-id* access to any topic in the topic tree, alternatively, you can define a topic object using **DEFINE TOPIC** and grant accesses only to the part of the topic tree referenced by that topic object. See [Controlling user access to topics](#) for details.

What to do next

Your client application can now connect to the queue manager and put or get messages using the queue.

Related tasks

[Giving access to a WebSphere MQ object on UNIX or Linux systems and Windows](#)

Related reference

[SET CHLAUTH](#)

[DEFINE CHANNEL](#)

[DEFINE QLOCAL](#)

[SET AUTHREC](#)

Preparing and running sample programs on UNIX systems

Table 18. Where to find the samples for WebSphere MQ on UNIX and Linux systems	
Content	Directory
source files	<i>MQ_INSTALLATION_PATH</i> /samp
dead-letter queue handler source files	<i>MQ_INSTALLATION_PATH</i> /samp/dlq
executable files	<i>MQ_INSTALLATION_PATH</i> /samp/bin
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which WebSphere MQ is installed.	

The WebSphere MQ on UNIX and Linux systems sample files are in the directories listed in [Table 18](#) on page 106 if the defaults were used at installation time. To run the samples, either use the executable versions supplied or compile the source versions as you would any other applications, using an ANSI compiler. For information about how to do so, see [“Running the sample programs”](#) on page 107.

Preparing and running sample programs on Windows systems

Table 19. Where to find the samples for WebSphere MQ for Windows	
Content	Directory
C source code	<i>MQ_INSTALLATION_PATH</i> \Tools\C\Samples
Source code for dead-letter handler sample	<i>MQ_INSTALLATION_PATH</i> \Tools\C\Samples\DLQ
COBOL source code	<i>MQ_INSTALLATION_PATH</i> \Tools\Cobol\Samples
C executable files ¹	<i>MQ_INSTALLATION_PATH</i> \Tools\C\Samples\Bin (32-bit versions) <i>MQ_INSTALLATION_PATH</i> \Tools\C\Samples\Bin64 (64-bit versions)
Sample MQSC files	<i>MQ_INSTALLATION_PATH</i> \Tools\MQSC\Samples
Visual Basic source code	<i>MQ_INSTALLATION_PATH</i> \Tools\VB\SampVB6
.NET samples	<i>MQ_INSTALLATION_PATH</i> \Tools\dotnet\Samples

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Note:

1. 64-bit versions are available of some C executable file samples.

The WebSphere MQ for Windows sample files are in the directories listed in [Table 19 on page 106](#) if the defaults were used at installation time; the installation drive defaults to <c:>. To run the samples, either use the executable versions supplied or compile the source versions as you would any other WebSphere MQ for Windows applications. For information on how to do this, see [“Running the sample programs” on page 107](#).

Running the sample programs

Consider using this topic when running sample programs across different platforms.

Before you can run any of the sample programs, create a queue manager and set up the default definitions. This is explained in [Administering](#).

On Windows, UNIX and Linux platforms

The samples need a set of queues to work with. Either use your own queues or run the sample MQSC file `amqscos0.tst` to create a set.

To do this on UNIX and Linux systems, enter:

- `runmqsc QManagerName <amqscos0.tst >/tmp/sampobj.out`

Check the `sampobj.out` file to ensure that there are no errors.

To do this on Windows systems enter:

- `runmqsc QManagerName <amqscos0.tst > sampobj.out`

Check the `sampobj.out` file to ensure that there are no errors. This file is in your current directory.

You can now run the sample applications. Enter the name of the sample application followed by any parameters, for example:

- `amqspout myqueue qmanagername`

where `myqueue` is the name of the queue on which the messages are going to be put, and `qmanagername` is the queue manager that owns `myqueue`.

See the description of the individual samples for information on the parameters that each of them expects.

Length of queue name

For the COBOL sample programs, when you pass queue names as parameters, you must provide 48 characters, padding with blank characters if necessary. Anything other than 48 characters causes the program to fail with reason code 2085.

Inquire, Set, and Echo examples

For the Inquire, Set, and Echo examples, the sample definitions trigger the C versions of these samples.

If you want the COBOL versions you must change the process definitions:

- `SYSTEM.SAMPLE.INQPROCESS`
- `SYSTEM.SAMPLE.SETPROCESS`
- `SYSTEM.SAMPLE.ECHOPROCESS`

On Windows, UNIX and Linux systems do this by editing the `amqscos0.tst` file and changing the C executable file names to the COBOL executable file names before using the **runmqsc** command, as shown previously.

The API exit sample program

The sample API exit generates an MQI trace to a user-specified file with a prefix defined in the MQAPI_TRACE_LOGFILE environment variable.

For more information about API exits, see [“Writing and compiling API exits” on page 370](#).

Source

amqsaxe0.c

Binary

amqsaxe

Configuring for the sample exit

1. Add the following to the qm.ini file.

Platforms other than Windows

```
ApiExitLocal:
  Sequence=100
  Function=EntryPoint
  Module=MQ_INSTALLATION_PATH/samp/bin/amqsaxe
  Name=SampleApiExit
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM WebSphere MQ is installed.

Windows

```
ApiExitLocal:
  Sequence=100
  Function=EntryPoint
  Module=MQ_INSTALLATION_PATH\Tools\c\Samples\bin\amqsaxe
  Name=SampleApiExit
```

where *MQ_INSTALLATION_PATH* represents the directory where IBM WebSphere MQ is installed.

2. Set the environment variable

```
MQAPI_TRACE_LOGFILE=/tmp/MqiTrace
```

3. Run your application.

Output files are created in the /tmp directory with names like: MqiTrace.<pid>.<tid>.log

The Asynchronous consumption sample program

The amqscbf sample program demonstrates the use of MQCB and MQCTL to consume messages from multiple queues asynchronously.

amqscbf is provided as C source code, and a binary client and server executable on Windows, UNIX and Linux platforms.

The program is started from the command line and takes the following optional parameters:

```
Usage: [Options] <Queue Name> { <Queue Name> }
  where Options are:
    -m <Queue Manager Name>
    -o <Open options>
    -r <Reconnect Type>
      d Reconnect Disabled
      r Reconnect
      m Reconnect Queue Manager
```

Provide more than one queue name to read messages from multiple queues (a maximum of ten queues are supported by the sample.)

Note: *Reconnect type* is only valid for client programs.

Example

The example shows amqscbf run as a server program reading one message from QL1 and then being stopped.

Use WebSphere MQ Explorer to put a test message on QL1. Stop the program by pressing enter.

```
C:\>amqscbf QL1
Sample AMQSCBF0 start

Press enter to end
Message Call (9 Bytes) :
Message 1

Sample AMQSCBF0 end
```

What amqscbf demonstrates

The sample shows how to read messages from multiple queues in the order of their arrival. This would require a lot more code using synchronous MQGET. In the case of asynchronous consumption, no polling is required, and thread and storage management is performed by WebSphere MQ. A "real world" example would need to deal with errors; in the sample errors are written out to the console.

The sample code has the following steps,

1. Define the single message consumption callback function,

```
void MessageConsumer(MQHCONN hConn,
                     MQMD * pMsgDesc,
                     MQGMO * pGetMsgOpts,
                     MQBYTE * Buffer,
                     MQCBC * pContext)
{ ... }
```

2. Connect to the queue manager,

```
MQCONN(XQMName, &cno, &Hcon, &CompCode, &CReason);
```

3. Open the input queues, and associate each one with the MessageConsumer callback function,

```
MQOPEN(Hcon, &od, 0_options, &Hobj, &OpenCode, &Reason);
cbd.CallbackFunction = MessageConsumer;
MQCB(Hcon, MQOP_REGISTER, &cbd, Hobj, &md, &gmo, &CompCode, &Reason);
```

cbd.CallbackFunction does not need to be set for each queue; it is an input-only field. But you could associate a different callback function with each queue.

4. Start consumption of the messages,

```
MQCTL(Hcon, MQOP_START, &ctlo, &CompCode, &Reason);
```

5. Wait until the user has pressed enter and then stop consumption of messages,

```
MQCTL(Hcon, MQOP_STOP, &ctlo, &CompCode, &Reason);
```

6. Finally disconnect from the queue manager,

```
MQDISC(&Hcon, &CompCode, &Reason);
```

The Asynchronous Put sample program

Learn about running the amqsapt sample and the design of the Asynchronous Put sample program.

The asynchronous put sample program puts messages on a queue using the asynchronous MQPUT call and then retrieves status information using the MQSTAT call. See [“Features demonstrated in the sample programs” on page 93](#) for the name of this program on different platforms.

Running the amqsapt sample

This program takes up to 6 parameters:

1. The name of the target queue (required)
2. The name of the queue manager (optional)
3. Open options (optional)
4. Close options (optional)
5. The name of the target queue manager (optional)
6. The name of the dynamic queue (optional)

If a queue manager is not specified, amqsapt connects to the default queue manager.

Design of the Asynchronous Put sample program

The program uses the MQOPEN call with the output options supplied, or with the MQOO_OUTPUT and MQOO_FAIL_IF_QUIESCING options to open the target queue for putting messages.

If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple, on this and on subsequent MQI calls, the program uses default values for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call with MQPMO_ASYNC_RESPONSE to create a datagram message containing the text of that line and asynchronously put it to the target queue. The program continues until it reaches the end of the input or the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

The program then issues the MQSTAT call, returning an MQSTS structure, and displays messages containing the number of messages put successfully, the number of messages put with a warning, and the number of failures.

The Browse sample programs

The Browse sample programs browse messages on a queue using the MQGET call.

See [“Features demonstrated in the sample programs”](#) on page 93 for the names of these programs.

Design of the Browse sample program

The program opens the target queue using the MQOPEN call with the MQOO_BROWSE option. If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to copy the message from the queue, then displays the data contained in the message. The MQGET call uses these options:

MQGMO_BROWSE_NEXT

After the MQOPEN call, the browse cursor is positioned logically before the first message in the queue, so this option causes the *first* message to be returned when the call is first made.

MQGMO_NO_WAIT

The program does not wait if there are no messages on the queue.

MQGMO_ACCEPT_TRUNCATED_MSG

The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the program displays the truncated message, together with a warning that the message has been truncated.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call, because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues to the end of the queue; the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code and the program displays a warning message. If the MQGET call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

UNIX, Linux and Windows systems

Consider using this topic when learning about Browse sample programs on UNIX, Linux and Windows systems.

The C version of the program takes 2 parameters

1. The name of the source queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it connects to the default one. For example, enter one of the following:

- amqsgbr myqueue qmanagername
- amqsgbrc myqueue qmanagername
- amq0gbr0 myqueue

where myqueue is the name of the queue that the messages will be viewed from, and qmanagername is the queue manager that owns myqueue.

If you omit the qmanagername, when running the C sample, it assumes that the default queue manager owns the queue.

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target queue
```

Only the first 50 characters of each message are displayed, followed by - - - truncated when this is the case.

The Browser sample program

The Browser sample program reads and writes both the message descriptor and the message content fields of all the messages on a queue.

The sample program is written as a utility, not just to demonstrate a technique. See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

This program takes these parameters:

1. The name of the source queue
2. The name of the queue manager
3. An optional parameter for properties.

The first two input parameters for this program are mandatory. For example, start the program in one of the following ways:

- amqsbcg myqueue qmanagername
- amqsbcgc myqueue qmanagername

where myqueue is the name of the queue on which the messages are going to be browsed, and qmanagername is the queue manager that owns myqueue.

It reads each message from the queue and writes the following to stdout:

- Formatted message descriptor fields
- Message data (dumped in hex and, where possible, character format)

Permissible values for the property parameter are:

Value	Behavior
0	Default behaviour, as it was for V6. The properties that get delivered to the application depend on the <i>PropertyControl</i> queue attribute that the message is retrieved from.
1	<p>A message handle is created and used with the MQGET. Properties of the message, except those contained in the message descriptor (or extension) are displayed in a similar fashion to the message descriptor. For example:</p> <pre>****Message properties**** <property name> : <property value></pre> <p>Or if no properties are available:</p> <pre>****Message properties**** None</pre> <p>Numeric values are displayed using printf, string values are surrounding in single quotation marks, and byte strings are surrounded with X and single quotation marks, as for the message descriptor.</p>
2	MQGMO_NO_PROPERTIES is specified, so that only message descriptor properties will be returned.
3	MQGMO_PROPERTIES_FORCE_MQRFH2 is specified, so that all properties are returned in the message data.
4	MQGMO_PROPERTIES_COMPATIBILITY is specified, so that all properties can be returned depending on whether a version 6 property is included, otherwise the properties are discarded.

The program is restricted to printing the first 65535 characters of the message, and fails with the reason *truncated msg* if a longer message is read.

See the [Administering](#) for an example of the output from this utility.

The CICS transaction sample

A sample CICS transaction program is provided, named amqscic0.ccs for source code and amqscic0 for the executable version. You can build transactions using the standard CICS facilities.

See “[Building an IBM WebSphere MQ application](#)” on page 409 for details on the commands needed for your platform.

The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue, the name of which is contained in the transmission header of the message. Any failures are sent to the queue SYSTEM.SAMPLE.CICS.DLQ.

Note: You can use a sample MQSC script amqscic0.tst to create these queues and sample input queues.

The Connect sample program

The Connect sample program allows you to explore the MQCONN call and its options from a client. The sample connects to the queue manager using the MQCONN call, inquires about the name of the queue manager using the MQINQ call, and displays it. Also, learn about running the amqscnxc sample.

Note: The Connect sample program is a client sample. You can compile and run it on a server but the function is meaningful only on a client, and only client-executable files are supplied.

Running the amqscnxc sample

The command-line syntax of the Connect sample program is:


```
amqscnxc [-x ConnName [-c SvrconnChannelName]] [QMgrName]
```

The parameters are optional and their order is not important except for QMgrName, which, if specified, must come last. The parameters are:

ConnName

The TCP/IP connection name of the server queue manager

SvrconnChannelName

The name of the server connection channel

QMgrName

The name of the target queue manager

If you do not specify the TCP/IP connection name, MQCONN is issued with the *ClientConnPtr* set to NULL. If you specify the TCP/IP connection name but not the server connection channel (the reverse is not allowed), the sample uses the name SYSTEM.DEF.SVRCONN. If you do not specify the target queue manager, the sample connects to whichever queue manager is listening at the given TCP/IP connection name.

Note: If you enter a question mark as the only parameter, or if you enter incorrect parameters, you get a message explaining how to use the program.

If you run the sample with no command-line options, the contents of the MQSERVER environment variable are used to determine the connection information. (In this example MQSERVER is set to SYSTEM.DEF.SVRCONN/TCP/machine.site.company.com.) You see output like this:

```
Sample AMQSCNXC start
Connecting to the default queue manager
with no client connection information specified.
Connection established to queue manager machine

Sample AMQSCNXC end
```

If you run the sample and provide a TCP/IP connection name and a server connection channel name but no target queue manager name, like this:

```
amqscnxc -x machine.site.company.com -c SYSTEM.ADMIN.SVRCONN
```

the default queue manager name is used and you see output like this:

```
Sample AMQSCNXC start
Connecting to the default queue manager
using the server connection channel SYSTEM.ADMIN.SVRCONN
on connection name machine.site.company.com.
Connection established to queue manager MACHINE

Sample AMQSCNXC end
```

If you run the sample and provide a TCP/IP connection name and a target queue manager name, like this:

```
amqscnxc -x machine.site.company.com MACHINE
```

you see output like this:

```
Sample AMQSCNXC start
Connecting to queue manager MACHINE
using the server connection channel SYSTEM.DEF.SVRCONN
on connection name machine.site.company.com.
Connection established to queue manager MACHINE

Sample AMQSCNXC end
```

The Data-Conversion sample program

The data-conversion sample program is a skeleton of a data conversion exit routine. Learn about the design of the data-conversion sample.

See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

Design of the data-conversion sample

Each data-conversion exit routine converts a single named message format. This skeleton is intended as a wrapper for code fragments generated by the data-conversion exit generation utility program.

The utility produces one code fragment for each data structure; several such structures make up a format, so several code fragments are added to this skeleton to produce a routine to do data conversion of the entire format.

The program then checks whether the conversion is a success or failure, and returns the values required to the caller.

Database coordination samples

Two samples are provided that demonstrate how WebSphere MQ can coordinate both WebSphere MQ updates and database updates within the same unit of work.

These samples are:

1. AMQXSAS0 (in C) or AMQ0XAS0 (in COBOL), which updates a single database within a WebSphere MQ unit of work.
2. AMQSXAG0 (in C) or AMQ0XAG0 (in COBOL), AMQSXAB0 (in C) or AMQ0XAB0 (in COBOL), and AMQSXAF0 (in C) or AMQ0XAF0 (in COBOL), which together update two databases within a WebSphere MQ unit of work, showing how multiple databases can be accessed. These samples are provided to show the use of the MQBEGIN call, mixed SQL and WebSphere MQ calls, and where and when to connect to a database.

[Figure 18 on page 115](#) shows how the samples provided are used to update databases:

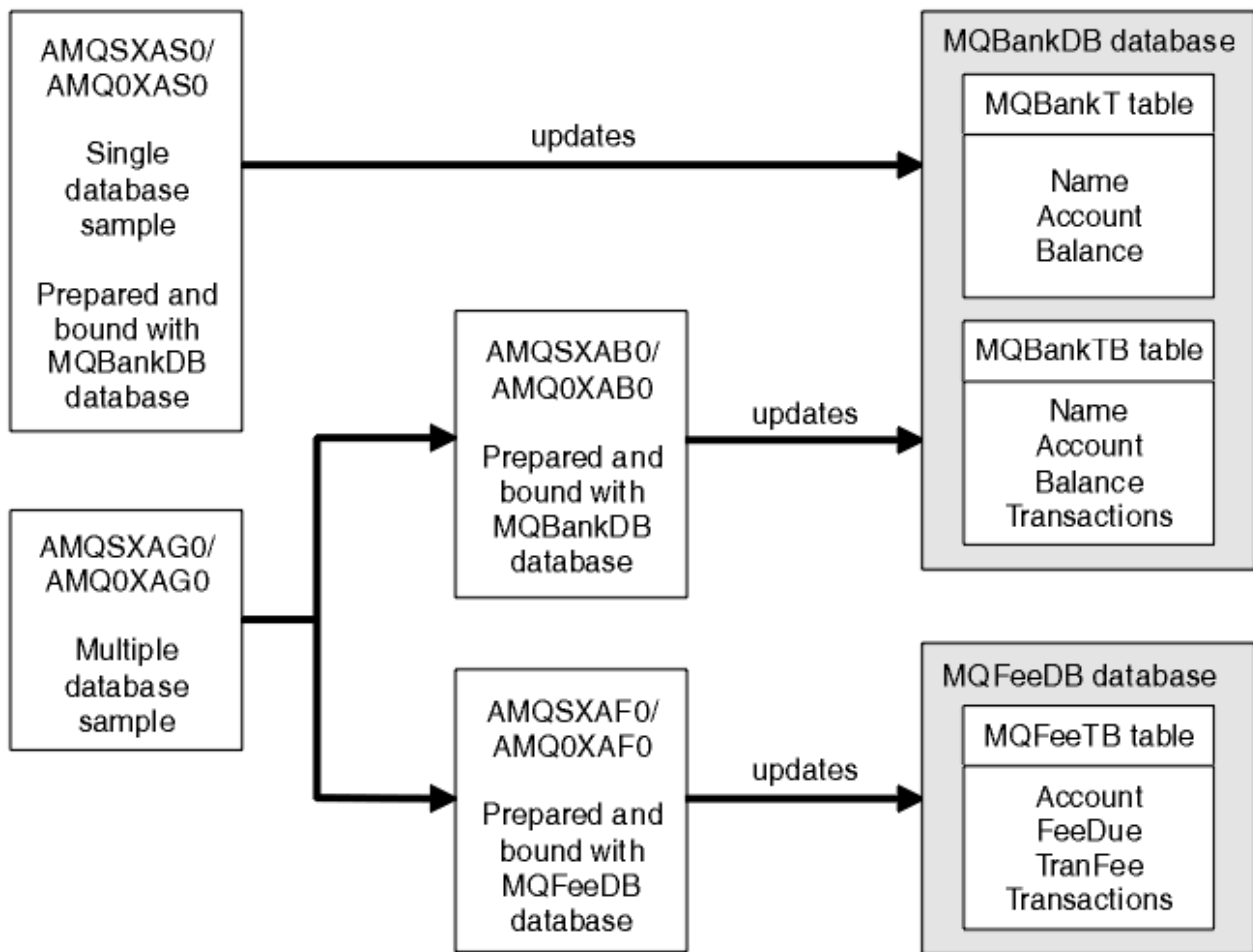


Figure 18. The database coordination samples

The programs read a message from a queue (under syncpoint), then, using the information in the message, obtain the relevant information from the database and update it. The new status of the database is then printed.

The program logic is as follows:

1. Use the name of the input queue from the program argument
2. Connect to the default queue manager (or optionally the supplied name in C) using MQCONN
3. Open a queue (using MQOPEN) for input while there are no failures
4. Start a unit of work using MQBEGIN
5. Get the next message (using MQGET) from the queue under syncpoint
6. Get information from databases
7. Update information from databases
8. Commit changes using MQCMIT
9. Print updated information (no message being available counts as a failure, and the loop ends)
10. Close the queue using MQCLOSE
11. Disconnect from the queue using MQDISC

SQL cursors are used in the samples, so that reads from the databases (that is, multiple instances) are locked while a message is being processed, allowing multiple instances of these programs to run simultaneously. The cursors are explicitly opened, but implicitly closed by the MQCMIT call.

The single database sample (AMQSXAS0 or AMQ0XAS0) has no SQL CONNECT statements and the connection to the database is implicitly made by WebSphere MQ with the MQBEGIN call. The multiple

database sample (AMQ SXAG0 or AMQ0XAG0, AMQ SXAB0 or AMQ0XAB0, and AMQ SXAF0 or AMQ0XAF0) has SQL CONNECT statements, as some database products allow only one active connection. If this is not the case for your database product, or if you are accessing a single database in multiple database products, the SQL CONNECT statements can be removed.

The samples are prepared with the IBM DB2 database product, so you might need to modify them to work with other database products.

The SQL error checking uses routines in UTIL.C and CHECKERR.CBL supplied by DB2. These must be compiled or replaced before compiling and linking.

Note: If you are using the Micro Focus COBOL source CHECKERR.MFC for SQL error checking, you must change the program ID to uppercase, that is CHECKERR, for AMQ0XAS0 to link correctly.

Creating the databases and tables

Create the databases and tables before compiling the samples.

To create the databases, use the usual method for your database product, for example:

```
DB2 CREATE DB MQBankDB
DB2 CREATE DB MQFeeDB
```

Create the tables using SQL statements as follows:

In C:

```
EXEC SQL CREATE TABLE MQBankT(Name          VARCHAR(40) NOT NULL,
                                Account        INTEGER   NOT NULL,
                                Balance         INTEGER   NOT NULL,
                                PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQBankTB(Name          VARCHAR(40) NOT NULL,
                                Account        INTEGER   NOT NULL,
                                Balance         INTEGER   NOT NULL,
                                Transactions     INTEGER,
                                PRIMARY KEY (Account));

EXEC SQL CREATE TABLE MQFeeTB(Account        INTEGER   NOT NULL,
                                FeeDue          INTEGER   NOT NULL,
                                TranFee         INTEGER   NOT NULL,
                                Transactions     INTEGER,
                                PRIMARY KEY (Account));
```

In COBOL:

```
EXEC SQL CREATE TABLE
MQBankT(Name          VARCHAR(40) NOT NULL,
         Account        INTEGER   NOT NULL,
         Balance         INTEGER   NOT NULL,
         PRIMARY KEY (Account))
END-EXEC.

EXEC SQL CREATE TABLE
MQBankTB(Name          VARCHAR(40) NOT NULL,
         Account        INTEGER   NOT NULL,
         Balance         INTEGER   NOT NULL,
         Transactions     INTEGER,
         PRIMARY KEY (Account))
END-EXEC.

EXEC SQL CREATE TABLE
MQFeeTB(Account        INTEGER   NOT NULL,
         FeeDue          INTEGER   NOT NULL,
         TranFee         INTEGER   NOT NULL,
         Transactions     INTEGER,
         PRIMARY KEY (Account))
END-EXEC.
```

Enter data into the tables using SQL statements as follows:

```

EXEC SQL INSERT INTO MQBankT VALUES ('Mr Fred Bloggs',1,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Mrs S Smith',2,0);
EXEC SQL INSERT INTO MQBankT VALUES ('Ms Mary Brown',3,0);
:
EXEC SQL INSERT INTO MQBankTB VALUES ('Mr Fred Bloggs',1,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Mrs S Smith',2,0,0);
EXEC SQL INSERT INTO MQBankTB VALUES ('Ms Mary Brown',3,0,0);
:
EXEC SQL INSERT INTO MQFeeTB VALUES (1,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (2,0,50,0);
EXEC SQL INSERT INTO MQFeeTB VALUES (3,0,50,0);
:

```

Note: For COBOL, use the same SQL statements but add END_EXEC at the end of each line.

Precompiling, compiling, and linking the samples

Learn about precompiling, compiling, and linking samples in C and COBOL.

Precompile the .SQC files (in C) and .SQB files (in COBOL), and bind them against the appropriate database to produce the .C or .CBL files. To do this, use the typical method for your database product.

Precompiling in C

```

db2 connect to MQBankDB
db2 prep AMQXSAS0.SQC
db2 connect reset

db2 connect to MQBankDB
db2 prep AMQXSAB0.SQC
db2 connect reset

db2 connect to MQFeeDB
db2 prep AMQXAF0.SQC
db2 connect reset

```

Precompiling in COBOL

```

db2 connect to MQBankDB
db2 prep AMQ0XAS0.SQB bindfile target ibmcob
db2 bind AMQ0XAS0.BND
db2 connect reset

db2 connect to MQBankDB
db2 prep AMQ0XAB0.SQB bindfile target ibmcob
db2 bind AMQ0XAB0.BND
db2 connect reset

db2 connect to MQFeeDB
db2 prep AMQ0XAF0.SQB bindfile target ibmcob
db2 bind AMQ0XAF0.BND
db2 connect reset

```

Compiling and linking

The following sample commands use the symbols <DB2TOP> and MQ_INSTALLATION_PATH. <DB2TOP> represents the installation directory for the DB2 product. MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

- On AIX, the directory path is:

```
/usr/lpp/db2_05_00
```

- On HP-UX and Solaris, the directory path is:

```
/opt/IBMdb2/V5.0
```

- On Windows systems, the directory path depends on the path chosen when installing the product. If you chose the default settings the path is:

```
c:\sqllib
```

Note: Before issuing the link command on Windows systems, ensure that the LIB environment variable contains paths to the DB2 and WebSphere MQ libraries.

Copy the following files into a temporary directory:

- The amqsxag0.c file from your WebSphere MQ installation

Note: This file can be found in the following directories:

- On UNIX and Linux systems:

```
MQ_INSTALLATION_PATH/samp/xatm
```

- On Windows systems:

```
MQ_INSTALLATION_PATH\tools\c\samples\xatm
```

- The .c files that you have obtained by precompiling the .sqc source files, amqsxas0.sqc, amqsxaf0.sqc, and amqsxab0.sqc
- The files util.c and util.h from your DB2 installation.

Note: These files can be found in the directory:

```
<DB2TOP>/samples/c
```

Build the object files for each .c file using the following compiler command for the platform that you are using:

- AIX

```
xlc_r -IMQ_INSTALLATION_PATH/inc -I  
<DB2TOP>/include -c -o  
<FILENAME>.o <FILENAME>.c
```

- HP-UX

```
cc -Aa +z -IMQ_INSTALLATION_PATH/inc -I  
<DB2TOP>/include -c -o  
<FILENAME>.o <FILENAME>.c
```

- Solaris

```
cc -Aa -KPIC -mt -IMQ_INSTALLATION_PATH  
/inc -I<DB2TOP>/include -c -o  
<FILENAME>.o <FILENAME>.c
```

- Windows systems

```
cl /c /IMQ_INSTALLATION_PATH\tools\c\include /I  
<DB2TOP>\include  
<FILENAME>.c
```

Build the amqsxag0 executable file using the following link command for the platform that you are using:

- AIX

```
xlc_r -H512 -T512 -L<DB2TOP>/lib -ldb2 -LMQ_INSTALLATION_PATH/lib
-lmqm util.o amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- HP-UX Revision 11i

```
ld -E -L<DB2TOP>/lib -ldb2 -LMQ_INSTALLATION_PATH/lib -lmqm -lc -lpthread -lcl
/lib/crt0.o util.o amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- Solaris

```
cc -mt -L<DB2TOP>/lib -ldb2 -LMQ_INSTALLATION_PATH/lib
-lmqm -lthread -lsocket -lc -lnsl -ldl util.o
amqsxaf0.o amqsxab0.o amqsxag0.o -o amqsxag0
```

- Windows systems

```
link util.obj amqsxaf0.obj amqsxab0.obj amqsxag0.obj mqm.lib db2api.lib
/out:amqsxag0.exe
```

Build the amqsxas0 executable file using the following compile and link commands for the platform that you are using:

- AIX

```
xlc_r -H512 -T512 -L<DB2TOP>/lib -ldb2
-LMQ_INSTALLATION_PATH/lib -lmqm util.o amqsxas0.o -o amqsxas0
```

- HP-UX Revision 11i

```
ld -E -L<DB2TOP>/lib -ldb2 -LMQ_INSTALLATION_PATH/lib -lmqm -lc -lpthread
-lcl /lib/crt0.o util.o amqsxas0.o -o amqsxas0
```

- Solaris

```
cc -mt -L<DB2TOP>/lib -ldb2 -LMQ_INSTALLATION_PATH/lib
-lmqm -lthread -lsocket -lc -lnsl -ldl util.o
amqsxas0.o -o amqsxas0
```

- Windows systems

```
link util.obj amqsxas0.obj mqm.lib db2api.lib /out:amqsxas0.exe
```

Additional information

If you are working on AIX or HP-UX and want to access Oracle, use the xlc_r compiler and link to libmqm_r.a.

Running the samples

Use this information to learn how to configure the queue manager before running database coordination samples on C and COBOL.

Before you run the samples, configure the queue manager with the database product that you are using. For information about how to do this, see [“Scenario 1: Queue manager performs the coordination” on page 41](#).

The following titles provide information about how to run samples in C and COBOL:

- [“C samples” on page 120](#)
- [“COBOL samples” on page 120](#)

C samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=nnn WHERE Account=nnn
```

AMQSPUT can be used to put the messages on the queue.

The database coordination samples take two parameters:

1. Queue name (required)
2. Queue manager name (optional)

Assuming that you have created and configured a queue manager for the single database sample called singDBQM, with a queue called singDBQ, you increment Mr Fred Bloggs's account by 50 as follows:

```
AMQSPUT singDBQ singDBQM
```

Then key in the following message:

```
UPDATE Balance change=50 WHERE Account=1
```

You can put multiple messages on the queue.

```
AMQSXAS0 singDBQ singDBQM
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have created and configured a queue manager for the multiple-database sample called multDBQM, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ multDBQM
```

Then key in the following message:

```
UPDATE Balance change=-75 WHERE Account=3
```

You can put multiple messages on the queue.

```
AMQSXAG0 multDBQ multDBQM
```

The updated status of Ms Mary Brown's account is then printed.

COBOL samples

Messages must be in the following format to be read from a queue:

```
UPDATE Balance change=snnnnnnnn WHERE Account=nnnnnnnn
```

For simplicity, the Balance change must be a signed eight-character number and the Account must be an eight-character number.

The sample AMQSPUT can be used to put the messages on the queue.

The samples take no parameters and use the default queue manager. It can be configured to run only one of the samples at any time. Assuming that you have configured the default queue manager for the single database sample, with a queue called singDBQ, you increment Mr Fred Bloggs's account by 50 as follows:


```
AMQSPUT singDBQ
```

Then key in the following message:

```
UPDATE Balance change=+000000050 WHERE Account=000000001
```

You can put multiple messages on the queue:

```
AMQ0XAS0
```

Type in the name of the queue:

```
singDBQ
```

The updated status of Mr Fred Bloggs's account is then printed.

Assuming that you have configured the default queue manager for the multiple database sample, with a queue called multDBQ, you decrement Ms Mary Brown's account by 75 as follows:

```
AMQSPUT multDBQ
```

Then key in the following message:

```
UPDATE Balance change=-000000075 WHERE Account=000000003
```

You can put multiple messages on the queue:

```
AMQ0XAG0
```

Type in the name of the queue:

```
multDBQ
```

The updated status of Ms Mary Brown's account is then printed.

Dead-letter queue handler sample

A sample dead-letter queue handler is provided, the name of the executable version is amqsdldq. If you want a dead-letter queue handler that is different from RUNMQDLQ, the source of the sample is available for you to use as your base.

The sample is similar to the dead-letter handler provided within the product but trace and error reporting are different. There are two environment variables available to you:

ODQ_TRACE

Set to YES or yes to switch tracing on

ODQ_MSG

Set to the name of the file containing error and information messages. The file provided is called amqsdldq.msg.

You need to make these variables known to your environment using either the **export** or **set** commands, depending on your platform; trace is turned off using the **unset** command.

You can modify the error message file, amqsdldq.msg, to suit your own requirements. The sample puts messages to stdout, **not** to the WebSphere MQ error log file.

The [Administering](#) or the *System Management Guide* for your platform explains how the dead-letter handler works, and how you run it.

The Distribution List sample program

The Distribution List sample `amqsptl0` gives an example of putting a message on several message queues. It is based on the MQPUT sample, `amqsput0`.

Running the Distribution List sample, `amqsptl0`

The Distribution List sample runs in a similar way to the Put samples.

It takes the following parameters:

- The names of the queues
- The names of the queue managers

These values are entered as pairs. For example:

```
amqsptl0 queue1 qmanagername1 queue2 qmanagername2
```

The queues are opened using MQOPEN and messages are put to the queues using MQPUT. Reason codes are returned if any of the queue or queue manager names are not recognized.

Remember to define channels between queue managers so that messages can flow between them. The sample program does not do that for you.

Design of the Distribution List sample

Put Message Records (MQPMRs) specify message attributes for each destination. The sample provides values for *MsgId* and *CorrelId*, and these override the values specified in the MQMD structure.

The *PutMsgRecFields* field in the MQPMO structure indicates which fields are present in the MQPMRs:

```
MQLONG PutMsgRecFields=MQPMRF_MSG_ID + MQPMRF_CORREL_ID;
```

Next, the sample allocates the response records and object records. The object records (MQORs) require at least one pair of names and an even number of names, that is, *ObjectName* and *ObjectQMGrName*.

The next stage involves connecting to the queue managers using MQCONN. The sample attempts to connect to the queue manager associated with the first queue in the MQOR; if this fails, it goes through the object records in turn. You are informed if it is not possible to connect to any queue manager and the program exits.

The target queues are opened using MQOPEN and the message is put to these queues using MQPUT. Any problems and failures are reported in the response records (MQRRs).

Finally, the target queues are closed using MQCLOSE and the program disconnects from the queue manager using MQDISC. The same response records are used for each call stating the *CompCode* and *Reason*.

The Echo sample programs

The Echo sample programs echo a message from a message queue to the reply queue.

See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

The programs are intended to run as triggered programs.

On UNIX, Linux, and Windows systems, their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue and the queue manager. The COBOL version uses the default queue manager.

When you have set the definition correctly, first start AMQSERV4 in one job, then start AMQSREQ4 in another. You could use AMQSTRG4 instead of AMQSERV4, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send messages to queue SYSTEM.SAMPLE.ECHO. The Echo sample programs send a reply message containing the data in the request message to the reply-to queue specified in the request message.

Design of the Echo sample programs

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG, MQGMO_CONVERT, and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each line of input, the program then reads the text into a buffer and uses the MQPUT1 call to put a request message, containing the text of that line, onto the reply-to queue.

If the MQGET call fails, the program puts a report message on the reply-to queue, setting the *Feedback* field of the message descriptor to the reason code returned by the MQGET.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The Get sample programs

The Get sample programs get messages from a queue using the MQGET call.

See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

Design of the Get sample program

The program opens the target queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each message on the queue, the program uses the MQGET call to remove the message from the queue, then displays the data contained in the message. The MQGET call uses the MQGMO_WAIT option, specifying a *WaitInterval* of 15 seconds, so that the program waits for this period if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code.

The program demonstrates how you must clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The MQGET call specifies a buffer of fixed size. If a message is longer than this buffer, the call fails and the program stops.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes the queue using the MQCLOSE call.

Running the amqsget and amqsgetc samples

These programs each take two parameters:

1. The name of the source queue (required)
2. The name of the queue manager (optional)

If a queue manager is not specified, amqsget connects to the default queue manager, and amqsgetc connects to the queue manager identified by an environment variable or the client channel definition file.

To run these programs, enter one of the following:

- `amqsget myqueue qmanagername`
- `amqsgetc myqueue qmanagername`

where `myqueue` is the name of the queue from which the program will get messages, and `qmanagername` is the queue manager that owns `myqueue`.

If you omit the `qmanagername`, the programs assume the default, or, in the case of the MQI client, the queue manager identified by an environment variable or the client channel definition file.

High availability sample programs

The **amqsgshac**, **amqsphac**, and **amqsmhac** high availability sample programs use automated client reconnection to demonstrate recovery following the failure of a queue manager. **amqsfhac** checks that a queue manager using networked storage maintains data integrity following a failure.

The **amqsgshac**, **amqsphac**, and **amqsmhac** programs are started from the command line, and can be used in combination to demonstrate reconnection after the failure of one instance of a multi-instance queue manager.

Alternatively, you can also use the **amqsgshac**, **amqsphac**, and **amqsmhac** samples to demonstrate client reconnection to single instance queue managers, typically configured into a queue manager group.

To keep the example simple, so it is easy to configure, you are shown the sample programs reconnecting to a single instance queue manager that is started, stopped and then restarted again; see [“Setup and control the queue manager” on page 126](#).

Use **amqsfhac** in parallel with **amqmfscck** to check file system integrity. See [amqmfscck \(file system check\)](#) and [Verifying shared file system behavior](#) for more information.

amqsphac *queueName [qMgrName]*

- **amqsphac** is a IBM WebSphere MQ MQI client application. It puts a sequence of messages to a queue with a two second delay between each message and displays events sent to its event handler.
- No sync point is used to put messages to the queue.
- Reconnection can be made to any queue manager in the same queue manager group.

amqsgshac *queueName [qMgrName]*

- **amqsgshac** is a IBM WebSphere MQ MQI client application. It gets messages from a queue and displays events sent to its event handler.
- No sync point is used to get messages from the queue.
- Reconnection can be made to any queue manager in the same queue manager group.

amqsmhac **-s** *sourceQueueName* **-t** *targetQueueName* [**-m** *qMgrName*] [**-w** *waitInterval*]

- **amqsmhac** is a IBM WebSphere MQ MQI client application. It copies messages from one queue to another with a default wait interval of 15 minutes after the last message that is received before the program finishes.
- The messages are copied within sync point.
- Reconnection can be made only to the same queue manager.

amqsfhac *QueueManagerName QueueName SideQueueName InTransactionCount RepeatCount (0|1|2)*

- **amqsfhac** is a IBM WebSphere MQ MQI client application. It checks that an IBM WebSphere MQ multi-instance queue manager using networked storage, such as a NAS or a cluster file system, maintains data integrity. Follow the steps to run **amqsfhac** in [Verifying shared file system behavior](#).

- It uses the MQCNO_RECONNECT_Q_MGR option when connecting to *QueueManagerName*. It automatically reconnects when the queue manager fails over.
- It puts *InTransactionCount*RepeatCount* persistent messages to *QueueName* during which time you cause the queue manager to fail over any number of times. **amqsfhac** reconnects to the queue manager each time, and continues. The test is to make sure that no messages are lost.
- *InTransactionCount* messages are put within each transaction. The transaction is repeated *RepeatCount* number of times. If a failure occurs within a transaction, **amqsfhac** rolls back and resubmits the transaction when **amqsfhac** reconnects to the queue manager.
- It also puts messages to *SideQueueName*. It uses *SideQueueName* to check whether the all the messages are committed or rolled back from *QueueName* successfully. If it detects an inconsistency, it writes out an error message.
- Vary the amount of output tracing from **amqsfhac** by setting the last parameter to (0|1|2).

0

Least output.

1

Middling output.

2

Most output.

Configuring a client connection

You need to configure a client and server connection channel to run the samples. The client verification procedure explains how to set up a client test environment. See [Verifying a client installation](#).

Alternatively, use the configuration provided in the following example.

Example using amqsgfhac, amqsfhac, and amqsmfhac

The example demonstrates reconnectable clients using a single instance queue manager.

Messages are placed on the queue SOURCE by **amqsfhac**, transferred to TARGET by **amqsmfhac**, and retrieved from TARGET by **amqsgfhac**; see [Figure 19 on page 125](#).




Figure 19. Reconnectable client samples

Follow these steps to run the samples.

1. Create a file `hasamples.txt` containing the commands:

```

DEFINE QLOCAL(SOURCE) REPLACE
DEFINE QLOCAL(TARGET) REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
MCAUSER(MUSR_MQADMIN) REPLACE
DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNAME('LOCALHOST(2345)') QMNAME(QM1) REPLACE
ALTER LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) +
PORT(2345)

```

```
START LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
START CHANNEL(CHANNEL1)
```

2. Type the following commands at a command prompt:

- a. `crtmqm QM1`
- b. `strmqm QM1`
- c. `runmqsc QM1 < hasamples.tst`

3. Set the environment variable **MQCHLLIB** to the path to the AMQCLCHL.TAB client channel definition file; for example, `SET MQCHLLIB=C:\IBM\MQ\MQ7\Data\qmgrs\QM1\@ipcc.`
4. Open three new windows with **MQCHLLIB** set; for example on Windows, type **start** three times at the previous command prompt starting each program in one of the windows. See step “5” on page 127 in “Setup and control the queue manager” on page 126.)
5. Type the command `endmqm -r -p QM1` to stop the queue manager, and then allow the clients to reconnect.
6. Type the command `strmqm QM1` to restart the queue manager.

The results from running the **amqsgbac**, **amqsphac**, and **amqsmhac** samples on Windows are shown in the following examples.

Setup and control the queue manager

1. Create the queue manager.

```
C:\>crtmqm QM1
WebSphere MQ queue manager created.
Directory 'C:\IBM\MQ\MQ7\Data\qmgrs\QM1' created.
Creating or replacing default objects for QM1.
Default objects statistics : 67 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
```

Remember the data directory to set the **MQCHLLIB** variable later.

2. Start the queue manager.

```
C:\>strmqm QM1
WebSphere MQ queue manager 'QM1' starting.
5 log records accessed on queue manager 'QM1' during the log replay phase.
Log replay for queue manager 'QM1' complete.
Transaction manager state recovered for queue manager 'QM1'.
WebSphere MQ queue manager 'QM1' started.
```

3. Create the queues and channels, modify the listener port, and start the listener and channel.

```
C:\>runmqsc QM1 < hasamples.tst
5724-H72 (C) Copyright IBM Corp. 1994, 2024. ALL RIGHTS RESERVED.
Starting MQSC for queue manager QM1.

      1 : DEFINE QLOCAL(SOURCE) REPLACE
AMQ8006: WebSphere MQ queue created.
      2 : DEFINE QLOCAL(TARGET) REPLACE
AMQ8006: WebSphere MQ queue created.
      3 : DEFINE CHANNEL(CHANNEL1) CHLTYPE(SVRCONN) TRPTYPE(TCP) MCAUSER(MUSR_MQADMIN)
REPLACE
AMQ8014: WebSphere MQ channel created.
      4 : DEFINE CHANNEL(CHANNEL1) CHLTYPE(CLNTCONN) TRPTYPE(TCP) CONNAME('LOCALHOST(2345)')
QMNAME(QM1) REPLACE
AMQ8014: WebSphere MQ channel created.
      5 : ALTER LISTENER(SYSTEM.DEFAULT.LISTENER.TCP) TRPTYPE(TCP) PORT(2345)
AMQ8623: WebSphere MQ listener changed.
      6 : START LISTENER(SYSTEM.DEFAULT.LISTENER.TCP)
AMQ8021: Request to start WebSphere MQ Listener accepted.
      7 : START CHANNEL(CHANNEL1)
AMQ8018: Start WebSphere MQ channel accepted.
7 MQSC commands read.
No commands have a syntax error.
All valid MQSC commands were processed.
```

4. Make the client channel table known to the clients.

Use the data directory returned from the **crtmqm** command in step “1” on page 126, and add the directory @ipcc to it to set the **MQCHLLIB** variable.

```
C:\>SET MQCHLLIB=C:\IBM\MQ\MQ7\Data\qmgrs\QM1\@ipcc
```

5. Start the sample programs in the other windows

```
C:\>start amqsphac SOURCE QM1
C:\>start amqsmhac -s SOURCE -t TARGET -m QM1
C:\>start amqsgnac TARGET QM1
```

6. End the queue manager and restart it again.

```
C:\>endmqm -r -p QM1
Waiting for queue manager 'QM1' to end.
WebSphere MQ queue manager 'QM1' ending.
WebSphere MQ queue manager 'QM1' ended.

C:\>stimqm QM1
WebSphere MQ queue manager 'QM1' starting.
5 log records accessed on queue manager 'QM1' during the log replay phase.
Log replay for queue manager 'QM1' complete.
Transaction manager state recovered for queue manager 'QM1'.
WebSphere MQ queue manager 'QM1' started.
```

amqsphac

```
Sample AMQSPHAC start
target queue is SOURCE
message <Message 1>
message <Message 2>
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnectedmessage
<Message 3>
message <Message 4>
message <Message 5>
```

amqsmhac

```
Sample AMQSMHA0 start
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnected
No more messages.
Sample AMQSMHA0 end
C:\>
```

amqsgnac

```
Sample AMQSGHAC start
message <Message 1>
message <Message 2>
16:25:22 : EVENT : Connection Reconnecting (Delay: 0ms)
16:25:45 : EVENT : Connection Reconnecting (Delay: 0ms)
16:26:02 : EVENT : Connection Reconnected
message <Message 3>
message <Message 4>
message <Message 5>
```

Related tasks

[Verifying shared file system behavior](#)

Related reference

[amqmfsc \(file system check\)](#)

The Inquire sample programs

The Inquire sample programs inquire about some of the attributes of a queue using the MQINQ call.

See [“Features demonstrated in the sample programs”](#) on page 93 for the names of these programs.

These programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure for IBM i, Windows, UNIX and Linux systems. This structure contains the name of a target queue with attributes that are to be inquired upon. The C version also uses the queue manager name. The COBOL version uses the default queue manager.

For the triggering process to work, ensure that the Inquire sample program that you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.INQ. To do this, specify the name of the Inquire sample program that you want to use in the *ApplicId* field of the process definition SYSTEM.SAMPLE.INQPROCESS. The sample queue has a trigger type of FIRST; if there are already messages on the queue before you run the request sample, the inquire sample is not triggered by the messages that you send.

When you have set the definition correctly:

- For UNIX, Linux and Windows systems, start the **runmqtrm** program in one session, then start the **amqsreq** program in another.

Use the Request sample programs to send request messages, each containing just a queue name, to queue SYSTEM.SAMPLE.INQ. For each request message, the Inquire sample programs send a reply message containing information about the queue specified in the request message. The replies are sent to the reply-to queue specified in the request message.

Design of the Inquire sample program

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the MQOPEN call to open this queue for shared input.

The program uses the MQGET call to remove messages from this queue. This call uses the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the *target queue*) contained in the data, and opens that queue using the MQOPEN call with the MQOO_INQ option. The program then uses the MQINQ call to inquire about the values of the *InhibitGet*, *CurrentQDepth*, and *OpenInputCount* attributes of the target queue.

If the MQINQ call is successful, the program uses the MQPUT1 call to put a reply message on the reply-to queue. This message contains the values of the three attributes.

If the MQOPEN or MQINQ call is unsuccessful, the program uses the MQPUT1 call to put a report message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the MQOPEN or MQINQ call, depending on which one failed.

After the MQINQ call, the program closes the target queue using the MQCLOSE call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The INQUIRE Properties of a Message Handle sample program

AMQSIQMA is a sample C program to inquire properties of a message handle from a message queue, and is an example of the use of the MQINQM API call.

This sample creates a message handle and puts it into the MsgHandle field of the MQGMO structure. The sample then gets one message and inquires and prints all properties with which the message handle was populated.

```
C:\Program Files\IBM\WebSphere MQ\tools\c\Samples\Bin >amqsiqm Q QM1
Sample AMQSIQMA start
property name <MyProp> value <MyValue>
message text <Hello world!>
Sample AMQSIQMA end
```

The Publish/Subscribe sample programs

The publish/subscribe sample programs demonstrate the use of the publish and subscribe features in WebSphere MQ.

There are three C language sample programs illustrating how to program to the WebSphere MQ publish/subscribe interface. There are some C samples that use older interfaces, and there are Java samples. The Java samples use the WebSphere MQ publish/subscribe interface in com.ibm.mq.jar and the JMS publish/subscribe interface in com.ibm.mqjms. The JMS samples are not covered in this topic.

C

Find the publisher sample amqspub in the C samples folder. Run it with any topic name you like as the first parameter, followed by an optional queue manager name. For example, `amqspub mytopic QM3`. There is also a client version called `amqspubc`. If you choose to run the client version, first see [“Preparing and running the sample programs”](#) on page 104 for details.

The publisher connects to the default queue manager and responds with the output, `target topic is mytopic`. Every line you enter into this window from now onwards is published to `mytopic`.

Open another command window in the same directory, and run the subscriber program, `amqssub`, supplying it with the same topic name, and an optional queue manager name. For example, `amqssub mytopic QM3`.

The subscriber responds with the output, `Calling MQGET : 30 seconds wait time`. From now onwards, lines you type into the publisher appear in the output of the subscriber.

Start another subscriber in another command window, and watch both subscribers receive publications.

For full documentation of the parameters, including setting options, refer to the sample source code. The values for the subscriber options field is described in the following topic: [Options \(MQLONG\)](#).

There is another subscriber sample `amqssbx`, which offers additional subscription options as command line switches.

Type `amqssbx -d mysub -t mytopic -k` to invoke the subscriber using durable subscriptions that are retained after the subscriber has terminated.

Test the subscription by publishing another item using the publisher. Wait for 30 seconds for the subscriber to terminate. Publish some more items under the same topic. Restart the subscriber. The last item published while the subscriber was not running is displayed by the subscriber immediately it is restarted.

C legacy

There is an additional set of C samples which demonstrate queued commands. Some of these samples were originally shipped as part of the MQOC Supportpac. The capabilities the samples demonstrate are fully supported, for compatibility reasons.

We discourage you from using the queued command interface. It is much more complex than the publish/subscribe API, and there is no compelling functional reason to program complex queued commands. However, you might find the queued approach more suitable, perhaps because you are already using the interface, or because your programming environment makes it easier to build a complex message and call a generic MQPUT, rather than constructing different calls to MQSUB.

The additional samples are located in the pubsub subdirectory in the samples folder.

There are six types of sample listed in [Table 20 on page 130](#).

<i>Table 20. Categories of legacy publish/subscribe sample C programs</i>		
Category	Programs	Comments
RFH1	amqssr1a.c amqspr1a.c	Simple publish/subscribe example built using RFH1 format messages.
RFH2	amqssr2a.c amqspr2a.c	Simple publish/subscribe example built using RFH2 format messages.
MQAI samples	amqsppca.c amqsspca.c	Simple publish/subscribe example built using PCF commands and the MQAI command interface.
MAOC Results service using RFH1	amqsgama.c amqsresa.c	Results service built using RFH1 headers 1. Requires the queues defined in amqsgama.tst and amqsresa.tst 2. amqsresa must be started before amqsgama
MAOC Results service using RFH2	amqsgr2a.c amqsrr2a.c	Results service built using RFH2 headers 1. Requires the queues defined in amqsgama.tst and amqsresa.tst 2. amqsresa must be started before amqsgama
Routing exit publish/subscribe sample	amqspdra.c	Demonstrates how to change the queue or queue manager destination for a publish/subscribe message in a routing exit.

Java

The Java sample MQPubSubApiSample.java combines publisher and subscribers in a single program. Its source and compiled class files are found in the wmqjava samples folder.

If you choose to run in client mode, first see [“Preparing and running the sample programs” on page 104](#) for details.

Run the sample from the command line using the Java command, if you have a Java environment configured. You can also run the sample from the WebSphere MQ Explorer Eclipse workspace that has a Java programming workbench already set up.

You might need to change some of the sample program's properties to run it. You do this by providing parameters to the JVM, or editing the source.

The instructions in [“Running the MQPubSubApiSample Java sample” on page 131](#) show how to run the sample from the Eclipse workspace.

Running the MQPubSubApiSample Java sample

How to run the MQPubSubApiSample using the Java Development Tools from the Eclipse platform.

Before you begin

Open the Eclipse workbench. Create a new workspace directory and select it. Close the welcome window.

Follow the steps in [“Preparing and running the sample programs” on page 104](#) prior to running as a client.

About this task

The Java publish/subscribe sample program is a WebSphere MQ MQI client Java program. The sample runs without modification using a default queue manager listening on port 1414. The task describes this simple case, and indicates in general terms how to provide parameters and modify the sample to suit different WebSphere MQ configurations. The example is illustrated running on Windows. The file paths will differ on other platforms.

Procedure

1. Import the Java sample programs
 - a) In the workbench, click **Window > Open perspective > Other > Java** and click **OK**.
 - b) Switch to the **Package Explorer** view.
 - c) Right-click in the white-space in the **Package Explorer** view. Click **New > Java project**.
 - d) In the **Project name** field type `MQ Java Samples`. Click **Next**.
 - e) In the **Java Settings** panel, switch to the **Libraries** tab.
 - f) Click **Add External JARs**.
 - g) Browse to `MQ_INSTALLATION_PATH\java\lib` where `MQ_INSTALLATION_PATH` is the WebSphere MQ installation folder and select `com.ibm.mq.jar` and `com.ibm.mq.jmqi.jar`
 - h) Click **Open > Finish**.
 - i) Right-click `src` in the **Package Explorer** view.
 - j) Select **Import... > General > File System > Next > Browse...** and browse to the path `MQ_INSTALLATION_PATH\tools\wmqjava\samples` where `MQ_INSTALLATION_PATH` is the WebSphere MQ installation directory.
 - k) On the **Import** panel, [Figure 20 on page 132](#), click `samples` (do not select the check box).
 - l) Select `MQPubSubApiSample.java`. The **Into folder** field should contain `MQ Java Samples/src`. Click **Finish**.

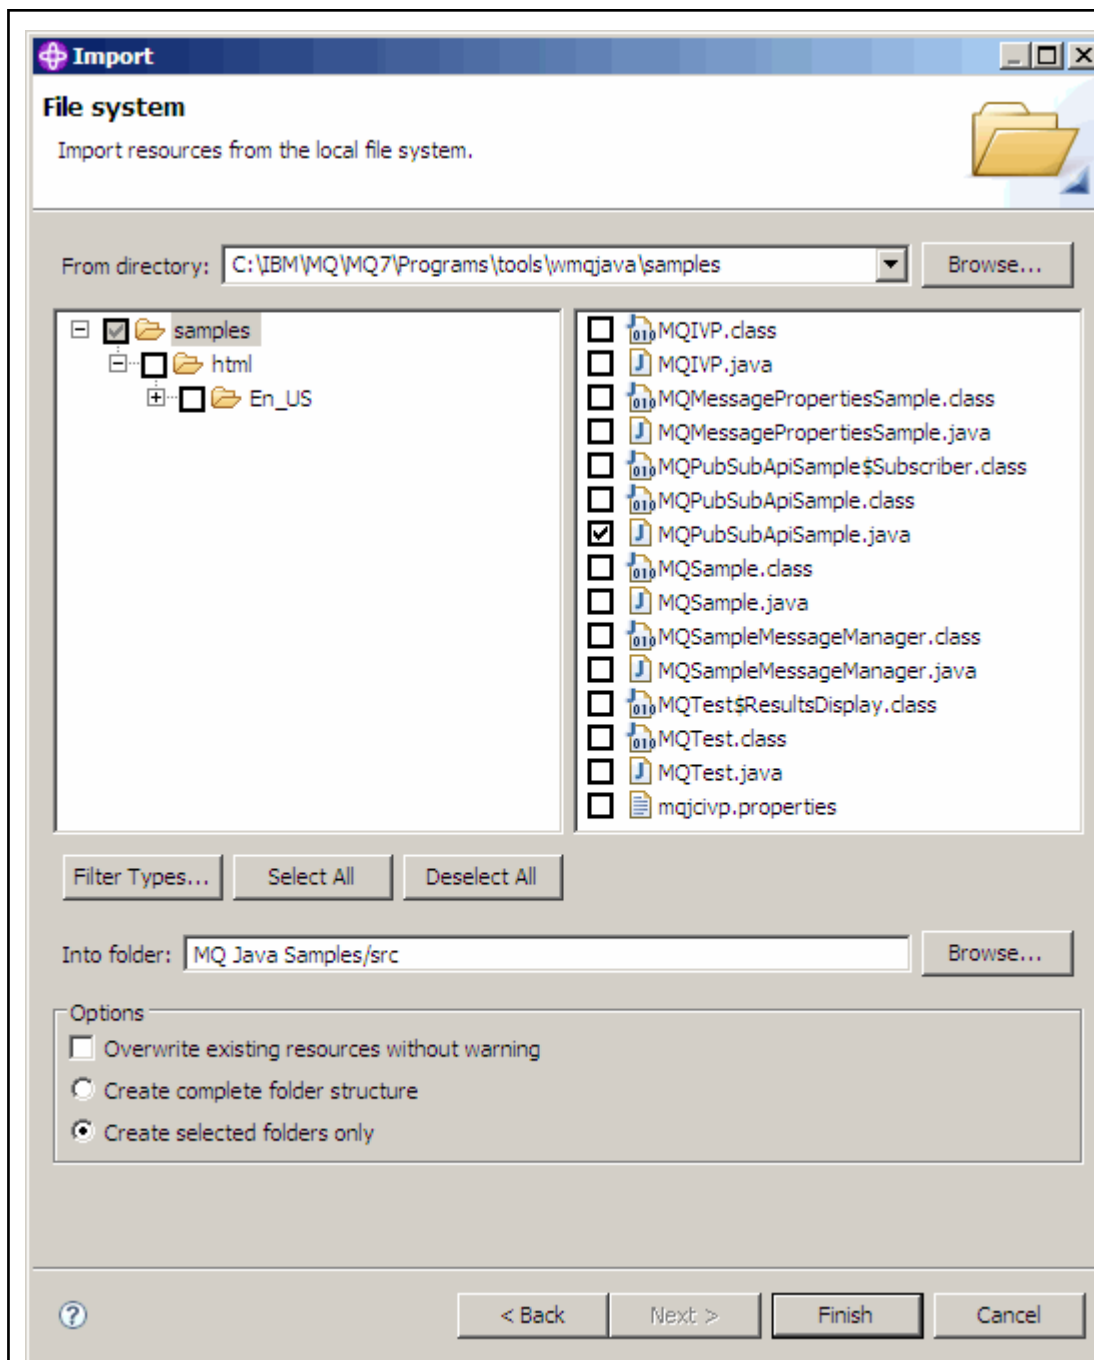


Figure 20. File system import

2. Run the publish/subscribe sample program.

There are two ways to run the program, depending on whether you need to change the default parameters.

- The first choice runs the program without making any changes:
 - In the workspace main menu, expand the `src` folder. Right-click **MQPubSubApiSample.java** **Run-as > 1. Java Application**
- The second choice runs the program with parameters or with modified source code for your environment:
 - Open `MQPubSubApiSample.java` and study the `MQPubSubApiSample` constructor.

- Modify the attributes of the program.

These attributes are modifiable using the -D JVM switch, or by providing a default value for the System property by editing the source code.

- topicObject
- queueManagerName
- subscriberCount

These attributes are changeable only by editing the source code in the constructor.

- hostname
- port
- channel

To set System properties, code a default value in the accessor, for example:

```
queueManagerName = System.getProperty("com.ibm.mq.pubSubSample.queueManagerName",
"QM3");
```

Or provide the parameter to the JVM using the -D option, as shown in the following steps:

- Copy the full name of the System.Property you want to set, for example:
com.ibm.mq.pubSubSample.queueManagerName.
- In the workspace, right-click **Run > Open Run Dialog**. Double click Java Application in **Create, Manage and Run applications** and click the **(x) = Arguments** tab.
- In the **VM arguments:** pane, type -D and paste the System.property name, com.ibm.mq.pubSubSample.queueManagerName, followed by =QM3. Click **Apply > Run**.
- Add further arguments as a comma separated list, or as additional lines in the pane, without comma separators.

For example: -Dcom.ibm.mq.pubSubSample.queueManagerName=QM3,
-Dcom.ibm.mq.pubSubSample.subscriberCount=6.

The Publish Exit sample program

AMQSPSE0 is a sample C program of an exit to intercept a publication before it is delivered to a subscriber. The exit can then, for example, alter the message headers, payload or destination, or prevent the message being published to a subscriber.

To run the sample, perform the following tasks:

1. Configure the queue manager:
 - On UNIX and Linux systems add a stanza like this to the qm.ini file:

```
PublishSubscribe:
    PublishExitPath=<Module>
    PublishExitFunction=EntryPoint
```

where the module is *MQ_INSTALLATION_PATH/samp/bin/amqspse*. *MQ_INSTALLATION_PATH* represents the high-level directory in which WebSphere MQ is installed. On Windows set the equivalent attributes in the registry.

2. Make sure the Module is accessible to WebSphere MQ.
3. Restart the Queue Manager to pick up the configuration.
4. In the application process to be traced, describe where the trace files should be written to. For example:
 - On UNIX and Linux systems, ensure that the directory */var/mqm/trace* exists and export the following environment variable:

```
export MQPSE_TRACE_LOGFILE=/var/mqm/trace/PubTrace
```

- On Windows, ensure that the directory C:\temp exists and set the following environment variable:

```
set MQPSE_TRACE_LOGFILE=C:\temp\PubTrace
```

The Put sample programs

The Put sample programs put messages on a queue using the MQPUT call.

See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

Design of the Put sample program

The program uses the MQOPEN call with the MQOO_OUTPUT option to open the target queue for putting messages.

If it cannot open the queue, the program outputs an error message containing the reason code returned by the MQOPEN call. To keep the program simple, on this and on subsequent MQI calls, the program uses default values for many of the options.

For each line of input, the program reads the text into a buffer and uses the MQPUT call to create a datagram message containing the text of that line. The program continues until either it reaches the end of the input or the MQPUT call fails. If the program reaches the end of the input, it closes the queue using the MQCLOSE call.

Running the Put sample programs

Running the amqspout and amqsputc samples

These programs each take 2 parameters:

1. The name of the target queue (required)
2. The name of the queue manager (optional)

If a queue manager is not specified, amqspout connects to the default queue manager and amqsputc connects to the queue manager identified by an environment variable or the client channel definition file. To run these programs, enter one of the following:

- amqspout myqueue qmanagername
- amqsputc myqueue qmanagername

where myqueue is the name of the queue on which the messages are going to be put, and qmanagername is the queue manager that owns myqueue.

Running the amq0put sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target queue
```

It takes input from StdIn and adds each line of input to the target queue. A blank line indicates there is no more data.

The Reference Message sample programs

The Reference Message samples allow a large object to be transferred from one node to another (usually on different systems) without the need for the object to be stored on WebSphere MQ queues at either the source or the destination nodes.

A set of sample programs is provided to demonstrate how Reference Messages can be put to a queue, received by message exits, and taken from a queue. The sample programs use Reference Messages to

move files. If you want to move other objects such as databases, or if you want to perform security checks, define your own exit, based on our sample, amqsxrm. The following sections describe the Reference Message sample programs.

There version of the Reference Message exit sample program to use depends on the platform on which the channel is running. On all platforms, use amqsxrma at the sending end. Use amqsxrma at the receiving end if the receiver is running under any WebSphere MQ product except WebSphere MQ for IBM i.

Running the Reference Message samples

Use this information to learn how to run Reference Message sample programs.

The Reference Message samples run as follows:

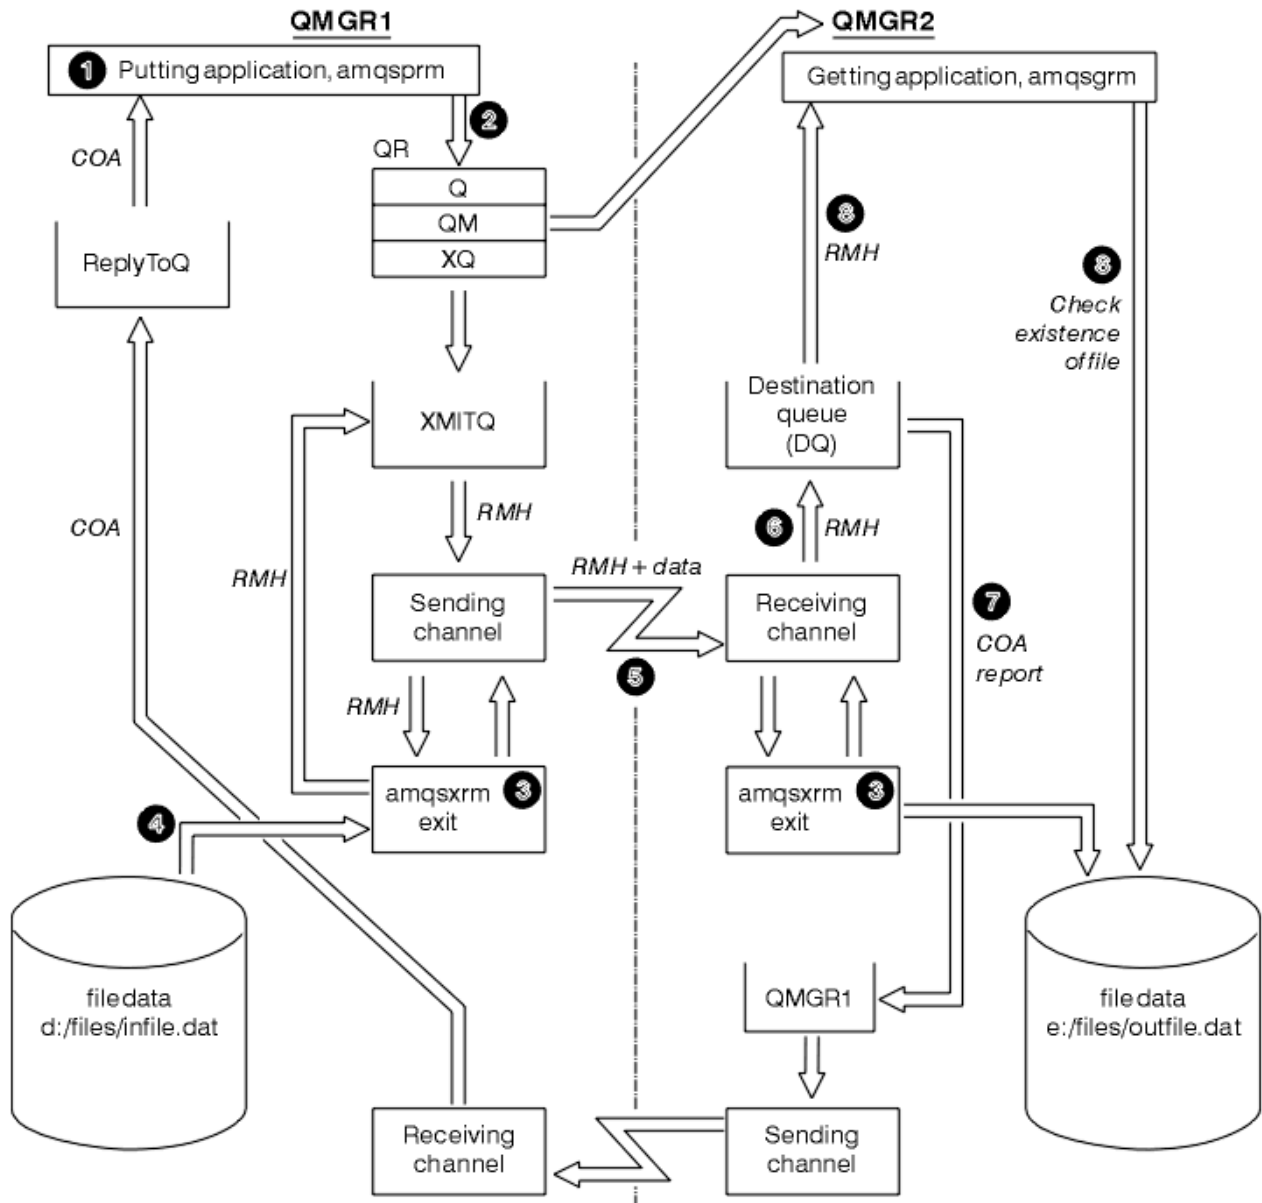


Figure 21. Running the Reference Message samples

1. Set up the environment to start the listeners, channels, and trigger monitors, and define your channels and queues.

For the purposes of describing how to set up the Reference Message example this refers to the sending machine as MACHINE1 with a queue manager called QMGR1 and the receiving machine as MACHINE2 with a queue manager called QMGR2.

Note: The following definitions allow a Reference Message to be built to send a file with an object type of FLATFILE from queue manager QMGR1 to QMGR2 and to re-create the file as defined in the call to AMQSPRM (or AMQSPRMA on IBM i). The Reference Message (including the file data) is sent using channel CHL1 and transmission queue XMITQ and placed on queue DQ. Exception and COA reports are sent back to QMGR1 using the channel REPORT and transmission queue QMGR1.

The application that receives the Reference Message (AMQSGRM) is triggered using the initiation queue INITQ and process PROC. Ensure that the CONNAME fields are set correctly and the MSGEXIT field reflects your directory structure, depending on machine type and where the WebSphere MQ product is installed.

The MQSC definitions have used an AIX style for defining the exits. It is important to note that the message data FLATFILE is case sensitive and the sample will not work unless it is in uppercase.

On machine MACHINE1, queue manager QMGR1

MQSC syntax

```
define chl(chl1) chltype(sdr) trptype(tcp) conname('machine2') xmitq(xmitq)
msgdata(FLATFILE) msgexit('/usr/lpp/mqm/samp/bin/amqsxrm(MsgExit)')

define ql(xmitq) usage(xmitq)

define chl(report) chltype(rcvr) trptype(tcp) replace

define qr(qr) rname(dq) rqmname(qmgr2) xmitq(xmitq) replace
```

Note: If you do not specify a queue manager name the system uses the default queue manager.

```
CRTMQMCHL CHLNAME(CHL1) CHLTYPE(*SDR) MQMNAME(QMGR1) +
          REPLACE(*YES) TRPTYPE(*TCP) +
          CONNAME('MACHINE2(60501)') TMQNAME(XMITQ) +
          MSGEXIT(QMQM/AMQSXRM4) MSGUSRDATA(FLATFILE)

CRTMQMQ QNAME(XMITQ) QTYPE(*LCL) MQMNAME(QMGR1) +
        REPLACE(*YES) USAGE(*TMQ)

CRTMQMCHL CHLNAME(REPORT) CHLTYPE(*RCVR) +
          MQMNAME(QMGR1) REPLACE(*YES) TRPTYPE(*TCP)

CRTMQMQ QNAME(QR) QTYPE(*RMT) MQMNAME(QMGR1) +
        REPLACE(*YES) RMTQNAME(DQ) +
        RMTMQMNAME(QMGR2) TMQNAME(XMITQ)
```

On machine MACHINE2, queue manager QMGR2

MQSC syntax

```
define chl(chl1) chltype(rcvr) trptype(tcp)
msgexit('/usr/lpp/mqm/samp/bin/amqsxrm(MsgExit)')
msgdata(flatfile)

define chl(report) chltype(sdr) trptype(tcp) conname('MACHINE1')
xmitq(qmgr1)

define ql(initq)

define ql(qmgr1) usage(xmitq)

define pro(proc) applicid('/usr/lpp/mqm/samp/bin/amqsgrm')

define ql(dq) initq(initq) process(proc) trigger trigtype(first)
```

2. Once the WebSphere MQ objects have been created:
 - a. Where applicable to the platform, start the listener for the sending and receiving queue managers
 - b. Start the channels CHL1 and REPORT
 - c. On the receiving queue manager start the trigger monitor for the initiation queue INITQ

3. Invoke the put Reference Message sample program AMQSPRM from the command line using the following parameters:

- m Name of the local queue manager; this defaults to the default queue manager
- i Name and location of source file
- o Name and location of destination file
- q Name of queue
- g Name of queue manager where the queue, defined in the -q parameter exists This defaults to the queue manager specified in the -m parameter
- t Object type
- w Wait interval, that is, the waiting time for exception and COA reports from the receiving queue manager

For example, to use the sample with the objects defined previously you would use the following parameters:

```
-mQMGR1 -iInput File -oOutput File -qQR -tFLATFILE -w120
```

Increasing the waiting time allows time for a large file to be sent across a network before the program putting the messages times out.

```
amqsprpm -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
```

Note: For UNIX and Linux platforms, you must use two backslashes (\\) instead of one to denote the destination file directory. Therefore, the **amqsprpm** command looks like this:

```
amqsprpm -i /files/infile.dat -o e:\\files\\outfile.dat -q QR  
-m QMGR1 -w 30 -t FLATFILE
```

Running the put Reference Message program does the following:

- The Reference Message is put to queue QR on queue manager QMGR1.
 - The source file and path are d:\files\infile.dat and exists on the system where the example command is issued.
 - If the queue QR is a remote queue, the Reference Message is sent to another queue manager, on a different system, where a file is created with the name and path e:\files\outfile.dat. The contents of this file are the same as the source file.
 - amqsprpm waits for 30 seconds for a COA report from the destination queue manager.
 - The object type is flatfile, so the channel used to move messages from the queue QR must specify this in the *MsgData* field.
4. When you define your channels, select the message exit at both the sending and receiving ends to be amqsxrm. This is defined on WebSphere MQ for Windows as follows:

```
msgexit('pathname\amqsxrm.dll(MsgExit)')
```

This is defined on WebSphere MQ for AIX, WebSphere MQ for HP-UX, and WebSphere MQ for Solaris as follows:

```
msgexit('pathname/amqsxrm(MsgExit)')
```

If you specify a path name, specify the complete name. If you omit the path name, it is assumed that the program is in the path specified in the `qm.ini` file (or, on WebSphere MQ for Windows, the path specified in the registry).

5. The channel exit reads the Reference Message header and finds the file that it refers to.
6. The channel exit can then segment the file before sending it down the channel along with the header. On WebSphere MQ for AIX, WebSphere MQ for HP-UX, and WebSphere MQ for Solaris, change the group owner of the target directory to 'mqm' so that the sample message exit can create the file in that directory. Also, change the permissions of the target directory to allow mqm group members to write to it. The file data is not stored on the WebSphere MQ queues.
7. When the last segment of the file is processed by the receiving message exit, the Reference Message is put to the destination queue specified by `amqsprmq`. If this queue is triggered (that is, the definition specifies *Trigger*, *InitQ*, and *Process* queue attributes), the program specified by the `PROC` parameter of the destination queue is triggered. The program to be triggered must be defined in the `ApplId` field of the *Process* attribute.
8. When the Reference Message reaches the destination queue (DQ), a COA report is sent back to the putting application (`amqsprmq`).
9. The Get Reference Message sample, `amqsgmr`, gets messages from the queue specified in the input trigger message and checks the existence of the file.

Design of the Put Reference Message sample (amqsprma.c, AMQSPRM4)

This topic gives a detailed description of a Put Reference Message sample.

This sample creates a Reference Message that refers to a file and puts it on a specified queue:

1. The sample connects to a local queue manager using `MQCONN`.
2. It then opens (`MQOPEN`) a model queue that is used to receive report messages.
3. The sample builds a Reference Message containing the values required to move the file, for example, the source and destination file names and the object type. As an example, the sample shipped with WebSphere MQ builds a Reference Message to send the file `d:\x\file.in` from QMGR1 to QMGR2 and to re-create the file as `d:\y\file.out` using the following parameters:

```
amqsprmq -q QR -m QMGR1 -i d:\x\file.in -o d:\y\file.out -t FLATFILE
```

Where `QR` is a remote queue definition that refers to a target queue on QMGR2.

Note: For UNIX and Linux platforms, use two backslashes (`\\`) instead of one to denote the destination file directory. Therefore, the `amqsprmq` command looks like this:

```
amqsprmq -q QR -m QMGR1 -i /x/file.in -o d:\\y\\file.out -t FLATFILE
```

4. The Reference Message is put (without any file data) to the queue specified by the `/q` parameter. If this is a remote queue, the message is put to the corresponding transmission queue.
5. The sample waits, for the duration of time specified in the `/w` parameter (which defaults to 15 seconds), for COA reports, which, along with exception reports, are sent back to the dynamic queue created on the local queue manager (QMGR1).

Design of the Reference Message Exit sample (amqsxrma.c, AMQSXRMA4)

This sample recognizes Reference Messages with an object type that matches the object type in the message exit user data field of the channel definition.

For these messages, the following happens:

- At the sender or server channel, the specified length of data is copied from the specified offset of the specified file into the space remaining in the agent buffer after the Reference Message. If the end of the file is not reached, the Reference Message is put back on the transmission queue after updating the `DataLogicalOffset` field.

- At the requester or receiver channel, if the *DataLogicalOffset* field is zero and the specified file does not exist, it is created. The data following the Reference Message is added to the end of the specified file. If the Reference Message is not the last one for the specified file, it is discarded. Otherwise, it is returned to the channel exit, without the appended data, to be put on the target queue.

For sender and server channels, if the *DataLogicalLength* field in the input Reference Message is zero, the remaining part of the file, from *DataLogicalOffset* to the end of the file, is to be sent along the channel. If it is not zero, only the length specified is sent.

If an error occurs (for example, if the sample cannot open a file), MQCXP.*ExitResponse* is set to MQXCC_SUPPRESS_FUNCTION so that the message being processed is put to the dead-letter queue instead of continuing to the destination queue. A feedback code is returned in MQCXP.*Feedback* and returned to the application that put the message in the *Feedback* field of the message descriptor of a report message. This is because the putting application requested exception reports by setting MQRO_EXCEPTION in the *Report* field of the MQMD.

If the encoding or *CodedCharacterSetId* (CCSID) of the Reference Message is different from that of the queue manager, the Reference Message is converted to the local encoding and CCSID. In our sample, amqsprm, the format of the object is MQFMT_STRING, so amqsxrm converts the object data to the local CCSID at the receiving end before the data is written to the file.

Do not specify the format of the file being transferred as MQFMT_STRING if the file contains multibyte characters (for example, DBCS or Unicode). This is because a multibyte character could be split when the file is segmented at the sending end. To transfer and convert such a file, specify the format as something other than MQFMT_STRING so that the Reference Message exit does not convert it and convert the file at the receiving end when the transfer is complete.

Compiling the Reference Message Exit sample

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

To compile amqsxrma, use the following commands:

On AIX

```
xlc_r -q64 -e MsgExit -bE:amqsxrm.exp -bM:SRE -o amqsxrm_64_r
-LMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -lmqm_r amqsqrma.c
```

On HP-UX

```
$ cc89 +DD64 +z -c -D_HPUX_SOURCE -o amqsxrma.o amqsqrma.c -IMQ_INSTALLATION_PATH/inc
$ ld -b amqsxrma.o -o /var/mqm/exits64/amqsxrma -LMQ_INSTALLATION_PATH/lib64
-L/usr/lib/pa20_64 -lmqm_r -lpthread
```

On Linux

```
$ gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsxrma amqsqrma.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm_r
```

On Solaris

```
$ cc -xarch=v9 -mt -G -o /var/mqm/exits64/amqsxrma amqsqrma.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm
-lsocket
-lnsl -ldl
```

On Windows

WebSphere MQ now supplies the mqm library with client packages as well as server packages, so the following example uses `mqm.lib` instead of `mqmvx.lib`:

```
cl amqsqrma.c /link /out:amqsxrm.dll /dll mqm.lib mqm.lib /def:amqsxrm.def
```

For general information about writing and compiling channel exits, see [“Writing channel-exit programs” on page 381](#)

Design of the Get Reference Message sample (amqsgrma.c, AMQSGRM4)

This topic explains the design of the Get Reference Message sample.

The program logic is as follows:

1. The sample is triggered and extracts the queue and queue manager names from the input trigger message.
2. It then connects to the specified queue manager using MQCONN and opens the specified queue using MQOPEN.
3. The sample issues MQGET with a wait interval of 15 seconds within a loop to get messages from the queue.
4. If a message is a Reference Message, the sample checks the existence of the file that has been transferred.
5. It then closes the queue and disconnects from the queue manager.

The Request sample programs

The Request sample programs demonstrate client/server processing. The samples are the clients that put request messages on a target server queue that is processed by a server program. They wait for the server program to put a reply message on a reply-to queue.

The Request samples put a series of request messages on the target server queue using the MQPUT call. These messages specify the local queue, SYSTEM.SAMPLE.REPLY as the reply-to queue, which can be a local or remote queue. The programs wait for reply messages, then display them. Replies are sent only if the target server queue is being processed by a server application, or if an application is triggered for that purpose (the Inquire, Set, and Echo sample programs are designed to be triggered). The C sample waits 1 minute (the COBOL sample waits 5 minutes), for the first reply to arrive (to allow time for a server application to be triggered), and 15 seconds for subsequent replies, but both samples can end without getting any replies. See [“Features demonstrated in the sample programs” on page 93](#) for the names of the Request sample programs.

Running the Request sample programs

Running the amqsreq0.c, amqsreq, and amqsreqc samples

The C version of the program takes three parameters:

1. The name of the target server queue (necessary)
2. The name of the queue manager (optional)
3. The reply queue (optional)

For example, enter one of the following:

- `amqsreq myqueue qmanagername replyqueue`
- `amqsreqc myqueue qmanagername`
- `amq0req0 myqueue`

where `myqueue` is the name of the target server queue, `qmanagername` is the name of the queue manager that owns `myqueue`, and `replyqueue` is the name of the reply queue.

If you omit the name of the queue manager, it is assumed that the default queue manager owns the queue. If you omit the name of the reply queue, the default reply queue is provided.

Running the amq0req0.cbl sample

The COBOL version does not have any parameters. It connects to the default queue manager and when you run it you are prompted:

```
Please enter the name of the target server queue
```

The program takes its input from StdIn and adds each line to the target server queue, taking each line of text as the content of a request message. The program ends when a null line is read.

Running the AMQSREQ4 sample

The C program creates messages by taking data from stdin (the keyboard) with a blank time terminating input. The program takes up to three parameters: the name of the target queue (required), the queue manager name (optional), and the reply-to queue name (optional). If no queue manager name is specified, the default queue manager is used. If no reply-to queue is specified, the SYSTEM.SAMPLE.REPLY queue is used.

Here is an example of how to call the C sample program, specifying the reply-to queue, but letting the queue manager default:

```
CALL PGM(QMQM/AMQSREQ4) PARM('SYSTEM.SAMPLE.LOCAL' ' ' 'SYSTEM.SAMPLE.REPLY')
```

Note: Remember that queue names are case sensitive. All the queues created by the sample file create program AMQSAMP4 have names created in uppercase characters.

Running the AMQOREQ4 sample

The COBOL program creates messages by accepting data from the keyboard. To start the program, call the program and specify the name of your target queue as a parameter. The program accepts input from the keyboard into a buffer and creates a request message for each line of text. The program stops when you enter a blank line at the keyboard.

Running the Request sample using triggering

If the sample is used with triggering and one of the Inquire, Set, or Echo sample programs, the line of input must be the queue name of the queue that you want the triggered program to access.

UNIX, Linux and Windows systems

To run the samples using triggering:

1. Start the trigger monitor program RUNMQTRM in one session (the initiation queue SYSTEM.SAMPLE.TRIGGER is available for you to use).
2. Start the amqsreq program in another session.
3. Make sure that you have defined a target server queue.

The sample queues available to you to use as the target server queue for the request sample to put messages are:

- SYSTEM.SAMPLE.INQ - for the Inquire sample program
- SYSTEM.SAMPLE.SET - for the Set sample program
- SYSTEM.SAMPLE.ECHO - for the Echo sample program

These queues have a trigger type of FIRST, so if there are already messages on the queues before you run the Request sample, server applications are not triggered by the messages you send.

4. Make sure that you have defined a queue for the Inquire, Set or Echo sample program to use.

This means that the trigger monitor is ready when the request sample sends a message.

Note: The sample process definitions created using RUNMQSC and the amqscos0.tst file trigger the C samples. Change the process definitions in amqscos0.tst and use RUNMQSC with this updated file to use COBOL versions.

Figure 22 on page 142 demonstrates how to use the Request and Inquire samples together.

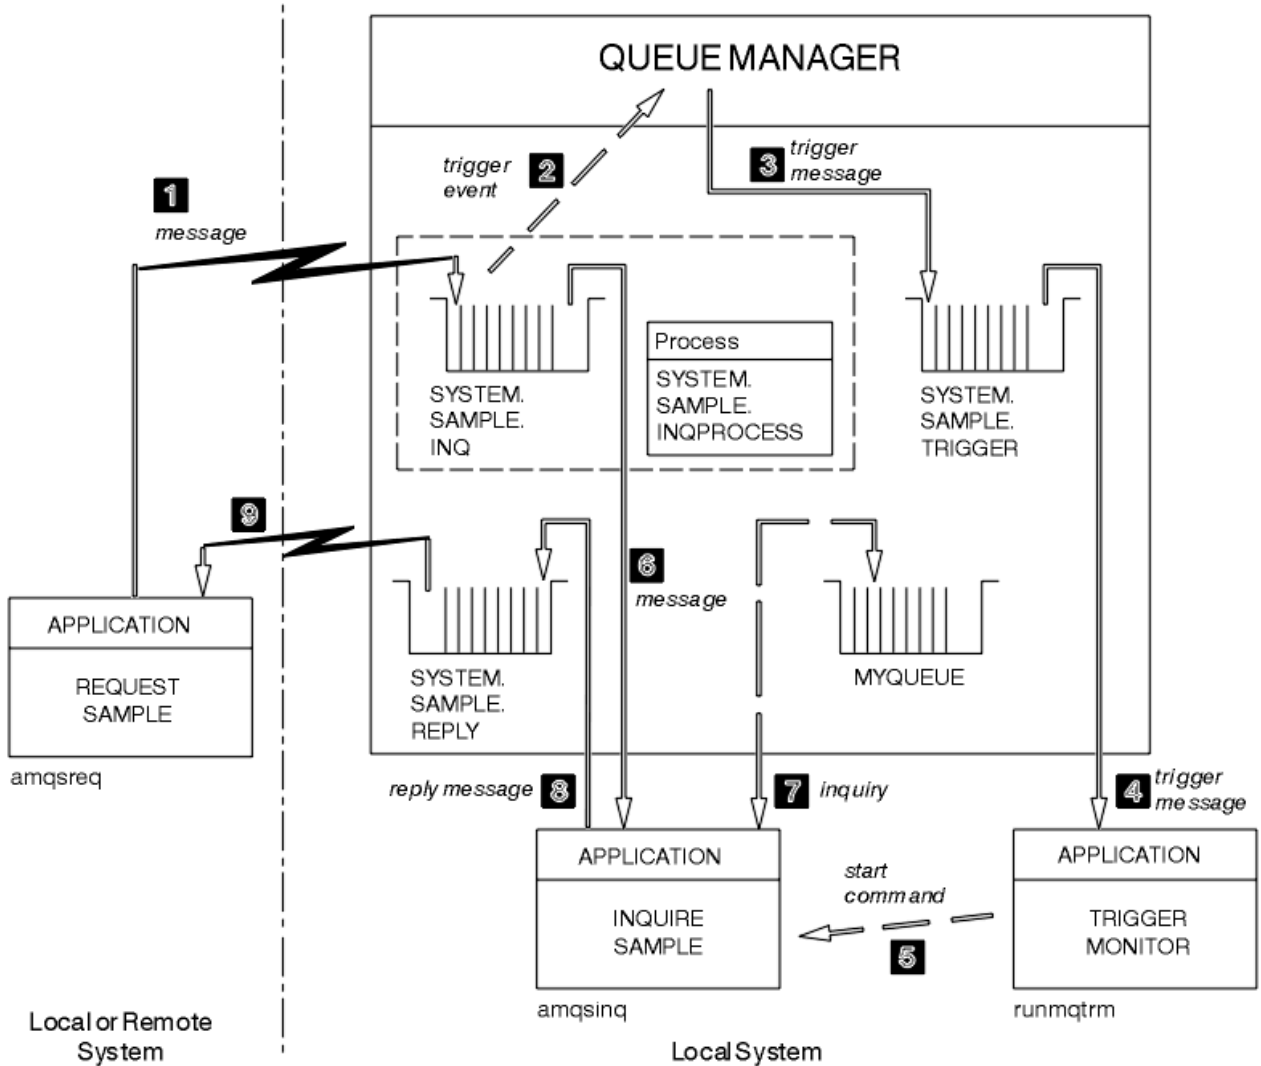


Figure 22. Request and Inquire samples using triggering

In Figure 22 on page 142 the Request sample puts messages onto the target server queue, SYSTEM.SAMPLE.INQ, and the Inquire sample queries the queue, MYQUEUE. Alternatively, you can use one of the sample queues defined when you ran amqscos0.tst, or any other queue that you have defined, for the Inquire sample.

Note: The numbers in Figure 22 on page 142 show the sequence of events.

To run the Request and Inquire samples, using triggering:

1. Check that the queues that you want to use are defined. Run amqscos0.tst, to define the sample queues, and define a queue MYQUEUE.
2. Run the trigger monitor command RUNMQTRM:

```
RUNMQTRM -m qmanagername -q SYSTEM.SAMPLE.TRIGGER
```

3. Run the request sample

```
amqsreq SYSTEM.SAMPLE.INQ
```

Note: The process object defines what is to be triggered. If the client and server are not running on the same platform, any processes started by the trigger monitor must define *ApplType*, otherwise the server takes its default definitions (that is, the type of application that is normally associated with the server machine) and causes a failure.

For a list of application types, see [ApplType](#).

4. Enter the name of the queue that you want the Inquire sample to use:

```
MYQUEUE
```

5. Enter a blank line (to end the Request program).

6. The request sample will then display a message, containing the data the Inquire program obtained from MYQUEUE.

You can use more than one queue; in this case, enter the names of the other queues at step “4” on [page 143](#).

For more information on triggering see [“Starting IBM WebSphere MQ applications using triggers”](#) on [page 315](#).

Design of the Request sample program

The program opens the target server queue so that it can put messages. It uses the MQOPEN call with the MQOO_OUTPUT option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

The program then opens the reply-to queue called SYSTEM.SAMPLE.REPLY so that it can get reply messages. For this, the program uses the MQOPEN call with the MQOO_INPUT_EXCLUSIVE option. If it cannot open the queue, the program displays an error message containing the reason code returned by the MQOPEN call.

For each line of input, the program then reads the text into a buffer and uses the MQPUT call to create a request message containing the text of that line. On this call the program uses the MQRO_EXCEPTION_WITH_DATA report option to request that any report messages sent about the request message will include the first 100 bytes of the message data. The program continues until either it reaches the end of the input or the MQPUT call fails.

The program then uses the MQGET call to remove reply messages from the queue, and displays the data contained in the replies. The MQGET call uses the MQGMO_WAIT, MQGMO_CONVERT, and MQGMO_ACCEPT_TRUNCATED options. The *WaitInterval* is 5 minutes in the COBOL version, and 1 minute in the C version, for the first reply (to allow time for a server application to be triggered), and 15 seconds for subsequent replies. The program waits for these periods if there is no message on the queue. If no message arrives before this interval expires, the call fails and returns the MQRC_NO_MSG_AVAILABLE reason code. The call also uses the MQGMO_ACCEPT_TRUNCATED_MSG option, so that messages longer than the declared buffer size are truncated.

The program demonstrates how to clear the *MsgId* and *CorrelId* fields of the MQMD structure after each MQGET call because the call sets these fields to the values contained in the message it retrieves. Clearing these fields means that successive MQGET calls retrieve messages in the order in which the messages are held in the queue.

The program continues until either the MQGET call returns the MQRC_NO_MSG_AVAILABLE reason code or the MQGET call fails. If the call fails, the program displays an error message that contains the reason code.

The program then closes both the target server queue and the reply-to queue using the MQCLOSE call.

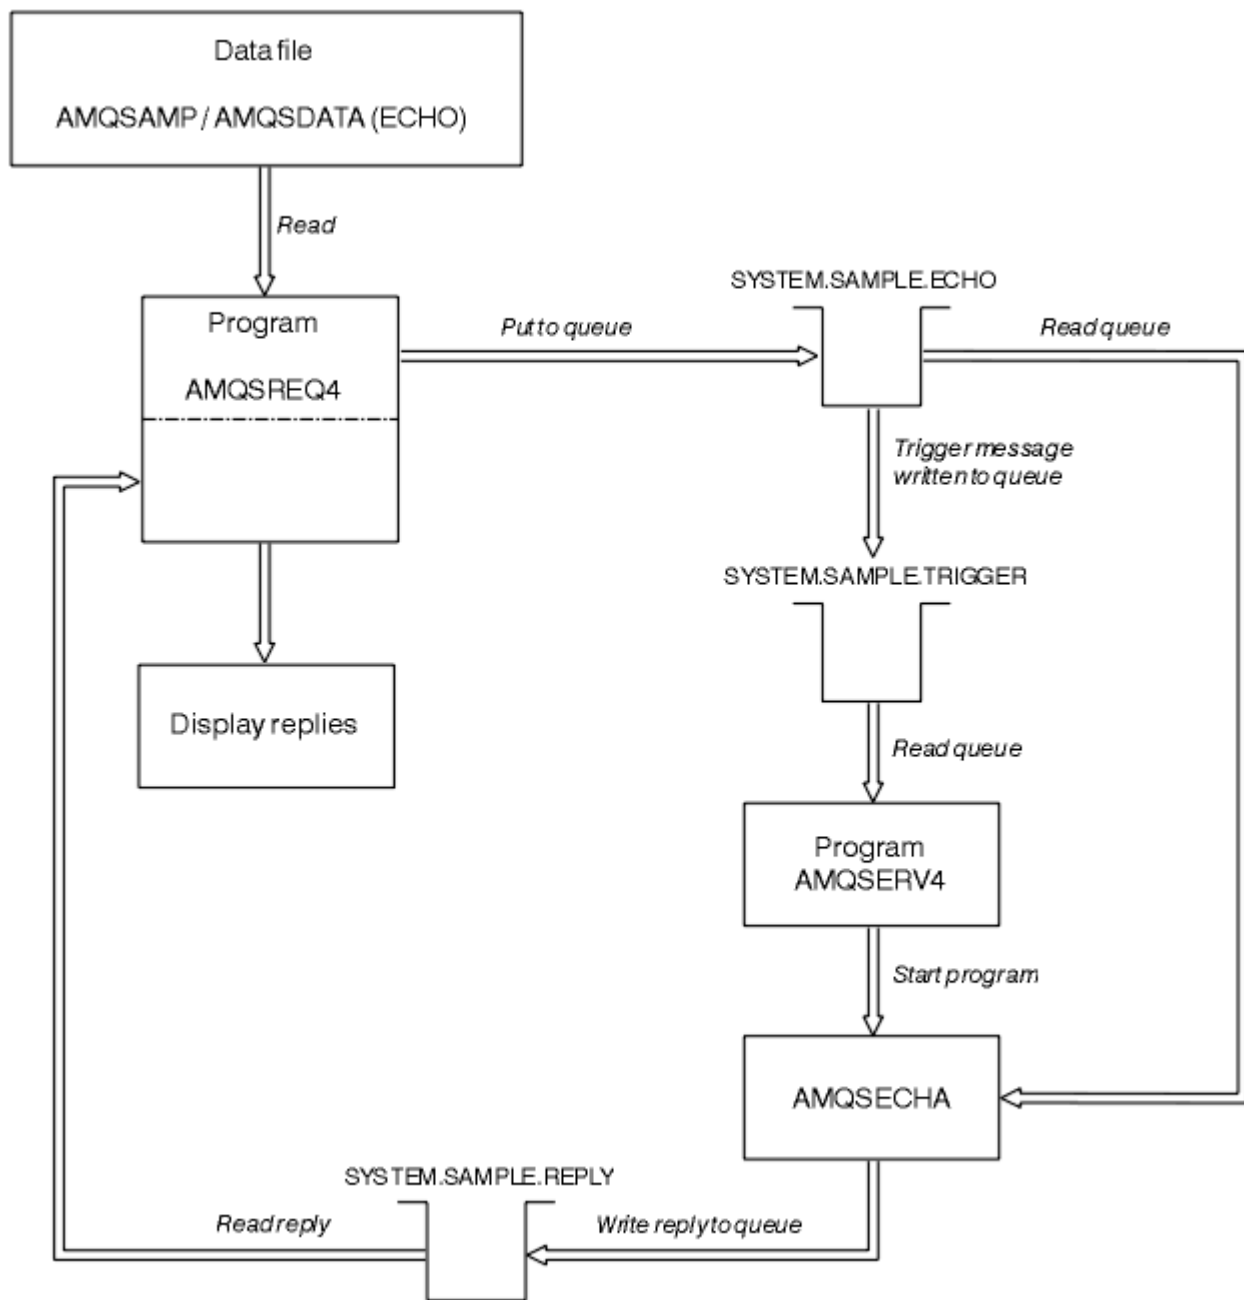


Figure 23. Sample IBM i Client/Server (Echo) program flowchart

The Set sample programs

The Set sample programs inhibit put operations on a queue by using the MQSET call to change the queue's *InhibitPut* attribute. Also, learn about the design of Set sample programs.

See [“Features demonstrated in the sample programs”](#) on page 93 for the names of these programs.

The programs are intended to run as triggered programs, so their only input is an MQTMC2 (trigger message) structure that contains the name of a target queue with attributes that are to be inquired upon. The C version also uses the queue manager name. The COBOL version uses the default queue manager.

For the triggering process to work, ensure that the Set sample program that you want to use is triggered by messages arriving on queue SYSTEM.SAMPLE.SET. To do this, specify the name of the Set sample program that you want to use in the *AppLicId* field of the process definition SYSTEM.SAMPLE.SETPROCESS. The sample queue has a trigger type of FIRST; if there are already

messages on the queue before you run the Request sample, the Set sample is not triggered by the messages that you send.

When you have set the definition correctly:

- For UNIX, Linux and Windows systems, start the **runmqtrm** program in one session, then start the **amqsreq** program in another.
- For IBM i, start the **AMQSERV4** program in one session, then start the **AMQSREQ4** program in another. You could use **AMQSTRG4** instead of **AMQSERV4**, but potential job submission delays could make it less easy to follow what is happening.

Use the Request sample programs to send request messages, each containing just a queue name, to queue **SYSTEM.SAMPLE.SET**. For each request message, the Set sample programs send a reply message containing a confirmation that put operations have been inhibited on the specified queue. The replies are sent to the reply-to queue specified in the request message.

Design of the Set sample program

The program opens the queue named in the trigger message structure that it was passed when it started. (For clarity, we will call this the *request queue*.) The program uses the **MQOPEN** call to open this queue for shared input.

The program uses the **MQGET** call to remove messages from this queue. This call uses the **MQGMO_ACCEPT_TRUNCATED_MSG** and **MQGMO_WAIT** options, with a wait interval of 5 seconds. The program tests the descriptor of each message to see if it is a request message; if it is not, the program discards the message, and displays a warning message.

For each request message removed from the request queue, the program reads the name of the queue (which we will call the *target queue*) contained in the data and opens that queue using the **MQOPEN** call with the **MQOO_SET** option. The program then uses the **MQSET** call to set the value of the *InhibitPut* attribute of the target queue to **MQQA_PUT_INHIBITED**.

If the **MQSET** call is successful, the program uses the **MQPUT1** call to put a reply message on the reply-to queue. This message contains the string **PUT inhibited**.

If the **MQOPEN** or **MQSET** call is unsuccessful, the program uses the **MQPUT1** call to put a report message on the reply-to queue. In the *Feedback* field of the message descriptor of this report message is the reason code returned by either the **MQOPEN** or **MQSET** call, depending on which one failed.

After the **MQSET** call, the program closes the target queue using the **MQCLOSE** call.

When there are no messages remaining on the request queue, the program closes that queue and disconnects from the queue manager.

The SSL/TLS sample program

AMQSSLC is a sample C program that demonstrates how to use the **MQCNO** and **MQSCO** structures to supply SSL/TLS client connection information on the **MQCONN** call. This enables a client MQI application to provide the definition of its client connection channel and SSL/TLS settings at run time without a client channel definition table (CCDT).

If a connection name is supplied, the program constructs a client connection channel definition in an **MQCD** structure.

If the stem name of the key repository file is supplied, the program constructs an **MQSCO** structure; if an OCSP responder URL is also supplied, the program constructs an authentication information record **MQAIR** structure.

The program then connects to the queue manager using **MQCONN**. It inquires and prints out the name of the queue manager to which it connected.

This program is intended to be linked as an MQI client application. However, it can be linked as a regular MQI application. Then, it simply connects to a local queue manager and ignores the client connection information

AMQSSLSC accepts the following parameters, all of which are optional:

-m QmgrName

Name of the queue manager to connect to

-c ChannelName

Name of the channel to use

-x ConnName

Server connection name

SSL/TLS parameters:

-k KeyReposStem

The stem name of the key repository file. This is the full path to the file without the .kdb suffix. For example:

```
/home/user/client  
C:\User\client
```

-s CipherSpec

The SSL/TLS channel CipherSpec string corresponding to the SSLCIPH on the SVRCONN channel definition on the queue manager.

-f

Specifies that only FIPS 140-2 certified algorithms must be used.

-b VALUE1[,VALUE2...]

Specifies that only Suite B compliant algorithms must be used. This parameter is a comma-separated list of one or more of the following values: NONE,128_BIT,192_BIT. These values have the same meaning as those for the MQSUITEB environment variable, and the equivalent EncryptionPolicySuiteB setting in the client configuration file SSL stanza.

-p Policy

Specifies the certificate validation policy to be used. This can be one of the following values:

ANY

Apply each of the certificate validation policies supported by the secure sockets library and accept the certificate chain if any of the policies considers the certificate chain valid. This setting can be used for maximum backwards compatibility with older digital certificates which do not comply with the modern certificate standards.

RFC5280

Apply only the RFC 5280 compliant certificate validation policy. This setting provides stricter validation than the ANY setting, but rejects some older digital certificates.

The default value is ANY.

OCSP certificate revocation parameter:

-o URL

The OCSP Responder URL

Running the SSL/TLS sample program

To run the SSL/TLS sample program you must first set up your SSL or TLS environment. You then run the sample from the command line, supplying a number of parameters.

About this task

The following instructions run the sample program using personal certificates. By varying the command you can, for example, use CA certificates and check their status using an OCSP responder. See the instructions within the sample.

Procedure

1. Create a queue manager with the name QM1. For more information, see [crtmqm](#).

2. Create a key repository for the queue manager. For more information, see [Setting up a key repository on UNIX, Linux, and Windows systems](#).
3. Create a key repository for the client. Call it *clientkey.kdb*.
4. Create a personal certificate for the queue manager. For more information, see [Creating a self-signed personal certificate on UNIX, Linux, and Windows systems](#).
5. Create a personal certificate for the client.
6. Extract the personal certificate from the server key repository and add it to the client repository. For more information, see [Extracting the public part of a self-signed certificate from a key repository on UNIX, Linux and Windows systems](#), and [Adding a CA certificate \(or the public part of a self-signed certificate\) into a key repository, on UNIX, Linux or Windows systems](#).
7. Extract the personal certificate from the client key repository and add it to the server key repository.
8. Create a server connection channel using the MQSC command:

```
DEFINE CHANNEL(QM1SVRCONN) CHLTYPE(SVRCONN) TRPTYPE(TCP) SSLCIPH(NULL_SHA)
```

For more information see [Server-connection channel](#)

9. Define and start a channel listener on the queue manager. For more information see [DEFINE LISTENER](#) and [START LISTENER](#).
10. Run the sample program using the following command:

```
AMQSSSLC -m QM1 -c QM1SVRCONN -x localhost  
-k "c:\Program Files\IBM\WebSphere MQ\clientkey" -s NULL_SHA  
-o http://dummy.OCSP.responder
```

Results

The sample program performs the following actions:

1. Connects to any specified queue manager, or to the default queue manager, using any options specified.
2. Opens the queue manager and inquires on its name.
3. Closes the queue manager.
4. Disconnects from the queue manager.

If the sample program runs successfully, it displays output similar to the following example:

```
Sample AMQSSSLC start  
Connecting to queue manager QM1  
Using the server connection channel QM1SVRCONN  
on connection name localhost.  
Using SSL CipherSpec NULL_SHA  
Using SSL key repository stem c:\Program Files\IBM\WebSphere MQ\clientkey  
Using OCSP responder URL http://dummy.OCSP.responder  
Connection established to queue manager QM1
```

```
Sample AMQSSSLC end
```

If the sample program encounters a problem, it displays an appropriate error message, for example if you specify an invalid OCSP responder URL, you receive the following message:

```
MQCONN ended with reason code 2553
```

For a list of reason codes see [API reason codes](#).

The Triggering sample programs

The function provided in the triggering sample is a subset of that provided in the trigger monitor in the **runmqtrm** program.

See [“Features demonstrated in the sample programs” on page 93](#) for the names of these programs.

Design of the triggering sample

The triggering sample program opens the initiation queue using the MQOPEN call with the MQOO_INPUT_AS_Q_DEF option. It gets messages from the initiation queue using the MQGET call with the MQGMO_ACCEPT_TRUNCATED_MSG and MQGMO_WAIT options, specifying an unlimited wait interval. The program clears the *MsgId* and *CorrelId* fields before each MQGET call to get messages in sequence.

When it has retrieved a message from the initiation queue, the program tests the message by checking the size of the message to make sure that it is the same size as an MQTM structure. If this test fails, the program displays a warning.

For valid trigger messages, the triggering sample copies data from these fields: *ApplicId*, *EnvrData*, *Version*, and *ApplType*. The last two of these fields are numeric, so the program creates character replacements to use in an MQTMC2 structure for UNIX, Linux, and Windows systems.

The triggering sample issues a start command to the application specified in the *ApplicId* field of the trigger message, and passes an MQTMC2 or MQTMC (a character version of the trigger message) structure. In UNIX, Linux and Windows systems, the *EnvrData* field is used as an extension to the invoking command string.

Finally, the program closes the initiation queue.

Running the Triggering sample programs

This topic contains information about running Triggering sample programs.

Running the amqstrg0.c, amqstrg, and amqstrgc samples

The program takes 2 parameters:

1. The name of the initiation queue (necessary)
2. The name of the queue manager (optional)

If a queue manager is not specified, it connects to the default one. A sample initiation queue will have been defined when you ran amqscos0.tst; the name of that queue is SYSTEM.SAMPLE.TRIGGER, and you can use it when you run this program.

Note: The function in this sample is a subset of the full triggering function that is supplied in the **runmqtrm** program.

Design of the trigger server

The design of the trigger server is similar to that of the trigger monitor, except that the trigger server:

- Allows MQAT_CICS as well as MQAT_OS400 applications
- For CICS applications, substitutes the *EnvrData*, for example, to specify the CICS region, from the trigger message in the STRCICSUSR command
- Opens the initiation queue for shared input, so that many trigger servers can run at the same time

Note: Programs started by AMQSERV4 must not use the MQDISC call because this stops the trigger server. If programs started by AMQSERV4 use the MQCONN call, they get the MQRC_ALREADY_CONNECTED reason code.

TUXEDO samples

Learn about the Put and Get sample programs for TUXEDO, and building the server environment in TUXEDO.

Before running these samples, you must build the server environment.

Note: Throughout this topic, the backslash (\) character is used to split long commands over more than one line. Do not enter this character. Enter each command as a single line.

Building the server environment

Information about building the server environment for WebSphere MQ for different platforms.

It is assumed that you have a working TUXEDO environment.

Building the server environment for WebSphere MQ for AIX (32-bit)

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO, and MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed:

```
$ export CFLAGS="-I MQ_INSTALLATION_PATH/inc -I /<APPDIR> -L MQ_INSTALLATION_PATH/lib"
$ export LDOPTS="-lmqm"
$ export FIELDTBLS=MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=/<<APPDIR>/amqstxvx.V
$ export LIBPATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:lib
```

3. Add the following to the TUXEDO file udataobj/RM

```
MQSeries_XA_RMI:MQRMIASwitchDynamic: -lmqmx -lmqm
```

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
$ buildtms -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.a \
-r MQSeries_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.a \
-r MQSeries_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.a
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.a
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmlodcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> cirdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ stimqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for AIX (64-bit)

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR represents the root directory for TUXEDO, and MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.:

```
$ export CFLAGS="-I MQ_INSTALLATION_PATH/inc -I /<APPDIR> -L  
MQ_INSTALLATION_PATH/lib64"  
$ export LDOPTS="-lmqm"  
$ export FIELDTBLS=MQ_INSTALLATION_PATH/samp/amqstxvx.flds  
$ export VIEWFILES=/<APPDIR>/amqstxvx.V  
$ export LIBPATH=$TUXDIR/lib64:MQ_INSTALLATION_PATH/lib64:lib64
```

3. Add the following to the TUXEDO file udataobj/RM

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: -lmqmx64 -lmqm
```

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds  
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v  
$ buildtms -o MQXA -r MQSeries_XA_RMI  
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a \  
-r MQSeries_XA_RMI -s MPUT1:MPUT \  
-s MGET1:MGET \  
-v -bshm  
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a \  
-r MQSeries_XA_RMI -s MPUT2:MPUT  
-s MGET2:MGET \  
-v -bshm  
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a  
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \  
-f MQ_INSTALLATION_PATH/lib64/libmqm.a
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmlodcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ stimqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for Solaris (32-bit)

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

1. Create a directory (for example, `APPDIR`) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where `TUXDIR` is the root directory for TUXEDO:

```
$ export CFLAGS="-I /APPDIR"
$ export FIELDTBLS=amqstxvx.flds
$ export VIEWFILES=amqstxvx.V
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:/lib
$ export LD_LIBRARY_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:/lib
```

3. Add the following to the TUXEDO file `udataobj/RM`.

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib/libmqmxa.a MQ_INSTALLATION_PATH/lib/libmqm.so \
/opt/tuxedo/lib/libtux.a
```

4. Run the commands:

```
$ mkfldhdr amqstxvx.flds
$ viewc amqstxvx.v
$ buildtms -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSeries_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
-l -ldl
$ buildserver -o MQSERV2 -f amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSeries_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
-l -ldl
$ buildclient -o doputs -f amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
$ buildclient -o dogets -f amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so
```

5. Edit `ubbstxcx.cfg` and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmlloadcf -y ubbstxcx.cfg
```

6. Create the `TLOGDEVICE`:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> cd1 -z /APPDIR/TLOG1
```

7. Start the queue manager:

```
$ stimqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for Solaris (64-bit)

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-I /<APPDIR>"
$ export FIELDTBLS=amqstxvx.flds
$ export VIEWFILES=amqstxvx.V
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:lib64
$ export LD_LIBRARY_PATH=$TUXDIR/lib64:MQ_INSTALLATION_PATH/lib64:lib64
```

3. Add the following to the TUXEDO file udataobj/RM.

```
MQSeries_XA_RMI:MORMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib64/libmqmxa64.a MQ_INSTALLATION_PATH/lib64/libmqm.so \
/opt/tuxedo/lib64/libtux.a
```

4. Run the commands:

```
$ mkfldhdr    amqstxvx.flds
$ viewc      amqstxvx.v
$ buildtms    -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f amqstxsx.c \
    -f MQ_INSTALLATION_PATH/lib64/libmqm.so \
    -r MQSeries_XA_RMI -s MPUT1:MPUT \
    -s MGET1:MGET \
    -v -bshm
    -l -ldl
$ buildserver -o MQSERV2 -f amqstxsx.c \
    -f MQ_INSTALLATION_PATH/lib64/libmqm.so \
    -r MQSeries_XA_RMI -s MPUT2:MPUT \
    -s MGET2:MGET \
    -v -bshm
    -l -ldl
$ buildclient -o doputs -f amqstxpx.c \
    -f MQ_INSTALLATION_PATH/lib64/libmqm.so \
$ buildclient -o dogets -f amqstxgx.c \
    -f MQ_INSTALLATION_PATH/lib64/libmqm.so
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmlodacf -y ubbstxcx.cfg
```

6. Create the TLOGDEVICE:


```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ stmqm
```

8. Start Tuxedo:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for HP-UX (32-bit)

Note: The 32-bit TUXEDO server environment can only be built on the Itanium platform.

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-Aa -D_HPUX_SOURCE"
$ export FIELDTBLs=MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=$APPDIR/amqstxvx.V
$ export TUXCONFIG=$APPDIR/tuxconfig
$ export PATH=$TUXDIR/bin:/usr/bin:/sbin:MQ_INSTALLATION_PATH/bin:$PATH
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib:/lib
$ export FLDTBLDIR=$APPDIR:$TUXDIR/udataobj
```

3. Add the following to the TUXEDO file udataobj/RM

```
MQSeries_XA_RMI:MQRMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib/libmqmxa.so MQ_INSTALLATION_PATH/lib/libmqm.so \
/opt/tuxedo/lib/libtux.sl
```

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
```

After running the mkfldhdr and viewc commands, the amqstxvx.h header file is created in the TUXEDO application directory. Copy this file from the TUXEDO application directory into the TUXEDO include directory, and then run the following commands.

```
$ buildtms -o MQXA -r MQSeries_XA_RMI
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSeries_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so \
-r MQSeries_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so
```

```
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib/libmqm.so
```

5. Edit ubbstxcx.cfg and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ stimqm
```

8. Start TUXEDO:

```
$ tmboot -y
```

You can now use the doputs and dogets programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for HP-UX (64-bit)

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

1. Create a directory (for example, <APPDIR>) in which the server environment is built and execute all commands in this directory.
2. Export the following environment variables, where TUXDIR is the root directory for TUXEDO:

```
$ export CFLAGS="-Aa -D_HPUX_SOURCE"
$ export FIELDTBL=MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ export VIEWFILES=$APPDIR/amqstxvx.V
$ export TUXCONFIG=$APPDIR/tuxconfig
$ export PATH=$TUXDIR/bin:/usr/bin:/sbin:MQ_INSTALLATION_PATH/bin:$PATH
$ export SHLIB_PATH=$TUXDIR/lib:MQ_INSTALLATION_PATH/lib64:/lib64
$ export FLDTBLDIR=$APPDIR:$TUXDIR/udataobj
```

3. Add the following to the TUXEDO file udataobj/RM

On the HP-UX IA64 (IPF) platform:

```
MQSeries_XA_RMI:MORMIXASwitchDynamic: \
MQ_INSTALLATION_PATH/lib64/libmqmxa64.so MQ_INSTALLATION_PATH/lib64/libmqm.so \
/opt/tuxedo/lib/libtux.sl
```

Note: The WebSphere MQ libraries shipped on the HP-UX IA64 (IPF) platform have a .so file name extension.

4. Run the commands:

```
$ mkfldhdr MQ_INSTALLATION_PATH/samp/amqstxvx.flds
$ viewc MQ_INSTALLATION_PATH/samp/amqstxvx.v
```

After running the `mkfldhdr` and `viewc` commands, the `amqstxvx.h` header file is created in the TUXEDO application directory. Copy this file from the TUXEDO application directory into the TUXEDO include directory, and then run the following commands.

```
$ buildtms -o MQXA -r MQSeries_XA_RMI
```

On the HP-UX IA64 (IPF) platform:

```
$ buildserver -o MQSERV1 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-r MQSeries_XA_RMI -s MPUT1:MPUT \
-s MGET1:MGET \
-v -bshm
$ buildserver -o MQSERV2 -f MQ_INSTALLATION_PATH/samp/amqstxsx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so \
-r MQSeries_XA_RMI -s MPUT2:MPUT \
-s MGET2:MGET \
-v -bshm
$ buildclient -o doputs -f MQ_INSTALLATION_PATH/samp/amqstxpx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so
$ buildclient -o dogets -f MQ_INSTALLATION_PATH/samp/amqstxgx.c \
-f MQ_INSTALLATION_PATH/lib64/libmqm.so
```

5. Edit `ubbstxcx.cfg` and add details of the machine name, working directories, and queue manager as necessary:

```
$ tmloadcf -y MQ_INSTALLATION_PATH/samp/ubbstxcx.cfg
```

6. Create the TLOGDEVICE:

```
$tmadmin -c
```

A prompt is displayed. At this prompt, enter:

```
> crdl -z /<APPDIR>/TLOG1
```

7. Start the queue manager:

```
$ stmqm
```

8. Start TUXEDO:

```
$ tmboot -y
```

You can now use the `doputs` and `dogets` programs to put messages to a queue and retrieve them from a queue.

Building the server environment for WebSphere MQ for Windows (32-bit)

Note: Change the fields identified by `<>` in the following, to the directory paths:

<code><MQMDIR></code>	the directory path specified when WebSphere MQ was installed, for example g:\Program Files\IBM\WebSphere MQ
<code><TUXDIR></code>	the directory path specified when TUXEDO was installed, for example f:\tuxedo
<code><APPDIR></code>	the directory path to be used for the sample application, for example f:\tuxedo\apps\mqapp

To build the server environment and samples:

1. Create an application directory in which to build the sample application, for example:

```
f:\tuxedo\apps\mqapp
```

2. Copy the following sample files from the WebSphere MQ sample directory to the application directory:

```
amqstxmn.mak  
amqstxen.env  
ubbstxcn.cfg
```

3. Edit each of these files to set the directory names and directory paths used on your installation.
4. Edit ubbstxcn.cfg (see [Figure 24 on page 157](#)) to add details of the machine name and the queue manager that you want to connect to.
5. Add the following line to the TUXEDO file <TUXDIR>udataobj\rm

```
MQSeries_XA_RMI;MQRMIXASwitchDynamic;  
<MQMDIR>\tools\lib\mqmxa.lib <MQMDIR>\tools\lib\mqm.lib
```

where <MQMDIR> is replaced as shown in the previous example. Although shown here as two lines, the new entry must be one line in the file.

6. Set the following environment variables:

```
TUXDIR=<TUXDIR>  
TUXCONFIG=<APPDIR>\tuxconfig  
FIELDTBLS=<MQMDIR>\tools\c\samples\amqstxvx.fld  
LANG=C
```

7. Create a TLOG device for TUXEDO. To do this, invoke `tmadmin -c`, and enter the command:

```
crdl -z <APPDIR>\TLOG
```

where <APPDIR> is replaced.

8. Set the current directory to <APPDIR>, and invoke the sample makefile (amqstxmn.mak) as an external project makefile. For example, with Microsoft Visual C++ , issue the command:

```
msvc amqstxmn.mak
```

Select **build** to build all the sample programs.

```

*RESOURCES
IPCKEY      99999
UID         0
GID         0
MAXACCESSERS 20
MAXSERVERS  20
MAXSERVICES 50
MASTER     SITE1
MODEL       SHM
LDBAL       N

*MACHINES
<MachineName> LMID=SITE1
               TUXDIR="f:\tuxedo"
               APPDIR="f:\tuxedo\apps\mqapp;g:\Program Files\IBM\WebSphere MQ\bin"
               ENVFILE="f:\tuxedo\apps\mqapp\amqstxen.env"
               TUXCONFIG="f:\tuxedo\apps\mqapp\tuxconfig"
               ULOGPFX="f:\tuxedo\apps\mqapp\ULOG"
               TLOGDEVICE="f:\tuxedo\apps\mqapp\TLOG"
               TLOGNAME=TLOG
               TYPE="i386NT"
               UID=0
               GID=0

*GROUPS
GROUP1      LMID=SITE1  GRPNO=1
             TMSNAME=MQXA
             OPENINFO="MQSeries_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT: CLOPT="-A -- -m MYQUEUEMANAGER"

MQSERV1     SRVGRP=GROUP1 SRVID=1
MQSERV2     SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2

```

Figure 24. Example of ubbstxcn.cfg file for WebSphere MQ for Windows

Note: Change the directory names and directory paths to match your installation. Also change the queue manager name MYQUEUEMANAGER to the name of the queue manager that you want to connect to. Other information that you need to add is identified by <> characters.

The sample ubbconfig file for WebSphere MQ for Windows is listed in [Figure 24 on page 157](#). It is supplied as ubbstxcn.cfg in the WebSphere MQ samples directory.

The sample makefile (see [Figure 25 on page 158](#)) supplied for WebSphere MQ for Windows is called ubbstxmnmak, and is held in the WebSphere MQ samples directory.

```

TUXDIR   = f:\tuxedo
MQMDIR   = g:\Program Files\IBM\WebSphere MQ
APPDIR   = f:\tuxedo\apps\mqapp
MQMLIB   = $(MQMDIR)\tools\lib
MQMINC   = $(MQMDIR)\tools\c\include
MQMSAMP  = $(MQMDIR)\tools\c\samples
INC      = -f "-I$(MQMINC) -I$(APPDIR)"
DBG      = -f "/Zi"

amqstx.exe:
$(TUXDIR)\bin\mkfldhdr -d$(APPDIR) $(MQMSAMP)\amqstvx.fld
$(TUXDIR)\bin\viewc -d$(APPDIR) $(MQMSAMP)\amqstvx.v
$(TUXDIR)\bin\buildtms -o MQXA -r MQSeries_XA_RMI
$(TUXDIR)\bin\buildserver -o MQSERV1 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSeries_XA_RMI \
-s MPUT1:MPUT -s MGET1:MGET
$(TUXDIR)\bin\buildserver -o MQSERV2 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSeries_XA_RMI \
-s MPUT2:MPUT -s MGET2:MGET
$(TUXDIR)\bin\buildclient -o doputs -f $(MQMSAMP)\amqstxpx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)
$(TUXDIR)\bin\buildclient -o dogets -f $(MQMSAMP)\amqstngx.c \
-f $(MQMLIB)\mqm.lib $(INC) -v $(DBG)
$(TUXDIR)\bin\tmloadcf -y $(APPDIR)\ubbstxcn.cfg

```

Figure 25. Sample TUXEDO makefile for WebSphere MQ for Windows

Building the server environment for WebSphere MQ for Windows (64-bit)

Note: Change the fields identified by <> in the following, to the directory paths:

<MQMDIR>	the directory path specified when WebSphere MQ was installed, for example g:\Program Files\IBM\WebSphere MQ
<TUXDIR>	the directory path specified when TUXEDO was installed, for example f:\tuxedo
<APPDIR>	the directory path to be used for the sample application, for example f:\tuxedo\apps\mqapp

To build the server environment and samples:

1. Create an application directory in which to build the sample application, for example:

```
f:\tuxedo\apps\mqapp
```

2. Copy the following sample files from the WebSphere MQ sample directory to the application directory:

```
amqstxmn.mak
amqstxen.env
ubbstxcn.cfg
```

3. Edit each of these files to set the directory names and directory paths used on your installation.
4. Edit ubbstxcn.cfg (see Figure 26 on page 159) to add details of the machine name and the queue manager that you want to connect to.
5. Add the following line to the TUXEDO file <TUXDIR>udataobj\rm

```
MQSeries_XA_RMI;MQRMIXASwitchDynamic;
<MQMDIR>\tools\lib64\mqmxa64.lib <MQMDIR>\tools\lib64\mqm.lib
```

where <MQMDIR> is replaced. Although shown here as two lines, the new entry must be one line in the file.

6. Set the following environment variables:

```
TUXDIR=<TUXDIR>
TUXCONFIG=<APPDIR>\tuxconfig
FIELDTBLS=<MQMDIR>\tools\c\samples\amqstvx.fld
LANG=C
```

7. Create a TLOG device for TUXEDO. To do this, invoke `tmadmin -c`, and enter the command:

```
crdl -z <APPDIR>\TLOG
```

where `<APPDIR>` is replaced as shown in the previous example.

8. Set the current directory to `<APPDIR>`, and invoke the sample makefile (`amqstxmn.mak`) as an external project makefile. For example, with Microsoft Visual C++ , issue the command:

```
msvc amqstxmn.mak
```

Select **build** to build all the sample programs.

```
*RESOURCES
IPCKEY      99999
UID         0
GID         0
MAXACCESSERS 20
MAXSERVERS  20
MAXSERVICES 50
MASTER      SITE1
MODEL        SHM
LDBAL        N

*MACHINES
<MachineName> LMID=SITE1
               TUXDIR="f:\tuxedo"
               APPDIR="f:\tuxedo\apps\mqapp;g:\Program Files\IBM\WebSphere MQ\bin"
               ENVFILE="f:\tuxedo\apps\mqapp\amqstxen.env"
               TUXCONFIG="f:\tuxedo\apps\mqapp\tuxconfig"
               ULOGPFX="f:\tuxedo\apps\mqapp\ULOG"
               TLOGDEVICE="f:\tuxedo\apps\mqapp\TLOG"
               TLOGNAME=TLOG
               TYPE="i386NT"
               UID=0
               GID=0

*GROUPS
GROUP1
      LMID=SITE1 GRPNO=1
      TMSNAME=MQXA
      OPENINFO="MQSeries_XA_RMI:MYQUEUEMANAGER"

*SERVERS
DEFAULT: CLOPT="-A -- -m MYQUEUEMANAGER"

MQSERV1      SRVGRP=GROUP1 SRVID=1
MQSERV2      SRVGRP=GROUP1 SRVID=2

*SERVICES
MPUT1
MGET1
MPUT2
MGET2
```

Figure 26. Example of `ubbstxcn.cfg` file for WebSphere MQ for Windows

Note: Change the directory names and directory paths to match your installation. Also change the queue manager name `MYQUEUEMANAGER` to the name of the queue manager that you want to connect to. Other information that you need to add is identified by `<>` characters.

The sample `ubbconfig` file for WebSphere MQ for Windows is listed in [Figure 26 on page 159](#). It is supplied as `ubbstxcn.cfg` in the WebSphere MQ samples directory.

The sample makefile (see Figure 27 on page 160) supplied for WebSphere MQ for Windows is called ubbstxmn.mak, and is held in the WebSphere MQ samples directory.

```
TUXDIR = f:\tuxedo
MQMDIR = g:\Program Files\IBM\WebSphere MQ
APPDIR = f:\tuxedo\apps\mqapp
MQMLIB = $(MQMDIR)\tools\lib64
MQMINC = $(MQMDIR)\tools\c\include
MQMSAMP = $(MQMDIR)\tools\c\samples
INC = -f "-I$(MQMINC) -I$(APPDIR)"
DBG = -f "/Zi"

amqstx.exe:
$(TUXDIR)\bin\mkfldhdr -d$(APPDIR) $(MQMSAMP)\amqstxvx.fld
$(TUXDIR)\bin\viewc -d$(APPDIR) $(MQMSAMP)\amqstxvx.v
$(TUXDIR)\bin\builddtms -o MQXA -r MQSeries_XA_RMI
$(TUXDIR)\bin\buildserver -o MQSERV1 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSeries_XA_RMI \
-s MPUT1:MPUT -s MGET1:MGET
$(TUXDIR)\bin\buildserver -o MQSERV2 -f $(MQMSAMP)\amqstxsx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG) \
-r MQSeries_XA_RMI \
-s MPUT2:MPUT -s MGET2:MGET
$(TUXDIR)\bin\buildclient -o doputs -f $(MQMSAMP)\amqstxpx.c \
-f $(MQMLIB)\mqm.lib -v $(INC) $(DBG)
$(TUXDIR)\bin\buildclient -o dogets -f $(MQMSAMP)\amqstxgx.c \
-f $(MQMLIB)\mqm.lib $(INC) -v $(DBG)
$(TUXDIR)\bin\tmloadcf -y $(APPDIR)\ubbstxcn.cfg
```

Figure 27. Sample TUXEDO makefile for WebSphere MQ for Windows

Sample server program for TUXEDO

The sample server program (amqstxsx) is designed to run with the Put (amqstxpx.c) and the Get (amqstxgx.c) sample programs. The sample server program runs automatically when TUXEDO is started.

Note: You must start your queue manager *before* you start TUXEDO.

The sample server provides two TUXEDO services, MPUT1 and MGET1:

- The MPUT1 service is driven by the PUT sample and uses MQPUT1 in syncpoint to put a message in a unit of work controlled by TUXEDO. It takes the parameters QName and Message Text, which are supplied by the PUT sample.
- The MGET1 service opens and closes the queue each time that it gets a message. It takes the parameters QName and Message Text, which are supplied by the GET sample.

Any error messages, reason codes, and status messages are written to the TUXEDO log file.




Figure 28. How TUXEDO samples work together

Put sample program for TUXEDO

This sample allows you to put a message on a queue multiple times, in batches, demonstrating syncpointing using TUXEDO as the resource manager.

The sample server program `amqstxsx` must be running for the put sample to succeed; the server sample program connects to the queue manager and uses the XA interface. To run the sample enter:

- `doputs -n queueName -b batchSize -c tranCount -t message`

For example:

- `doputs -n myqueue -b 5 -c 6 -t "Hello World"`

This puts 30 messages onto the queue named `myqueue`, in six batches, each with five messages in it. If there are any problems it backs a batch of messages out, otherwise it commits them.

Any error messages are written to the TUXEDO log file and to `stderr`. Any reason codes are written to `stderr`.

Get sample for TUXEDO

This sample allows you to get messages from a queue in batches.

The sample server program `amqstxsx` must be running for the put sample to succeed; the server sample program connects to the queue manager and uses the XA interface. To run the sample enter:

- `dogets -n queueName -b batchSize -c tranCount`

For example:

- `dogets -n myqueue -b 6 -c 4`

This takes 24 messages off the queue named `myqueue`, in six batches, each with four messages in it. If you run this after the `put` example, which puts 30 messages on `myqueue`, you have only six messages on `myqueue`. The number of batches and the batch size can vary between putting the messages and getting them.

Any error messages are written to the TUXEDO log file and to `stderr`. Any reason codes are written to `stderr`.

Using the SSPI security exit on Windows systems

This topic describes how to use the SSPI channel-exit programs on Windows systems. The supplied exit code is in two formats: object and source.

Object code

The object code file is called `amqrspin.dll`. For both client and server, it is installed as a standard part of WebSphere MQ for Windows in the `MQ_INSTALLATION_PATH/exits/INSTALLATION_NAME` folder. For example, `C:\Program Files\IBM\WebSphere MQ\exits\installation2`. It is loaded as a standard user exit. You can run the supplied security channel exit and use authentication services in your definition of the channel.

To do this, specify either of the following:

```
SCYEXIT('amqrspin(SCY_KERBEROS)')
SCYEXIT('amqrspin(SCY_NTLM)')
```

To provide support for a restricted channel, specify the following on the `SVRCONN` channel:

```
SCYDATA('remote_principal_name')
```

where *remote_principal_name* is in the form `DOMAIN\user`. The secure channel is established only if the name of the remote principal matches *remote_principal_name*.

To use the supplied channel-exit programs between systems that operate within a Kerberos security domain, create a `servicePrincipalName` for the queue manager.

Source code

The exit source code file is called `amqsspin.c`. It is in `C:\Program Files\IBM\WebSphere MQ\Tools\c\Samples`.

If you modify the source code, you must recompile the modified source.

You compile and link it in the same way as any other channel exit for the relevant platform, except that SSPI headers need to be accessed at compile time, and the SSPI security libraries, together with any recommended associated libraries, need to be accessed at link time.

Before you execute the following command, make sure that `cl.exe`, and the Visual C++ library and the `include` folder are available in your path. For example:

```
cl /VERBOSE /LD /MT /I<path_to_Microsoft_platform_SDK\include>
/I<path_to_WebSphere MQ\tools\c\include> amqsspin.c /DSECURITY_WIN32
-link /DLL /EXPORT:SCY_KERBEROS /EXPORT:SCY_NTLM STACK:8192
```

Note: The source code does not include any provision for tracing or error handling. If you modify and use the source code, add your own tracing and error-handling routines.

Running the samples using remote queues

You can demonstrate remote queuing by running the samples on connected queue managers.

Program `amqscos0.tst` provides a local definition of a remote queue (`SYSTEM.SAMPLE.REMOTE`) that uses a remote queue manager named `OTHER`. To use this sample definition, change `OTHER` to the name of the second queue manager that you want to use. You must also set up a message channel between your two queue managers; for information on how to do this, see [Defining the channels](#).

The Request sample programs put their own local queue manager name in the *ReplyToQMGr* field of messages that they send. The Inquire and Set samples send reply messages to the queue and message queue manager named in the *ReplyToQ* and *ReplyToQMGr* fields of the request messages that they process.

The Cluster Queue Monitoring sample program (AMQSCLM)

This sample uses the built-in IBM WebSphere MQ cluster workload balancing features to direct messages to instances of queues that have consuming applications attached. This automatic direction prevents the build-up of messages on an instance of a cluster queue to which no consuming application is attached.

Overview

You can set up a cluster that has more than one definition for the same queue on different queue managers. This configuration provides the benefit of increased availability and workload balancing. However, there is no capability built into IBM WebSphere MQ to dynamically modify the distribution of messages across a cluster based on the state of attached applications. For this reason, a consuming application must always be attached to every instance of a queue to ensure that messages are processed.

The cluster queue monitoring sample program monitors the state of attached applications. The program dynamically adjusts the built-in workload balancing configuration to direct messages to instances of a clustered queue with consuming applications attached. In certain situations this program can be used to relax the need for a consuming application to always be connected to every instance of a queue. It also resends messages that become queued on an instance of a queue with no consuming applications attached. Resending messages enables messages to be routed around a consuming application that is temporarily shut down.

The program is designed to be used where the consuming applications are long running applications, rather than frequently attaching and detaching applications.

The cluster queue monitoring sample program is the compiled executable program of the C sample file `amqsc1ma.c`.

Further information about clusters and workload can be found in [Using clusters for workload management](#)

AMQSCLM: Design and Planning for using the sample

Information about how the cluster queue monitoring sample program works, points to consider when setting up a system for the sample program to run on, and modifications that can be made to the sample source code.

Design

The cluster queue monitoring sample program monitors local clustered queues that have consuming applications attached. The program monitors queues specified by the user. The name of the queue might be specific, for example `APP.TEST01`, or generic. Generic names must be in a format that conforms to PCF (Programmable Command Format). Examples of generic names are `APP.TEST*`, or `APP*`.

Each queue manager in a cluster that owns an instance of a local queue to be monitored, requires an instance of the cluster queue monitoring sample program to be connected to it.

Dynamic message routing

The cluster queue monitoring sample program uses the **IPPROCS** (open for input process count) value of a queue to determine whether that queue has any consumers. A value greater than 0 indicates that the queue has at least one consuming application attached. Such queues are active. A value of 0 indicates that the queue has no attached consuming programs. Such queues are inactive.

For a clustered queue with multiple instances in a cluster, WebSphere MQ uses the cluster workload priority property **CLWLPRTY** of each queue instance to determine which instances to send messages to. WebSphere MQ sends messages to the available instances of a queue with the highest **CLWLPRTY** value.

The cluster queue monitoring sample program activates a cluster queue by setting the local **CLWLPRTY** value to 1. The program deactivates a cluster queue by setting its **CLWLPRTY** value to 0.

WebSphere MQ clustering technology propagates the updated **CLWLPRTY** property of a clustered queue to all relevant queue managers in the cluster. For example,

- A queue manager with a connected application that puts messages to the queue.
- A queue manager that owns a local queue of the same name in the same cluster.

The propagation is done using the full repository queue managers of the cluster. New messages for the cluster queue are directed to the instances with the highest **CLWLPRTY** value within the cluster.

Queued message transfer

The dynamic modification of the value of **CLWLPRTY** influences the routing of new messages. This dynamic modification does not affect messages already queued on a queue instance with no attached consumers, or messages that have been through the workload balancing mechanism before a modified **CLWLPRTY** value was propagated across the cluster. As a result, messages remain on any inactive queue and not be processed by a consuming application. To solve this, the cluster queue monitoring sample program is able to get messages from a local queue with no consumers, and send these messages to remote instances of the same queue where consumers are attached.

The cluster queue monitoring sample program transfers messages from an inactive local queue to one or more active remote queues by getting messages (using **MQGET**) and putting messages (using **MQPUT**) to the same clustered queue. This transfer causes the WebSphere MQ cluster workload management to select a different target instance, based on a higher **CLWLPRTY** value than that of the local queue instance. Message persistence and context are preserved during the message transfer. Message order, and any binding options are not preserved.

Planning

The cluster queue monitoring sample program modifies the cluster configuration when there is a change in the connectivity of consuming applications. Modifications are transmitted from the queue managers where the cluster queue monitoring sample program is monitoring queues, to the full repository queue managers in the cluster. The full repository queue managers process the configuration updates and resend them to all relevant queue managers in the cluster. Relevant queue managers include those queue managers that own clustered queues of the same name (where an instance of the cluster queue monitoring sample program is running), and any queue manager where an application opened the cluster queue to put messages to it in the last 30 days.

Changes are asynchronously processed across the cluster. Therefore, after each change, different queue managers in the cluster might have different views of the configuration for a period of time.

The cluster queue monitoring sample program is only suitable for systems where consuming applications infrequently attach or detach; for example, long running consuming applications. When used to monitor systems where consuming applications are only attached for short periods, the latency incurred when distributing the configuration updates might result in queue managers in the cluster having an incorrect view of the queues where consumers are attached. This latency might result in incorrectly routed messages.

When monitoring many queues, a relatively low rate of change in attached consumers across all queues might increase cluster configuration traffic across the cluster. Increased cluster configuration traffic can result in excessive load on one or more of the following queue managers.

- The queue managers where the cluster queue monitoring sample program is running
- The full repository queue managers
- A queue manager with a connected application that puts messages to the queue
- A queue manager that owns a local queue of the same name in the same cluster

Processor usage on the full repository queue managers must be assessed. Additional processor usage is visible as message traffic on the full repository queue `SYSTEM.CLUSTER.COMMAND.QUEUE`. If messages build up on that queue, it indicates that the full repository queue managers are unable to keep up with the rate of cluster configuration change in the system.

When many queues are being monitored by the cluster queue monitoring sample program, there is an amount of work performed by the sample program and the queue manager. This work is performed, even when there are no changes to the attached consumers. The `-i` argument can be modified to reduce processor usage of the sample program on the local system, by decreasing the frequency of the monitoring cycle.

To help detect excessive activity, the cluster queue monitoring sample program reports average processing time per polling interval, elapsed processing time, and number of configuration changes. The reports are delivered in an information message, `CLM0045I`, every 30 minutes, or every 600 poll intervals, whichever is sooner.

Cluster queue monitoring usage requirements

The cluster queue monitoring sample program has requirements and restrictions. You can modify the sample source code provided to change some of these restrictions in how it can be used. Examples listed in this section detail modifications that can be made.

- The cluster queue monitoring sample program is designed to be used to monitor queues where consuming applications are either attached, or not attached. If the system has consuming applications that are frequently attaching and detaching, the sample program might generate excessive cluster configuration activity across the entire cluster. This might have an impact on the performance of the queue managers in the cluster.
- The cluster queue monitoring sample program depends upon the underlying WebSphere MQ system and cluster technology. The number of queues being monitored, the frequency of monitoring and the frequency of the change of the state of each queue affects the load on the overall system. These factors must be considered when selecting the queues to be monitored and the poll interval of the monitoring.
- An instance of the cluster queue monitoring sample program must be connected to every queue manager in the cluster that owns an instance of a queue to be monitored. It is not necessary to connect the sample program to queue managers in the cluster that do not own the queues.
- The cluster queue monitoring sample program must be run with suitable authorization to access all of the WebSphere MQ resources required. For example,
 - The queue manager to be connected to
 - The `SYSTEM.ADMIN.COMMAND.QUEUE`
 - All the queues to be monitored when message transfer is performed
- The command server must be running for each queue manager with the cluster queue monitoring sample program connected.
- Each instance of the cluster queue monitoring sample program requires exclusive use of a local (non-clustered) queue on the queue manager that it is connected to. This local queue is used to control the sample program, and receive reply messages from inquiries made to the command server of the queue manager.
- All queues to be monitored by a single instance of the cluster queue monitoring sample program must be in the same cluster. If a queue manager has queues in multiple clusters that require monitoring,

multiple instances of the sample program are required. Each instance needs a local queue for control and reply messages.

- All queues to be monitored must be in a single cluster. Queues configured to use a cluster namelist are not monitored.
- Enabling the transfer of messages from inactive queues is optional. It applies to all queues being monitored by the instance of the cluster queue monitoring sample program. If only a subset of the queues being monitored require message transfer enabled, two instances of the cluster queue monitoring sample program are needed. One sample program has message transfer enabled, and the other has message transfer disabled. Each instance of the sample program needs a local queue for control and reply messages.
- WebSphere MQ cluster workload balancing will, by default, send messages to instances of clustered queues that reside on the same queue manager that a putting application is connected to. This must be disabled while the local queue is inactive in the following circumstances:
 - Putting applications connect to queue managers that own instances of an inactive queue that are being monitored
 - Queued messages are being transferred from inactive queues to active queues.

The local workload balancing preference on the queue can be disabled statically, through setting the **CLWLUSEQ** value to **ANY**. In this configuration messages put on local queues are distributed to local and remote queue instances to balance workload, even when there are local consuming applications. Alternatively, the cluster queue monitoring sample program can be configured to temporarily set the **CLWLUSEQ** value to **ANY** while the queue has no attached consumers which results in only local messages going to local instances of a queue while that queue is active.

- The WebSphere MQ system and applications must not use **CLWLPRTY** for the queues to be monitored, or channels being used. Otherwise, the actions of the cluster queue monitoring sample program on **CLWLPRTY** queue attributes might have undesired effects.
- The cluster queue monitoring sample program logs runtime information to a set of report files. A directory to store these reports is required, and the cluster queue monitoring sample program must have authorization to write to it.

AMQSCLM: Preparing and running the sample

In order to run the cluster queue monitoring sample, you must configure the queue manager to securely accept incoming connection requests from applications running in client mode.

Before you begin

The following steps must be completed before running the cluster queue monitoring sample.

1. Create a working queue on each queue manager for the internal use of the sample.

Each instance of the sample needs a local non-cluster queue for exclusive internal use. You can choose the name of the queue. The example uses the name **AMQSCLM.CONTROL.QUEUE**. For example, on Windows, you can create this queue by using the **MQSC** command

```
DEFINE QLOCAL(AMQSCLM.CONTROL.QUEUE)
```

You can leave the values of **MAXDEPTH** and **MAXMSGL** as default.

2. Create a directory for error and information message logs.

The sample writes diagnostic messages to report files. You must choose a directory in which to store the files. For example, on Windows, you can create a directory using the following command:

```
mkdir C:\AMQSCLM\lpts
```

The report files created by the sample have the following naming convention:

```
QmgrName.ClusterName.RPT0n.LOG
```

3. (Optional) Define the cluster queue monitoring sample as an IBM WebSphere MQ service.

To monitor queues, the sample must always be running. To ensure that the cluster queue monitoring sample is always running, you can define the sample as a queue manager service. Defining the sample as a service means that AMQSCLM is started when the queue manager starts. You can use the following **RUNMQSC** example to define the cluster queue monitoring sample as an IBM WebSphere MQ service.

```
define service(AMQSCLM) +
  descr('Active Cluster Queue Message Distribution Monitor - AMQSCLM') +
  control(qmgr) +
  servtype(server) +
  startcmd('<Install Root>\tools\c\samples\Bin\AMQSCLM.exe') +
  startarg('-m +QMNAME+ -c CLUSTER1 -q ABC* -r AMQSCLM.CONTROL.QUEUE -l c:\AMQSCLM\reports') +
  stdout('C:\AMQSCLM\reports\+QMNAME+.TSTCLUS.stdout.log') +
  stderr('C:\AMQSCLM\reports\+QMNAME+.TSTCLUS.stderr.log')
```

where <Install Root> is the location of your installation.

Definition	Description
service	Specifies the service name. You can choose the service name.
descr	Specifies a textual description of the service.
control	Indicates that the service starts and stops at the same time as the queue manager.
servtype	Indicates a server service object, meaning only one instance, can be executed at a time for this queue manager.
startcmd	Specifies the location and name of the program.
startarg	Specifies the arguments of the sample. Note the use of +QMNAME+. The name of the queue manager is automatically substituted.
stdout	The fully qualified file name to which standard output is redirected. The sample writes to this file only messages confirming that the sample has terminated. The sample does this because the standard error file has already closed in an earlier stage of the sample termination process.
stderr	The fully qualified file name to which standard error output is redirected. The sample writes to the standard error file any error messages prior to termination of the sample.

About this task

This task enables you to start and stop the cluster queue monitoring sample in different ways. It also enables you to run the sample in a mode that generates report files containing statistical information about the queues being monitored.

The sample program can be run by using the following command.

```
AMQSCLM -m QMgrName -c ClusterName (-q QNameMask | -f QListFile) -r MonitorQName
[-l ReportDir] [-t] [-u ActiveVal] [-i Interval] [-d] [-s] [-v]
```

The table lists the arguments that can be used with the cluster queue monitoring sample, along with additional information about each.

Argument	Variable	Further Information
-m	QMgrName	The queue manager to monitor.
-c	ClusterName	The cluster containing the queues to monitor.
-q	QNameMask	The queue, or queues, to monitor. A trailing * monitors all queues with names that match zero or more trailing characters.
-f	QListFile	The full path and file name of a file containing a list of queue names or queue name masks to monitor. The file must contain one queue name/mask per line. You can specify -q or -f , but not both.
-r	MonitorQName	The local queue being used exclusively by the sample.
-l	ReportDir	The directory path in which to store logged information messages in a set of wrapping<fn>For each queue manager and queue combination a report file is generated that is capped at a certain size. The logger always writes into the same file, but it also keeps the two previous versions of the file.</fn> report files.
-t		(Optional) Enables the transfer of queued messages from inactive local queues to active queues. If not enabled, only new messages entering the cluster are dynamically routed to active instances of a queue.
-u	ActiveVal	(Optional) Automatically switches the CLWLUSEQ property of a monitored queue instance to ANY when it is inactive, and to the ActiveVal property when active. ActiveVal can be LOCAL or QMGR. If this argument is not set in a system where putting applications connect to the same queue manager, or where message transfer is enabled, then monitored queues must have a CLWLUSEQ value of ANY, or QMGR with the queue manager having a value of ANY.
-i	Interval	(Optional) The time interval in seconds, at which the monitor checks the queues. Default is 300 seconds (5 minutes).
-d		(Optional) Enables additional diagnostic output. Debug output might be useful when initially configuring the system, or when working with the sample code.
-s		(Optional) Enables minimal statistical output per interval.
-v		(Optional) Log report information to standard out, in addition to the report files.

Argument list examples:

```
-m QMGR1 -c CLUS1 -f c:\QList.txt -r CLMQ -l c:\amqsc1m\1rpts -s
-m QMGR2 -c CLUS1 -q ABC* -r CLMQ -l c:\amqsc1m\1rpts -i 600
-m QMGR1 -c CLUSDEV -q QUEUE.* -r CLMQ -l c:\amqsc1m\1rpts -t -u QMGR -d
```

Example queue list file:

```
Q1
QUEUE.*
ABC
ABD
```

Procedure

1. Start the cluster queue monitoring sample. You can start the sample in one of the following ways:

- Use a command prompt with the appropriate user authorizations.
- Use the MQSC **START SERVICE** command, if the sample is configured as an IBM WebSphere MQ service.

The argument list is the same in both cases.

The sample does not start monitoring the queues for 10 seconds after the program is initialized. This delay allows consuming applications to connect to the monitored queues first, preventing unnecessary changes to the active state of the queue.

2. Stop the cluster queue monitoring sample. The sample automatically stops when the queue manager is stopped, stopping, quiescing, or if the connection to the queue manager is broken. There are ways to stop the sample without ending the queue manager:

- Configure the local queue used exclusively by the sample to disable the Get function.
- Send a message with a **CorrelId** of "STOP CLUSTER MONITOR\0\0\0\0", to the local queue used exclusively by the sample.
- Terminate the sample process. This might result in the loss of non-persistent messages being transferred to active queues. It might also result in the local queue used by the sample being held open for a number of seconds after the termination. This situation prevents a new instance of the cluster queue monitoring sample from starting immediately.

If the sample has been started as an IBM WebSphere MQ service, **STOP SERVICE** has no effect. It is possible to use one of the termination methods described as a configured **STOP SERVICE** mechanism in the queue manager.

What to do next

Check the status of the sample.

If reporting is enabled, you can review the report files for status. Use the following command to review the most current report file.

```
QMgrName.ClusterName.RPT01.LOG
```

To review older report files, use the following commands.

```
QMgrName.ClusterName.RPT02.LOG
QMgrName.ClusterName.RPT03.LOG
```

Report files grow to a maximum size of approximately 1 MB. When the RPT01 file fills up, a new RPT01 file is created. The old RPT01 file is renamed to RPT02. RPT02 is renamed to RPT03. The old RPT03 is discarded.

The sample creates information messages in the following situations:

- At startup
- At termination
- When it marks a queue **ACTIVE** or **INACTIVE**
- when it requeues messages from an inactive queue to an active instance, or instances

The sample creates an error message *CLMnnnnE* to report a problem that requires attention.

Every 30 minutes, the sample reports average processing time per polling interval, and elapsed processing time. This information is held in message CLM0045I.

When statistical messages are enabled **-s**, the sample reports the following statistical information about each queue check:

- Time taken to process the queues (in milliseconds)
- Number of queues checked
- Number of active/inactive changes made

- Number of messages transferred

This information is reported in message CLM0048I.

Report files might grow rapidly in debug mode, and quickly wrap. In this situation, the 1-MB size limit for individual files might be exceeded.

AMQSCLM: Troubleshooting

The following sections contain information about scenarios that might be encountered while using the sample. Information about potential explanations for a scenario, and options on how to resolve it, are provided.

Scenario: AMQSCLM is not starting

Potential explanation: Incorrect syntax.

Action: Check standard error output for correct syntax

Potential explanation: Queue manager is not available.

Action: Check the report file for message id CLM0010E.

Potential explanation: Cannot open or create report file or files.

Action: Check standard error output for error messages during initialization.

Scenario: AMQSCLM is not changing a queue to ACTIVE or INACTIVE

Potential explanation: The queue is not in the list of queues to be monitored

Action: Check the **-q** and **-f** parameter values.

Potential explanation: The queue is not a local queue in the correct cluster.

Action: Check that the queue is local and in the correct cluster.

Potential explanation: AMQSCLM is not running for this queue manager and cluster.

Action: Start AMQSCLM for the relevant queue manager and cluster.

Potential explanation: The queue is left INACTIVE, **CLWLPRTY**=0, because it has no consumers. Alternatively, it is left ACTIVE **CLWLPRTY**>=1, because it has at least 1 consumer.

Action: Check whether consuming applications are attached to the queue.

Potential explanation: The command server of the queue manager is not running.

Action: Check the report files for errors.

Scenario: Messages are not being routed around INACTIVE queues

Potential explanation: Messages are put directly to the queue manager that owns the inactive queue, and the **CLWLUSEQ** value of the queue is not ANY, and the **-u** argument is not being used for AMQSCLM.

Action: Check the **CLWLUSEQ** value of the relevant queue manager, or ensure that the **-u** argument is used for AMQSCLM.

Potential explanation: There are no active queues on any queue managers. Messages are evenly workload balanced across all inactive queues until a queue becomes active.

Action: Check the status of the queues on all queue managers.

Potential explanation: Messages are put to a different queue manager in the cluster to the one that owns the inactive queue, and the updated **CLWLPRTY** value of 0 is not propagated to the queue manager of the putting application.

Action: Check that the cluster channels between the monitored queue manager, and the full repository queue manager, are running. Check that the channels between the putting queue manager, and the full

repository queue manager, are running. Check the error logs of the monitored, putting, and full repository queue managers.

Potential explanation: The remote queue instances are active (`CLWLPRTY=1`), but messages cannot be routed to those queue instances because the cluster sender channel from the local queue manager is not running.

Action: Check the status of the cluster sender channels from the local queue manager to the remote queue manager, or managers, with an active instance of the queue.

Scenario: AMQSCLM is not transferring messages from an inactive queue

Potential explanation: Message transfer is not enabled (`-t`).

Action: Ensure that message transfer is enabled (`-t`).

Potential explanation: The queue is not in the list of queues to be monitored.

Action: Check the `-q` and `-f` parameter values.

Potential explanation: AMQSCLM is not running for this, or other queue managers in the cluster, that own instances of the same queue.

Action: Start AMQSCLM.

Potential explanation: The queue has `CLWLUSEQ=LOCAL` or `CLWLUSEQ=QMGR`, and the `-u` argument is not set.

Action: Set the `-u` parameter, or change the queue, or queue manager configuration, to ANY.

Potential explanation: There are no active instances of the queue in the cluster.

Action: Check for instances of the queue with a `CLWLPRTY` value of 1, or greater.

Potential explanation: Remote queue instances have consumers (`IPPROCS>=1`) but are inactive on those queue managers (`CLWLPRTY=0`) because AMQSCLM is not monitoring those remote instances.

Action: Ensure that AMQSCLM is running on those queue managers, and / or the queue is in the list of queues to be monitored by checking the `-q` and `-f` parameter values.

Potential explanation: The remote queue instances are active (`CLWLPRTY=1`), but are seen as inactive on the local queue manager (`CLWLPRTY=0`). This situation is due to the updated `CLWLPRTY` value not being propagated to this queue manager.

Action: Ensure that the remote queue managers are connected to at least one of the full repository queue managers in the cluster. Ensure that the full repository queue managers are functioning correctly. Check that the channels between the full repository queue managers, and the monitored queue managers, are running.

Potential explanation: The messages are not committed, therefore they are not retrievable.

Action: Check that the sending application is functioning correctly.

Potential explanation: AMQSCLM does not have access to the local queue where messages are queued.

Action: This scenario could be due to AMQSCLM not running as a user with sufficient authorization to access the queue.

Potential explanation: The command server of the queue manager is not running.

Action: Start the command server of the queue manager.

Potential explanation: AMQSCLM encountered an error.

Action: Check the report files for errors.

Potential explanation: The remote queue instances are active (`CLWLPRTY=1`), but messages cannot be transferred to those queue instances because the cluster sender channel from the local queue manager is not running. This is often accompanied by a CLM0030W warning in the amqsclm report log.

Action: Check the status of the cluster sender channels from the local queue manager to the remote queue manager, or managers, with an active instance of the queue.

Sample program for Connection Endpoint Lookup (CEPL)

IBM WebSphere MQ Connection Endpoint Lookup sample provides a simple yet powerful exit module that offers WebSphere MQ users a way to retrieve connection definitions from an LDAP repository such as Tivoli Directory Server.

Tivoli Directory Server v6.3 Client must be installed in order to use CEPL.

A working knowledge of WebSphere MQ administration on the supported platforms is required to use this sample.


Introduction

Configure a global repository, for example, an LDAP (Lightweight Directory Access Protocol) directory, to store client connection definitions to aid maintenance and administration.

Using an IBM WebSphere MQ Client application to establish a connection to a Queue Manager via a Client Connection Definition Table (CCDT).

The CCDT is created through the standard WebSphere MQ MQSC Administration interface. The user must be connected to a Queue Manager in order to create client connection definitions, even though the data contained within the definition is not restricted to the Queue Manager.

The CCDT file generated must be manually distributed among client machines and applications.




The CCDT file must be distributed to each WebSphere MQ client. Where thousands of clients can exist either locally or globally, it would soon become difficult to maintain and administer. A more flexible approach is needed to help ensure that each client has the correct client definitions available to them.

One such approach is to store the client connection definitions in a global repository such as an LDAP (Lightweight Directory Access Protocol) directory. An LDAP directory can also provide additional security, indexing, and search facilities, thereby allowing each client access to only those connection definitions pertaining to them.

The LDAP directory can be configured so that only specific definitions are available to certain user groups. For example, the Test Clients can access both Queue Manager

#1 and #2, whereas the Development Clients can access Queue Manager #2 only.



The exit module can look up an LDAP repository, for example, IBM Tivoli Directory Server, to retrieve channel definitions. Using those connection definitions, a WebSphere MQ client application can establish connection to a queue manager.

The exit module is a pre-connect exit module which enables channel definition to be obtained during the MQCONN/MQCONN call from an LDAP repository.

The exit module and schema might be implemented by:

- Customers who have already built a skill base using the existing CCDT file based technology and want to ease the administration and distribution costs.
- Existing customers who already employ their own propriety technology for distributing client connection definitions.
- New or existing customers who currently do not employ any type of client connection solution and want to use the features offered by IBM WebSphere MQ.
- New or existing customers who want to directly use or tune their messaging model inline with any current LDAP business architecture.

Supported Environments

Verify that you have a supported operating system and the relevant software before running the Connection Endpoint Lookup sample.

The sample program for IBM WebSphere MQ Connection Endpoint Lookup requires the following software:

- IBM WebSphere MQ V7.0, or later
- Tivoli Directory Server V6.3 Client, or later

Operating Systems supported:

1. Windows (XP/2003/2008)
2. Solaris (SPARC and x86-64)
3. AIX
4. Linux
 - RHEL v4 and v5 on System p
 - SUSE v9 and v10 on System p
 - RHEL v4 and v5 System x32 bit and x64 bit
 - SUSE v9 and v10 System x32 bit and x64 bit
5. HP IA64.

Note: The sample program is not available for z/OS, i/5, and HP PARISC platforms.

Installing and Configuring

Installing and configuring the exit module and connection endpoint schema.

Installing exit module

During installation of WebSphere MQ, the exit module is installed under `tools/samples/c/preconnext/bin`. For 32 bit platforms, the exit module must be copied to `exit/<install name>/` before it can be used. For 64 bit platforms, the exit module must be copied to `exit64/<installation name>/` before it can be used.

Installing Connection Endpoint schema

The exit uses the Connection Endpoint schema, *ibm-amq.schema*. The schema file must be imported into any LDAP server before the exit can be used. After importing the schema, values for the attributes must be added.

Here is an example for importing Connection Endpoint schema. The example assumes IBM Tivoli Directory Server (ITDS) is being used.

- Ensure that IBM Tivoli Directory Server is running, then copy or FTP the *ibm-amq.schema* file to the ITDS server.
- On the ITDS server, enter the following command to install the schema into ITDS store, where LDAP ID and LDAP password are the root DN and password for the LDAP server:

```
ldapadd -D "LDAP ID" -w "LDAP password" -f ibm-amq.schema
```

- In a command window, enter the following command or use a third party tool to browse the schema for verification:

```
ldapsearch objectclass=ibm-amqClientConnection
```

Refer to your LDAP Server documentation for further details on importing the schema file.

Configuration

A new section called **PreConnect** must be added to the client configuration file say *mqclient.ini* file. The PreConnect section contains the following keywords:

Module : The name of the module containing the API exit code. If this field contains the full path of the module, it is used as is else *exit* or *exit64* folder in WebSphere MQ installation is searched.

Function : Name of the functional entry point into the library that contains the PreConnect exit code. The function definition adheres to the MQ_PRECONNECT_EXIT prototype.

Data : URI of the LDAP repository containing channel definitions.

The following snippet is an example of the changes required to *mqclient.ini* file.

```
PreConnect:
Module=amqlcelp
Function=PreConnectExit
Data=ldap://myLDAPServer.com:389/cn=wmq,ou=ibm,ou=com
Sequence=1
```

Overview of Exit and Schema

Syntax, and parameters used to establish a connection to a queue manager.

WebSphere MQ v7.5 defines the following syntax for an entry point in an exit module.

```
void MQENTRY MQ_PRECONNECT_EXIT ( PMQNX  pExitParms
                                   , PMQCHAR pQMgrName
                                   , PPMQCN  ppConnectOpts
                                   , PMQLONG pCompCode
                                   , PMQLONG pReason)
```

During the MQCONN/X call execution, WebSphere MQ C Client loads the exit module containing an implementation of the function syntax. It then invokes an exit function to retrieve channel definitions. The retrieved channel definitions are then used to establish connection to a queue manager.

Parameters

pExitParms

Type: PMQNX input/output

The PreConnection exit parameter structure. The structure is allocated and maintained by the caller of the exit.

```
struct tagMQNX
{
    MQCHAR4    StrucId;          /* Structure identifier */
    MQLONG     Version;         /* Structure version number */
    MQLONG     ExitId;          /* Type of exit */
    MQLONG     ExitReason;      /* Reason for invoking exit */
    MQLONG     ExitResponse;    /* Response from exit */
    MQLONG     ExitResponse2;   /* Secondary response from exit */
    MQLONG     Feedback;        /* Feedback code (reserved) */
    MQLONG     ExitDataLength;  /* Exit data length */
    PMQCHAR    pExitDataPtr;    /* Exit data */
    MQPTR      pExitUserAreaPtr; /* Exit user area */
    PMQCD *    ppMQCDArrayPtr;  /* Array of pointers to MQCDs */
    MQLONG     MQCDArrayCount;  /* Number of entries found */
    MQLONG     MaxMQCDVersion;  /* Maximum MQCD version */
};
```

pQMgrName

Type: PMQCHAR input/output

Name of the queue manager. On input, this parameter is the filter string supplied to the MQCONN API call through the **QMgrName** parameter. This field might be blank, explicit, or contain certain wildcard characters. The field is changed by the exit. The parameter is NULL when the exit is called with MQXR_TERM.

ppConnectOpts

Type: ppConnectOpts input/output

Options that control the action of MQCONN/X. This is a pointer to an MQCNO connection options structure that controls the action of the MQCONN API call. The parameter is NULL when the exit is called with MQXR_TERM. The MQI client always provides an MQCNO structure to the exit, even if it was not originally provided by the application. If an application provides an MQCNO structure, the client makes a duplicate to pass it to the exit where it is modified. The client retains the ownership of the MQCNO. An MQCD referenced through the MQCNO takes precedence over any connection definition provided through the array. The client uses the MQCNO structure to connect to the queue manager and the others are ignored.

pCompCode

Type: PMQLONG input/output

Completion code. Pointer to an MQLONG that receives the exits completion code. It must be one of the following values:

MQCC_OK - Successful completion

MQCC_WARNING - Warning (partial completion)

MQCC_FAILED - Call failed

pReason

Type: PMQLONG input/output

Reason qualifying pCompCode. Pointer to an MQLONG that receives the exit reason code. If the completion code is MQCC_OK, the only valid value is:

MQRC_NONE - (0, x'000') No reason to report.

If the completion code is MQCC_FAILED or MQCC_WARNING, the exit function can set the reason code field to any valid MQRC_* value.

MQ LDAP Context Information

The exit uses the following data structure for context information.

MQNLDACTX

The MQNLDACTX structure has the following C prototype.

```
typedef struct tagMQNLDACTX MQNLDACTX;
typedef MQNLDACTX MQPOINTER PMQNLDACTX;

struct tagMQNLDACTX
{
    MQCHAR4    StrucId;          /* Structure identifier */
    MQLONG     Version;         /* Structure version number */
    LDAP *     objectDirectory; /* LDAP Instance */
    MQLONG     ldapVersion;     /* Which LDAP version to use? */
    MQLONG     port;            /* Port number for LDAP server */
    MQLONG     sizeLimit;       /* Size limit */
    MQBOOL     ssl;             /* SSL enabled? */
    MQCHAR *   host;            /* Hostname of LDAP server */
    MQCHAR *   password;        /* Password of LDAP server */
    MQCHAR *   searchFilter;    /* LDAP search filter */
    MQCHAR *   baseDN;          /* Base Distinguished Name */
    MQCHAR *   charSet;         /* Character set */
};
```

Sample code for building the connection endpoint lookup exit

Code snippets for compiling the source on Windows and distributed platforms.

Compiling source

You can compile the source with any LDAP client libraries, for example, IBM Tivoli Directory Server 6.3 Client libraries. This documentation assumes that you are using Tivoli Directory Server 6.3 client libraries.

Note: The pre-connect exit library has been tested with the following LDAP servers:

- IBM Tivoli Directory Server V6.3
- Novell eDirectory V8.2

The following code snippets describe how to compile the exits on Windows, and other distributed platforms:

Compiling the exit on the Windows platform

You can use the following snippet for compiling the exit source on Windows:

```
CC=c1.exe
LL=link.exe
CCARGS=/c /I. /DWIN32 /W3 /DNDEBUG /EHsc /D_CRT_SECURE_NO_DEPRECATED /Z1
```

```
# The libraries to include
LDLIBS=ws2_32.lib Advapi32.lib libibmldapstatic.lib libibmldapbgstatic.lib \
kernel32.lib user32.lib gdi32.lib winspool.lib comdlg32.lib advapi32.lib \
shell32.lib ole32.lib oleaut32.lib uuid.lib odbc32.lib odbccp32.lib msvcrt.lib

OBJS=amqlcel0.obj

all: amqlcelp.dll

amqlcelp.dll: $(OBJS)
    $(LL) /OUT:amqlcelp.dll /INCREMENTAL /NOLOGO /DLL /SUBSYSTEM:WINDOWS /MACHINE: X86 /
DEF:amqlcelp.def $(OBJS) $(LDLIBS) /NODEFAULTLIB:msvcrt.lib

# The exit source
amqlcel0.obj: amqlcel0.c
    $(CC) $(CCARGS) $*.c
```

Note: You might get warnings while compiling the IBM Tivoli Directory Server 6.3 Client libraries with Microsoft Visual Studio 2005 or above compiler, if you are using the IBM Tivoli Directory Server 6.3 Client libraries compiled with Microsoft Visual Studio 2003 compiler.

Compiling the exit on other distributed platforms

You can use the following snippet for compiling the exit source on other distributed platforms, for example, Linux. Some compiler options might differ on other distributed platforms.

```
#Make file to build exit
CC=gcc

MQML=/opt/mqm/lib
MQMI=/opt/mqm/inc
TDSI=/opt/ibm/ldap/V6.3/include
XFLAG=-m32

TDSL=/opt/ibm/ldap/V6.3/lib
```

IBM Tivoli Directory Server ships both static and dynamic link libraries, but only one form of the libraries can be used. This script assumes that you are using the static libraries.

```
#Use static libraries.
LDLIBS=-L$(TDSL) -libibmldapstatic

CFLAGS=-I. -I$(MQMI) -I$(TDSI)

all:amqlcepl

amqlcepl: amqlcel0.c
    $(CC) -o cepl amqlcel0.c -shared -fPIC $(XFLAG) $(CFLAGS) $(LDLIBS)
```

Exit Module Invocation

The PreConnect exit module can be invoked with three different reason codes. This section describes each exit reason in greater depth.

MQXR_INIT

The exit is invoked with MQXR_INIT reason code for initializing and establishing connection to an LDAP server.

Before the *MQXR_INIT* call, the *pExitDataPtr* field of MQNXP structure would have been populated with the Data attribute from the PreConnect stanza within the *mqclient.ini* file (i.e. the LDAP).

An LDAP URL consists of at least the protocol, hostname, port number, and base DN for the search. The exit parses the LDAP URL contained within the *pExitDataPtr* field, allocates an MQNLDAPCTX LDAP Lookup Context structure and populates it accordingly. The address of this structure is stored in the *pExitUserAreaPtr* field. Failure to correctly parse the LDAP URL results in error MQCC_FAILED.

At this point the exit connects and binds to the LDAP server using the MQNLDAPCTX parameters. The resulting LDAP API handles are also stored within this structure.

MQXR_PRECONNECT

The exit module is invoked with MQXR_PRECONNECT reason code for retrieving channel definitions from an LDAP server.

The exit searches the LDAP server for channel definitions matching the given filter. If *QMgrName* parameter contains a specific queue manager name, then search returns all channel definitions whose *ibm-amqQueueManagerName* LDAP attribute value matches with the given queue manager name.

If the *QMgrName* parameter is '*' or ' '(blank), then the search returns all channel definitions whose *ibm-amqIsClientDefault* Connection endpoint attribute is set to true.

After a successful search, the exit prepares one or an array of MQCD definitions and returns back to the caller.

MQXR_TERM

The exit is invoked with this reason code when exit is to be cleaned. During this the exit disconnects from the LDAP server, releases all the memory allocated and maintained by the exit. This will include the *MQNLDPCTX* structure, the pointer array and every MQCD it references. Any other fields are set to the default values. The *pQMgrName* and *ppConnectOpts* exit parameters are unused during *MQXR_TERM* and may be NULL.

LDAP schemas

Client connection data is stored in a global repository called the LDAP (Lightweight Directory Access Protocol) directory. A WebSphere MQ client uses an LDAP directory to obtain the connection definitions. The structure of the WebSphere MQ client connection definitions within the LDAP directory is known as the LDAP schema. An LDAP schema is the collection of attribute type definitions, object class definitions, and other information which a server uses to determine whether a filter or attribute value assertion matches against the attributes of an entry, and whether to permit, add, and modify operations.

Storing data in the LDAP directory

The client connection definitions are located under a specific branch within the directory tree known as the connection point. Like all other nodes within an LDAP directory, the connection point has a Distinguished Name (DN) associated with it. You can use this node as the starting point for any queries you make on the directory. Use filtering when querying the LDAP directory to return a subset of client connection definitions. You can restrict access to sub-trees based on permissions granted in other parts of the directory tree - for example, to users, departments, or groups.

Defining your own attributes and classes

Store the client channel definition by modifying the LDAP schema. All LDAP data definitions require objects and attributes. The objects and attributes are identified by an object identifier (OID) number which uniquely identifies the object or attribute. All classes within an LDAP schema inherit either directly or indirectly from the top object. The client channel definition object contains the attributes of the top object. All LDAP data definitions require objects and attributes:

- Object definitions are collections of LDAP attributes.
- Attributes are LDAP data types.

The description of each attribute and how they map to the normal WebSphere MQ properties are described in [LDAP attributes](#).

LDAP attributes

LDAP attributes defined are specific to WebSphere MQ and maps directly to the client connection properties.

WebSphere MQ Client Channel Directory String Attributes

The character string attributes with their mapping to WebSphere MQ properties are listed in the following table. The attributes can hold values of directoryString (UTF-8 encoded Unicode, that is, a variable byte encoding system that includes IA5/ASCII as a subset) syntax. The syntax is specified by its object identification number (OID).

Table 21. WebSphere MQ client channel directory string attributes

LDAP Attribute	Description	WebSphere MQ Property
CN	The common name consisting of the channel name and the defining queue manager name.	
ibm-amqChannelName	The name of the channel definition.	CHANNEL
ibm-amqConnectionName	The communication connection identifier.	CONNAME
ibm-amqDescription	The channel description.	DESCR
ibm-amqLocalAddress	The local communication address of the channel.	LOCLADDR
ibm-amqModeName	s a bThe LU 6.2 mode name.	MODENAME
ibm-amqPassword	The password that can be used.	PASSWORD
ibm-amqQueueManagerName	The name of the queue manager or queue manager group to which a WebSphere MQ client application can request connection.	QMNAME
ibm-amqSecurityExitUserData	The user data that is passed to the security exit.	SCYDATA
ibm-amqSecurityExitName	The name of the exit program to be run by the channel security exit.	SCYEXIT
ibm-amqSslCipherSpec	A single CipherSpec for an SSL connection.	SSLCIPH
ibm-amqSslPeerName	Checks the Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of a WebSphere MQ channel.	SSLPEER
ibm-amqTransactionProgramName	The transaction program name.	TPNAME
ibm-amqUserID	The user ID to be used by the MCA when attempting to initiate a secure SNA session with a remote MCA.	USERID

WebSphere MQ client connection integer attributes

The attributes with predefined values (for example, an enumerated type) are stored as standard integers. These values are stored in the LDAP directory as integer values, and not by using the associated constant name.

Table 22. WebSphere MQ client channel directory integer attributes

LDAP attribute	Description	WebSphere MQ Property
ibm-amqConnectionAffinity	Determines whether client applications, which connect multiple times through the same queue manager name, use the same client channel.	AFFINITY
ibm-amqClientChannelWeight	A weighting to influence which client connection channel definition is used.	CLNTWGHT
ibm-amqHeartBeatInterval	The approximate time between heartbeat flows that are to be passed from a sending MCA when there are no messages on the transmission queue.	HBINT
ibm-amqKeepAliveInterval	A time-out value for a channel.	KAINT

Table 22. WebSphere MQ client channel directory integer attributes (continued)

LDAP attribute	Description	WebSphere MQ Property
ibm-amqMaximumMessageLength	The maximum length of a message that can be transmitted on the channel.	MAXMSGL
ibm-amqSharingConversations	The maximum number of conversations that share each TCP/IP channel instance.	SHARECNV
ibm-amqTransportType	The transport type to be used.	TRPTYPE

WebSphere MQ client channel boolean attribute

This Boolean attribute is not mapped to any WebSphere MQ property. The syntax of this attribute indicates a boolean value.

Table 23. WebSphere MQ client channel boolean attribute

LDAP attribute	Description
ibm-amqIsClientDefault	This boolean attribute is defined to resolve the problem of searching entries whose ibm-amqQueueManagerName attribute has not been defined.

WebSphere MQ client channel list attributes

WebSphere MQ properties are stored as single-value, comma-separated list attribute within the LDAP directory. The attributes are defined in the same manner as the other directory string attributes. The list attributes along with their mapping to the WebSphere MQ properties are described in the following table.

Table 24. WebSphere MQ client channel list attributes

LDAP attribute	Description	WebSphere MQ Property
ibm-amqHeaderCompression	A list of header data compression techniques supported by the channel.	COMPHDR
ibm-amqMessageCompression	A list of message data compression techniques supported by the channel.	COMPMSG
ibm-amqSendExitUserData	The user data that is passed to the send exit.	SENDDATA
ibm-amqSendExitUserName	The name of the exit program to be run by the channel send exit.	SENDEXIT
ibm-amqReceiveExitUserData	The user data that is passed to the receive exit.	RCVDATA
ibm-amqReceiveExitName	The name of the user exit program to be run by the channel receive user exit.	RCVEXIT

Common Name

The common name (CN) consists of the channel name and the defining queue manager name.

It is a preexisting attribute.

The format of the CN is:

```
CN=CHANNEL_NAME(DEFINING_Q_MGR_NAME)
```

For example:

```
CN=TC1(QM_T1)
```

You can specify only one value for this attribute.

This attribute is a string attribute and the values are not case-sensitive. The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter, using a substring (for example, CN=jim* where CN is an attribute) and contains one or more wild cards.

ibm-amqChannelName

This attribute specifies the name of the channel definition.

This attribute has a single string value with a maximum of 20 characters that are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter, using a substring and contains one or more wild cards.

ibm-amqDescription

This LDAP attribute provides the channel description.

This attribute has a single string value with a maximum of 64 bytes, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqConnectionName

This LDAP attribute is the communications connection identifier. It specifies the particular communications links to be used by this channel.

This attribute has a single string value with a maximum of 264 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqLocalAddress

This attribute specifies the local communications address for the channel.

This attribute has a single string value with a maximum of 48 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqModeName

This attribute is for use with LU 6.2 connections. It gives extra definition for the session characteristics of the connection when a communication session allocation is performed.

This attribute has a single string value of exactly 8 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqPassword

This LDAP attribute specifies a password that can be used by the MCA when attempting to initiate a secure LU 6.2 session with a remote MCA.

This attribute has a single integer value with a maximum of 12 digits. It is not a pre-existing attribute.

ibm-amqQueueManagerName

This attribute specifies the name of the queue manager or queue manager group to which a WebSphere MQ client application can request connection.

This attribute has a single string value with a maximum of 48 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqSecurityExitUserData

This LDAP attribute specifies user data that is passed to the security exit.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqSecurityExitName

This LDAP attribute specifies the name of the exit program to be run by the channel security exit.

Leave blank if no channel security exit is in effect.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. This attribute is not a pre-existing one.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqSslCipherSpec

This LDAP attribute specifies a single CipherSpec for an SSL connection.

This attribute has a single string value with a maximum of 32 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqSslPeerName

This LDAP attribute is used to check the Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of a WebSphere MQ channel.

This LDAP attribute has a single string value with a maximum of 1024 bytes, which are not case-sensitive. It is not a pre-existing one.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqTransactionProgramName

This LDAP attribute specifies the transaction program name. It is for use with LU 6.2 connections.

This attribute has a single string value with a maximum of 64 characters, which are not case-sensitive. It is not a pre-existing one.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqUserID

This LDAP attribute specifies the user ID to be used by the MCA when attempting to initiate a secure SNA session with a remote MCA.

This attribute has a single string value of exactly 12 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

ibm-amqConnectionAffinity

This LDAP attribute specifies whether client applications, which connect multiple times using the same queue manager name, use the same client channel.

This attribute has a single integer value. It is not a pre-existing attribute.

ibm-amqClientChannelWeight

This LDAP attribute specifies a weighting that influences which client connection channel definition is used.

The client channel weighting attribute is used to bias the selection of client channel definitions when more than one suitable definition is available.

This attribute has a single integer value. It is not a pre-existing attribute.

ibm-amqHeartBeatInterval

This LDAP attribute specifies the approximate time between heartbeat flows that are to be passed from a sending MCA when there are no messages on the transmission queue.

This attribute has a single integer value. It is not a pre-existing attribute. The default value is 1. The default is set in the current MQSERVER environment variable operation.

ibm-amqKeepAliveInterval

This LDAP attribute is used to specify a time-out value for a channel.

The value of this attribute is passed to the communications stack specifying the keepalive timing for the channel. You can use this to specify a different keepalive value for each channel.

This attribute has a single integer value. It is not a pre-existing attribute.

ibm-amqMaximumMessageLength

This LDAP attribute specifies the maximum length of a message that can be transmitted on the channel.

The default value of this attribute is 104857600 as per the current MQSERVER environment variable operation. This attribute has a single integer value and it is not a pre-existing attribute.

ibm-amqSharingConversations

This LDAP attribute specifies the maximum number of conversations that share each TCP/IP channel instance.

This attribute has a single integer value. This attribute is not a pre-existing attribute.

ibm-amqTransportType

This LDAP attribute specifies the transport type to be used.

This attribute has a single integer value. It is not a pre-existing attribute.

ibm-amqIsClientDefault

This Boolean attribute solves the problem of searching entries where the *ibm-amqQueueManagerName* attribute has not been defined.

Preconnect exit modules generally search the LDAP servers with the value of the *ibm-amqQueueManagerName* attribute as the search criteria. Such a query would return all entries where the *ibm-amqQueueManagerName* attribute value matches the name of the queue manager specified on the MQCONN/X call. However when using the client channel definition tables (CCDT), you can either set the queue manager name on a MQCONN/X call as blank or prefix the name with an asterisk (*). If the name of the queue manager is blank, the client connects to the default queue manager. If the name is prefixed with an asterisk (*) to the queue manager, then the client connects any queue manager.

Similarly, the `ibm-amqQueueManagerName` attribute in an entry can be left undefined. In this case, it is expected that the client using this endpoint information can connect to any queue manager. For example, an entry contains the following lines:

```
ibm-amqChannelName = "CHANNEL1"
ibm-amqConnectionName = myhost(1414)
```

In this example, the client attempts to connect to the specified queue manager running on `myhost`.

However in LDAP Servers, a search is not made on an attribute value that has not been defined. For example, if an entry contains the connection information except `ibm-amqQueueManagerName`, then the search results would not include this entry. To overcome this problem, you can set `ibm-amqIsClientDefault`. This is a Boolean attribute and is assumed to have a value of `FALSE` if not defined.

For entries where the `ibm-amqQueueManagerName` has not been defined and are expected to be part of the search, set `ibm-amqIsClientDefault` to `TRUE`. When a blank or asterisk (*) is specified as the queue manager name in a call to `MQCONN/X`, the preconnect exit searches the LDAP sever for all entries where `ibm-amqIsClientDefault` attribute value is set to `TRUE`.

Note: Do not set or define `ibm-amqQueueManagerName` attribute if `ibm-amqIsClientDefault` is set to `TRUE`.

ibm-amqHeaderCompression

This LDAP attribute is a list of header data compression techniques supported by the channel.

The maximum size of this attribute is of 48 characters. It is not a pre-existing attribute.

You can specify only one value for this attribute.

This list attribute is specified as directory strings using a comma-separated format. For example, The values specified for **`ibm-amqHeaderCompression`** is `0` which is mapped to `NONE`. Any values that exceed the maximum allowed limit will be ignored by the client. For example, `ibm-amqHeaderCompression` contains a maximum of 2 integers in the list.

ibm-amqMessageCompression

This LDAP attribute is a list of message data compression techniques supported by the channel.

The maximum size of this attribute is of 48 characters. It is not a pre-existing attribute.

This attribute does not support multiple values.

This list attribute is specified as directory strings using a comma-separated format. For example, the value specified for this attribute is `1,2,4`, which maps to the underlying compression sequence `RLE`, `ZLIBFAST`, and `ZLIBHIGH`.

Any values exceeding the maximum allowed limit is ignored by the client. For example, `ibm-amqMessageCompression` contains a maximum of 16 integers in the list.

ibm-amqSendExitUserData

This LDAP attribute specifies user data that is passed to the send exit.

This LDAP attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

Note: **`ibm-amqSendExitName`** and **`ibm-amqSendExitUserData`** need to be synchronized in pairs . The user data should be synchronized with the exit name. So if one is specified, the other also has to be symmetrically specified, even if it contains no data.

ibm-amqSendExitName

This LDAP attribute specifies the name of the exit program to be run by the channel send exit.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

Note: **ibm-amqSendExitName** and **ibm-amqSendExitUserData** must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other also must be symmetrically specified even if it contains no data.

ibm-amqReceiveExitUserData

This LDAP attribute specifies user data that is passed to the receive exit.

You can run a sequence of receive exits. The string of user data for a series of exits is separated by a comma, spaces, or both.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

Note: **ibm-amqReceiveExitName** and **ibm-amqReceiveExitUserData** must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other also must be symmetrically specified even if it contains no data.

ibm-amqReceiveExitName

This LDAP attribute specifies the name of the user exit program to be run by the channel receive user exit.

This attribute is a list of names of programs that are to be run in succession. Leave blank, if no channel receive user exit is in effect.

This attribute has a single string value with a maximum of 999 characters, which are not case-sensitive. It is not a pre-existing attribute.

The substring matching is ignored. The substring matching is a matching rule used in sub-schema that specifies the behavior of the attribute in a search filter.

Note: **ibm-amqReceiveExitName** and **ibm-amqReceiveExitUserData** must be synchronized in pairs. The user data must be synchronized with the exit name. So if one is specified, the other must also be symmetrically specified, even if it contains no data.

Writing a queuing application

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

Use the following links to find out more about writing applications:

Related concepts

[“Application development concepts” on page 8](#)

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Sample WebSphere MQ programs” on page 91](#)

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

[“Writing client applications” on page 335](#)

What you need to know to write client applications on WebSphere MQ.

[“Using Web services in WebSphere MQ” on page 904](#)

You can develop IBM WebSphere MQ applications for Web services using the IBM WebSphere MQ transport for SOAP or the IBM WebSphere MQ bridge for HTTP.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

[“Handling program errors” on page 523](#)

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

The Message Queue Interface overview

Learn about the Message Queue Interface (MQI) components.

The Message Queue Interface consists of the following:

- *Calls* through which programs can access the queue manager and its facilities
- *Structures* that programs use to pass data to, and get data from, the queue manager
- *Elementary data types* for passing data to, and getting data from, the queue manager

WebSphere MQ for Windows and WebSphere MQ on UNIX and Linux systems also supply:

- Calls through which WebSphere MQ for Windows and WebSphere MQ on UNIX and Linux systems programs can commit and back out changes.
- *Include files* that define the values of constants supplied on these platforms.
- *Library files* to link your applications.
- A suite of sample programs that demonstrate how to use the MQI on these platforms. For further information about these samples, see [“Sample programs for distributed platforms” on page 92](#).
- Sample source and executable code for bindings to external transaction managers.

Use the following links to find out more about the MQI:

- [“MQI calls” on page 188](#)
- [“Sync point calls” on page 189](#)
- [“Data conversion, data types, data definitions, and structures” on page 189](#)
- [“IBM WebSphere MQ stub programs and library files” on page 189](#)
- [“Parameters common to all the calls” on page 194](#)
- [“Specifying buffers” on page 195](#)
- [“UNIX and Linux signal handling” on page 195](#)

Related concepts

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

MQI calls

Use this information to learn about calls in MQI.

The calls in the MQI can be grouped as follows:

MQCONN, MQCONNX, and MQDISC

Use these calls to connect a program to (with or without options), and disconnect a program from, a queue manager. If you write CICS programs for z/OS, you do not need to use these calls. However, you are recommended to use them if you want to port your application to other platforms.

MQOPEN and MQCLOSE

Use these calls to open and close an object, such as a queue.

MQPUT and MQPUT1

Use these calls to put a message on a queue.

MQGET

Use this call to browse messages on a queue, or to remove messages from a queue.

MQSUB, MQSUBRQ

Use these calls to register a subscription to a topic, and to request publications matching the subscription.

MQINQ

Use this call to inquire about the attributes of an object.

MQSET

Use this call to set some of the attributes of a queue. You cannot set the attributes of other types of object.

MQBEGIN, MQCMIT, and MQBACK

Use these calls when WebSphere MQ is the coordinator of a unit of work. MQBEGIN starts the unit of work. MQCMIT and MQBACK end the unit of work, either committing or rolling back the updates made during the unit of work. Native start commitment control, commit, and rollback commands are used.

MQCRTMH, MQBUFMH, MQMHBUF, MQDLTMH

Use these calls to create a message handle, to convert a message handle to a buffer or a buffer to a message handle, and to delete a message handle.

MQSETMP, MQINQMP, MQDLTMP

Use these calls to set a message property on a message handle, inquire on a message property, and delete a property from a message handle.

MQCB, MQCB_FUNCTION, MQCTL

Use these calls to register and control a callback function.

MQSTAT

Use this call to retrieve status information about previous asynchronous put operations.

See [Call descriptions](#) for a description of the MQI calls.

Sync point calls

Use this information to find out about sync point calls on different platforms.

Sync point calls are available as follows:

IBM WebSphere MQ calls on Windows, UNIX and Linux platforms



The following products provide the MQCMIT and MQBACK calls:

- IBM WebSphere MQ for Windows
- IBM WebSphere MQ on UNIX and Linux systems

Use sync point calls in programs to tell the queue manager that all the MQGET and MQPUT operations since the last sync point are to be made permanent (committed) or are to be backed out. To commit and back out changes in the CICS environment, use commands such as EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK.

Data conversion, data types, data definitions, and structures

Use this information to learn about data conversions, elementary data types, WebSphere MQ data definitions, and structures when using the Message Queue Interface.

Data conversion

The MQXCNVC (convert characters) call converts message character data from one character set to another. Except on WebSphere MQ for z/OS, this call is used only from a data-conversion exit.

See [MQXCNVC - Convert characters](#) for the syntax used with the MQXCNVC call, and [“Writing data-conversion exits”](#) on page 397 for guidance on writing and invoking data conversion exits.

Elementary data types

For the supported programming languages, the MQI provides elementary data types or unstructured fields.

These data types are described fully in [Elementary data types](#).

WebSphere MQ data definitions

The data definition files supplied with WebSphere MQ contain:

- Definitions of all the WebSphere MQ constants and return codes
- Definitions of the WebSphere MQ structures and data types
- Constant definitions for initializing the structures
- Function prototypes for each of the calls (for PL/I and the C language only)

For a full description of WebSphere MQ data definition files, see [“IBM WebSphere MQ data definition files”](#) on page 77.

Structures

Structures, used with the MQI calls listed in [“MQI calls”](#) on page 188, are supplied in data definition files for each of the supported programming languages.

See [Structure data types summary](#) for a summary of the structures.

IBM WebSphere MQ stub programs and library files

The stub programs and library files provided are listed here, for each platform.

For more information about how to use stub programs and library files when you build an executable application, see [“Building an IBM WebSphere MQ application” on page 409](#). For information about linking to C++ library files, see [Using C++ WebSphere MQ Using C++](#).

IBM WebSphere MQ for Windows

On IBM WebSphere MQ for Windows, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system:

<i>Table 25. Library files for Windows applications</i>	
Library File	Environment
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqm.lib</code>	Server for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqic.lib</code>	Client for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmxa.lib</code>	Server XA interface for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqcxa.lib</code>	Client XA interface for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicxa.lib</code>	Client MTS for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcics4.lib</code>	Server TXSeries CICS support for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqccics4.lib</code>	Client TXSeries CICS support for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmzf.lib</code>	Installable services exits for C (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcbb.lib</code>	Server for IBM COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcb.lib</code>	Server for Micro Focus COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicbb.lib</code>	Client for IBM COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicb.lib</code>	Client for Micro Focus COBOL (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqs23vn.lib</code>	Server for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqc23vn.lib</code>	Client for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqb23vn.lib</code>	Base for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib\imqx23vn.lib</code>	Client MTS for C++ (32-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqm.lib</code>	Server for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqic.lib</code>	Client for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmxa.lib</code>	Server XA interface for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqcxa.lib</code>	Client XA interface for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicxa.lib</code>	Client MTS for C (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcbb.lib</code>	Server for IBM COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcb.lib</code>	Server for Micro Focus COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicbb.lib</code>	Client for IBM COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicb.lib</code>	Client for Micro Focus COBOL (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqs23vn.lib</code>	Server for C++ (64-bit)

<i>Table 25. Library files for Windows applications (continued)</i>	
Library File	Environment
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqc23vn.lib</code>	Client for C++ (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqb23vn.lib</code>	Base for C++ (64-bit)
<code>MQ_INSTALLATION_PATH\Tools\Lib64\imqx23vn.lib</code>	Client MTS for C++ (64-bit)

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Use `amqmdnet.dll` for compiling .NET programs. See “[Compiling WebSphere MQ .NET programs](#)” on page 567 within the section “[Using .NET](#)” on page 534 for more information.

These files are shipped for compatibility with previous releases:

`mqc32.lib`
`mqc32xa.lib`

IBM WebSphere MQ for AIX

On IBM WebSphere MQ for AIX, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

In a non-threaded application:

<i>Table 26. Library files for non-threaded AIX applications</i>	
Library file	Environment
<code>libmqm.a</code>	Server for C
<code>libmqic.a</code> & <code>libmqm.a</code>	Client for C
<code>libmqmzf.a</code>	Installable service exits for C
<code>libmqmxa.a</code>	Server XA interface
<code>libmqmxa64.a</code>	Server alternative XA interface
<code>libmqcxa.a</code>	Client XA interface
<code>libmqcxa64.a</code>	Client alternative XA interface
<code>libmqmcbt.o</code>	WebSphere MQ runtime library for Micro Focus COBOL support
<code>libmqmcb.a</code>	Server for COBOL
<code>libmqicb.a</code>	Client for COBOL
<code>libimqc23ia.a</code>	Client for C++
<code>libimqs23ia.a</code>	Server for C++

In a threaded application:

<i>Table 27. Library files for threaded AIX applications</i>	
Library file	Environment
<code>libmqm_r.a</code>	Server for C

<i>Table 27. Library files for threaded AIX applications (continued)</i>	
Library file	Environment
libmqic_r.a & libmqm_r.a	Client for C
libmqmzf_r.a	Installable service exits for C
libmqmxa_r.a	Server XA interface
libmqmxa64_r.a	Server alternative XA interface
libmqcxa_r.a	Client XA interface
libmqcxa64_r.a	Client alternative XA interface
libimqc23ia_r.a	Client for C++
libimqs23ia_r.a	Server for C++

IBM WebSphere MQ for HP-UX

On IBM WebSphere MQ for HP-UX, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

IA64 (IPF) platform

In a non-threaded application:

<i>Table 28. Library files for non-threaded HP-UX applications</i>	
Library file	Environment
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqi23ah.so	C++
libmqmcbirt.o	WebSphere MQ runtime library for Micro Focus COBOL support
libmqmcb.so	Server for COBOL
libmqicb.so	Client for COBOL

In a threaded application:

<i>Table 29. Library files for threaded HP-UX applications</i>	
Library file	Environment
libmqm_r.so	Server for C
libmqmzf_r.so & libmqm_r.so	Installable service exits for C
libmqmxa_r.so	Server XA interface

<i>Table 29. Library files for threaded HP-UX applications (continued)</i>	
Library file	Environment
libmqmxa64_r.so	Server alternative XA interface
libmqcxa_r.so	Client XA interface
libmqcxa64_r.so	Client alternative XA interface
libimqi23ah_r.so	C++

IBM WebSphere MQ for Linux

On IBM WebSphere MQ for Linux, you must link your program to the MQI library files supplied for the environment in which you are running your application, in addition to those provided by the operating system.

In a non-threaded application:

<i>Table 30. Library files for non-threaded Linux applications</i>	
Library file	Environment
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqc23gl.so	Client for C++
libimqs23gl.so	Server for C++

In a threaded application:

<i>Table 31. Library files for threaded Linux applications</i>	
Library file	Environment
libmqm_r.so	Server for C
libmqic_r.so & libmqm_r.so	Client for C
libmqmzf_r.so	Installable service exits for C
libmqmxa_r.so	Server XA interface
libmqmxa64_r.so	Server alternative XA interface
libmqcxa_r.so	Client XA interface
libmqcxa64_r.so	Client alternative XA interface
libimqc23gl_r.so	Client for C++
libimqs23gl_r.so	Server for C++

IBM WebSphere MQ for Solaris

On IBM WebSphere MQ for Solaris, you must link your program to the MQI library files supplied for the environment in which you are running your application in addition to those provided by the operating system.

Table 32. Library files for Solaris applications	
Library File	Environment
libmqm.so	Server and client for C
libmqmzse.so	For C
libmqic.so	Client for C
libmqmcs.so	Common services for C
libmqmzf.so	Installable service exits for C
libmqmxa.so	Server XA interface
libmqmxa64.so	Server alternative XA interface
libmqcxa.so	Client XA interface
libmqcxa64.so	Client alternative XA interface
libimqc23as.a	Client for C++
libimqs23as.a	Server for C++

Parameters common to all the calls

There are two types of parameter common to all the calls: handles and return codes.

Using handles

All MQI calls use one or more *handles*. These identify the queue manager, queue or other object, message, or subscription, as appropriate to the call.

For a program to communicate with a queue manager, the program must have a unique identifier by which it knows that queue manager. This identifier is called a *connection handle*, sometimes referred to as a *Hconn*. For CICS programs, the connection handle is always zero. For all other platforms or styles of programs, the connection handle is returned by the MQCONN or MQCONNX call when the program connects to the queue manager. Programs pass the connection handle as an input parameter when they use the other calls.

For a program to work with a WebSphere MQ object, the program must have a unique identifier by which it knows that object. This identifier is called an *object handle*, sometimes referred to as an *Hobj*. The handle is returned by the MQOPEN call when the program opens the object to work with it. Programs pass the object handle as an input parameter when they use subsequent MQPUT, MQGET, MQINQ, MQSET, or MQCLOSE calls.

Similarly, the MQSUB call returns a *subscription handle* or *Hsub*, which is used to identify the subscription in subsequent MQGET, MQCB or MQSUBRQ calls, and certain calls processing message properties use a *message handle* or *Hmsg*.

Understanding return codes

A completion code and a reason code are returned as output parameters by each call. These are known collectively as *return codes*.

To show whether a call is successful, each call returns a *completion code* when the call is complete. The completion code is typically either MQCC_OK indicating success, or MQCC_FAILED indicating failure. Some calls can return an intermediate state, MQCC_WARNING, indicating partial success.

Each call also returns a *reason code* that shows the reason for the failure, or partial success, of the call. There are many reason codes, covering such circumstances as a queue being full, get operations not being allowed for a queue, and a particular queue not being defined for the queue manager. Programs can use the reason code to decide how to proceed. For example, they can prompt users to change their input data, then make the call again, or they can return an error message to the user.

When the completion code is MQCC_OK, the reason code is always MQRC_NONE.

The completion and reason codes for each call are listed with the description of that call. See [Call descriptions](#) and select the appropriate call from the list.

For more detailed information, including ideas for corrective action, see:

- [Reason codes](#) for all other WebSphere MQ platforms

Specifying buffers

The queue manager refers to buffers only if they are required. If you do not require a buffer on a call or the buffer is zero in length, you can use a null pointer to a buffer.

Always use `datalength` when specifying the size of the buffer that you require.

When you use a buffer to hold the output from a call (for example, to hold the message data for an MQGET call, or the values of attributes queried by the MQINQ call), the queue manager attempts to return a reason code if the buffer you specify is not valid or is in read-only storage. However, it might not always be able to return a reason code.

UNIX and Linux considerations

Considerations that you need to be aware of.

Take note of the following points when developing UNIX and Linux applications.

The fork system call in UNIX and Linux systems

Note these considerations when using a fork system call in IBM WebSphere MQ applications.

If your application wants to use `fork`, the parent process of that application should call `fork` before making any IBM WebSphere MQ calls, for example, MQCONN, or creating an IBM WebSphere MQ object using **ImqQueueManager**.

If your application wants to create a child process after making any IBM WebSphere MQ calls, the application code must use a `fork()` with `exec()` to ensure that the child is a new instance, and not an exact copy of the parent.

If your application does not use `exec()`, the IBM WebSphere MQ API call made within the child process returns MQRC_ENVIRONMENT_ERROR.

UNIX and Linux signal handling

This does not apply to WebSphere MQ for z/OS or WebSphere MQ for Windows.

In general, UNIX, Linux and IBM i systems have moved from a nonthreaded (process) environment to a multithreaded environment. In the nonthreaded environment, some functions could be implemented only by using signals, though most applications did not need to be aware of signals and signal handling. In the multithreaded environment, thread-based primitives support some of the functions that used to be implemented in the nonthreaded environments using signals.

In many instances, signals and signal handling, although supported, do not fit well into the multithreaded environment and various restrictions exist. This can be problematic when you are integrating application code with different middleware libraries (running as part of the application) in a multithreaded environment where each is trying to handle signals. The traditional approach of saving and restoring signal handlers (defined per process), which worked when there was only one thread of execution within a process, does not work in a multithreaded environment. This is because many threads of execution could be trying to save and restore a process-wide resource, with unpredictable results.

Unthreaded applications

Not applicable on Solaris as all applications are considered threaded even if they use only a single thread.

Each MQI function sets up its own signal handler for the signals:

SIGALRM
SIGBUS
SIGFPE
SIGSEGV
SIGILL

Users' handlers for these are replaced for the duration of the MQI function call. Other signals can be caught in the normal way by user-written handlers. If you do not install a handler, the default actions (for example, ignore, core dump, or exit) are left in place.

After WebSphere MQ handles a synchronous signal (SIGSEGV, SIGBUS, SIGFPE, SIGILL), it attempts to pass the signal to any registered signal handler before making the MQI function call.

Threaded applications

A thread is considered to be connected to WebSphere MQ from MQCONN (or MQCONNX) until MQDISC.

Synchronous signals

Synchronous signals arise in a specific thread.

UNIX and Linux systems safely allow the setting up of a signal handler for such signals for the whole process. However, WebSphere MQ sets up its own handler for the following signals, in the application process, while any thread is connected to WebSphere MQ:

SIGBUS
SIGFPE
SIGSEGV
SIGILL

If you are writing multithreaded applications, there is only one process-wide signal handler for each signal. When WebSphere MQ sets up its own synchronous signal handlers it saves any previously registered handlers for each signal. After WebSphere MQ handles one of the signals listed, WebSphere MQ attempts to call the signal handler that was in effect at the time of the first WebSphere MQ connection within the process. The previously registered handlers are restored when all application threads have disconnected from WebSphere MQ.

Because signal handlers are saved and restored by WebSphere MQ, application threads must not establish signal handlers for these signals while there is any possibility that another thread of the same process is also connected to WebSphere MQ.

Note: When an application, or a middleware library (running as part of an application), establishes a signal handler while a thread is connected to WebSphere MQ, the application's signal handler must call the corresponding WebSphere MQ handler during the processing of that signal.

When establishing and restoring signal handlers, the general principle is that the last signal handler to be saved must be the first to be restored:

- When an application establishes a signal handler after connecting to WebSphere MQ, the previous signal handler must be restored before the application disconnects from WebSphere MQ.
- When an application establishes a signal handler before connecting to WebSphere MQ, the application must disconnect from WebSphere MQ before restoring its signal handler.

Note: Failure to observe the general principle that the last signal handler to be saved must be the first to be restored can result in unexpected signal handling in the application and, potentially, the loss of signals by the application.

Asynchronous signals

WebSphere MQ does not use any asynchronous signals in threaded applications unless they are client applications.

Additional considerations for threaded client applications

WebSphere MQ handles the following signals during I/O to a server. These signals are defined by the communications stack. The application must not establish a signal handler for these signals while a thread is connected to a queue manager:

SIGPIPE (for TCP/IP)

Additional considerations

Note these considerations when using UNIX signal handling.

Fastpath (trusted) applications

Fastpath applications run in the same process as WebSphere MQ and so are running in the multithreaded environment.

In this environment WebSphere MQ handles the synchronous signals SIGSEGV, SIGBUS, SIGFPE, and SIGILL. All other signals must not be delivered to the Fastpath application while it is connected to WebSphere MQ. Instead they must be blocked or handled by the application. If a Fastpath application intercepts such an event, the queue manager must be stopped and restarted, or it may be left in an undefined state. For a full list of the restrictions for Fastpath applications under MQCONNX, see [“Connecting to a queue manager using the MQCONNX call” on page 200.](#)

MQI function calls within signal handlers

While you are in a signal handler, do not call an MQI function.

If you try to call an MQI function from a signal handler while another MQI function is active, MQRC_CALL_IN_PROGRESS is returned. If you try to call an MQI function from a signal handler while no other MQI function is active, it is likely to fail sometime during the operation because of the operating system restrictions where only selective calls can be issued from, or within, a handler.

For C++ destructor methods, which might be called automatically during program exit, you might not be able to stop the MQI functions from being called. Ignore any errors about MQRC_CALL_IN_PROGRESS. If a signal handler calls exit(), WebSphere MQ backs out uncommitted messages in sync point as usual and closes any open queues.

Signals during MQI calls

MQI functions do not return the code EINTR or any equivalent to application programs.

If a signal occurs during an MQI call, and the handler calls *return*, the call continues to run as if the signal had not happened. In particular, MQGET cannot be interrupted by a signal to return control immediately to the application. If you want to break out of an MQGET, set the queue to GET_DISABLED; alternatively, use a loop around a call to MQGET with a finite time expiry (MQGMO_WAIT with gmo.WaitInterval set), and use your signal handler (in a nonthreaded environment) or equivalent function in a threaded environment to set a flag which breaks the loop.

In the AIX environment, WebSphere MQ requires that system calls interrupted by signals are restarted. When establishing your own signal handler with sigaction(2), set the SA_RESTART flag in the sa_flags field of the new action structure otherwise WebSphere MQ might be unable to complete any call interrupted by a signal.

User exits and installable services

User exits and installable services that run as part of a WebSphere MQ process in a multithreaded environment have the same restrictions as for fastpath applications. Consider these to be permanently connected to WebSphere MQ and so not using signals or non-threadsafe operating system calls.

VMS exit handlers

Users can install exit handlers for an WebSphere MQ application using the **SYS\$DCLEXH** system service.

The exit handler receives control when an image exits. An image exit normally occurs when you call the Exit (\$EXIT) or Force Exit (\$FORCEX) service. The \$FORCEX interrupts the target process in user mode. Then all user-mode exit handlers (established by \$DCLEXH) begin to execute in reverse order of establishment. For more details on exit handlers and \$FORCEX, see the *VMS Programming Concepts Manual* and the *VMS System Services Manual*.

If you call an MQI function from within an exit handler, the behavior of the function depends on the way the image was terminated. If the image was terminated while another MQI function is active, an MQRC_CALL_IN_PROGRESS is returned.

It is possible to call an MQI function from within an exit handler if no other MQI function is active and upcalls are disabled for the WebSphere MQ application. If upcalls are enabled for the WebSphere MQ application, it fails with the reason code MQRC_HCONN_ERROR.

The scope of an MQCONN or MQCONNEX call is typically the thread that issued it. If upcalls are enabled, the exit handler runs as a separate thread and the connection handles cannot be shared.

Exit handlers are started within the interrupted context of the target process. It is up to the application to ensure that actions taken by a handler are safe and reliable, for the asynchronously interrupted context they are called from.

Connecting to and disconnecting from a queue manager

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

The way that this connection is made depends on the platform and the environment in which the program is operating:

z/OS batch, WebSphere MQ for IBM i, WebSphere MQ on UNIX systems, WebSphere MQ on Linux systems, and WebSphere MQ for Windows

Programs that run in these environments can use the MQCONN MQI call to connect to, and the MQDISC call to disconnect from, a queue manager. Alternatively, programs can use the MQCONNEX call.

z/OS batch programs can connect, consecutively or concurrently, to multiple queue managers on the same TCB.

IMS

The IMS control region is connected to one or more queue managers when it starts. This connection is controlled by IMS commands. However, writers of message queuing IMS programs must use the MQCONN MQI call to specify the queue manager to which they want to connect. They can use the MQDISC call to disconnect from that queue manager.

Following an IMS call that establishes a syncpoint, and before processing a message for another user, the IMS adapter ensures that the application closes handles and disconnects from the queue manager.

IMS programs can connect, consecutively or concurrently, to multiple queue managers on the same TCB.

CICS Transaction Server for z/OS and CICS for MVS/ESA

CICS programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. This connection is typically made automatically at initialization, but you can also use the CKQC transaction, which is supplied with WebSphere MQ for z/OS.

CICS tasks can connect only to the queue manager to which the CICS region, itself, is connected.

Note: CICS programs can also use the MQI connect and disconnect calls (MQCONN and MQDISC). You might want to do this so that you can port these applications to non-CICS environments with a minimum of recoding. However, these calls *always* complete successfully in a CICS environment. This means that the return code might not reflect the true state of the connection to the queue manager.

TXSeries for Windows and Open Systems

These programs do not need to do any work to connect to a queue manager because the CICS system itself is connected. Therefore, only one connection at a time is supported. CICS applications must issue an MQCONN call to obtain a connection handle, and an MQDISC call before they exit.

Use the following links to find out more about connecting and disconnecting from a queue manager:

- [“Connecting to a queue manager using the MQCONN call” on page 199](#)
- [“Connecting to a queue manager using the MQCONN call” on page 200](#)
- [“Disconnecting programs from a queue manager using MQDISC” on page 205](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Connecting to a queue manager using the MQCONN call

Use this information to learn how to connect to a queue manager using the MQCONN call.

In general, you can connect either to a specific queue manager, or to the default queue manager:

- For IBM WebSphere MQ for z/OS, in the batch environment, the default queue manager is specified in the CSQBDEFV module.
- For IBM WebSphere MQ for Windows, IBM i, UNIX, and Linux systems, the default queue manager is specified in the mq.ini file.

Alternatively, in the z/OS MVS™ batch, TSO, and RRS environments you can connect to any one queue manager within a queue-sharing group. The MQCONN or MQCONN request selects any one of the active members of the group.

When you connect to a queue manager it must be local to the task. It must belong to the same system as the IBM WebSphere MQ application.

In the IMS environment, the queue manager must be connected to the IMS control region and to the dependent region that the program uses. The default queue manager is specified in the CSQQDEFV module when IBM WebSphere MQ for z/OS is installed.

With the TXSeries CICS environment, and TXSeries for Windows and AIX, the queue manager must be defined as an XA resource to CICS.

To connect to the default queue manager, call MQCONN, specifying a name consisting entirely of blanks or starting with a null (X'00') character.

An application must be authorized for it to successfully connect to a queue manager. For more information, see [Security](#).

The output from MQCONN is:

- A connection handle (**Hconn**)
- A completion code
- A reason code

Use the connection handle on subsequent MQI calls.

If the reason code indicates that the application is already connected to that queue manager, the connection handle that is returned is the same as the one that was returned when the application first connected. The application must not issue the MQDISC call in this situation because the calling application expects to remain connected.

The scope of the connection handle is the same as the scope of the object handle (see [“Opening objects using the MQOPEN call” on page 207](#)).

Descriptions of the parameters are given in the description of the MQCONN call in [MQCONN](#).

The MQCONN call fails if the queue manager is in a quiescing state when you issue the call, or if the queue manager is shutting down.

Scope of MQCONN or MQCONNX

The scope of an MQCONN or MQCONNX call is typically the thread that issued it. That is, the connection handle returned from the call is valid only within the thread that issued the call. Only one call can be made at any one time using the handle. If it is used from a different thread, it is rejected as invalid. If you have multiple threads in your application and each wants to use IBM WebSphere MQ calls, each one must issue MQCONN or MQCONNX.

It is not necessary for each call to be made to the same queue manager when a process makes multiple MQCONN calls. However, only one WebSphere MQ connection can be made from a thread at a time. Alternatively, consider [“Shared \(thread independent\) connections with MQCONNX” on page 204](#) to allow multiple WebSphere MQ connections from a single thread and an WebSphere MQ connection to be used from any thread.¹

If your application is running as a client, it can connect to more than one queue manager within a thread.

Connecting to a queue manager using the MQCONNX call

The MQCONNX call is similar to the MQCONN call, but includes options to control the way that the call works.

As input to MQCONNX, you can supply a queue manager name, or a queue-sharing group name on z/OS shared queue systems. The output from MQCONNX is:

- A connection handle (Hconn)

¹ When using multithreaded applications with IBM WebSphere MQ on UNIX and Linux systems you need to ensure that the applications have a sufficient stack size for the threads. Consider using a stack size of 256 KB, or larger, when multithreaded applications are making MQI calls, either by themselves or, with other signal handlers (for example, CICS).

- A completion code
- A reason code

You use the connection handle on subsequent MQI calls.

A description of all the parameters of MQCONN is given in [MQCONN](#). The *Options* field allows you to set STANDARD_BINDING, FASTPATH_BINDING, SHARED_BINDING, or ISOLATED_BINDING for any version of MQCNO. You can also make shared (thread independent) connections using a MQCONN call. See [“Shared \(thread independent\) connections with MQCONN” on page 204](#) for more information about these.

MQCNO_STANDARD_BINDING

By default, MQCONN (like MQCONN) implies two logical threads where the WebSphere MQ application and the local queue manager agent run in separate processes. The WebSphere MQ application requests the WebSphere MQ operation and the local queue manager agent services the request. This is defined by the MQCNO_STANDARD_BINDING option on the MQCONN call.

If you specify MQCNO_STANDARD_BINDING, the MQCONN call uses either MQCNO_SHARED_BINDING or MQCNO_ISOLATED_BINDING, depending on the value of the DefaultBindType attribute of the queue manager, which is defined in qm.ini or the Windows registry.

This is the default value.

If you are linking to the mqm library, then a standard server connection using the default bind type is attempted first. If the underlying server library failed to load, a client connection is attempted instead.

- If the MQ_CONNECT_TYPE environment variable is specified, one of the following options can be supplied to change the behaviour of MQCONN or MQCONN if MQCNO_STANDARD_BINDING is specified. (The exception to this is if MQCNO_FASTPATH_BINDING is specified with MQ_CONNECT_TYPE set to LOCAL or STANDARD to allow fastpath connections to be downgraded by the administrator without a related change to the application:

Value	Meaning
CLIENT	A client connection only is attempted.
FASTPATH	This value was supported in previous releases, but is now ignored if specified.
LOCAL	A server connection only is attempted. Fastpath connections are downgraded to a standard server connection.
STANDARD	Supported for compatibility with previous releases. This value is now treated as LOCAL.

- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

MQCNO_FASTPATH_BINDING

Trusted applications imply that the WebSphere MQ application and the local queue manager agent become the same process. Because the agent process no longer needs to use an interface to access the queue manager, these applications become an extension of the queue manager. This is defined by the MQCNO_FASTPATH_BINDING option on the MQCONN call.

You need to link trusted applications to the threaded WebSphere MQ libraries. For instructions on how to set up a WebSphere MQ application to run as trusted, see [MQCNO Options](#).

This option gives the highest performance.

Note: This option compromises the integrity of the queue manager: there is no protection from overwriting its storage. This also applies if the application contains errors that can be exposed

to messages and other data in the queue manager too. Consider these issues before using this option.

MQCNO_SHARED_BINDING

Specify this option to make the application and the local queue manager agent run in separate processes. This maintains the integrity of the queue manager, that is, it protects the queue manager from errant programs. However, the application and the local-queue-manager agent share some resources.

This option is intermediate between MQCNO_FASTPATH_BINDING and MQCNO_ISOLATED_BINDING, both in terms of protecting the integrity of the queue manager, and in terms of the performance of MQI calls.

MQCNO_SHARED_BINDING is ignored if the queue manager does not support this type of binding. Processing continues as though the option had not been specified.

If an application has connected to the local queue manager using MQCNO_SHARED_BINDING, the queue manager can be stopped while the application is running. If you restart the queue manager while the application is still running, the attempt to start the queue manager fails with error AMQ7018 as the application is still holding on to resources needed by the queue manager.

In order to start the queue manager, you must stop the application.

MQCNO_ISOLATED_BINDING

Specify this option to make the application and the local queue manager agent run in separate processes, as for MQCNO_SHARED_BINDING. In this case, however, the application process and the local-queue-manager agent are isolated from each other in that they do not share resources.

This is the safest option for protecting the integrity of the queue manager, but it gives the slowest performance of MQI calls.

MQCNO_ISOLATED_BINDING is ignored if the queue manager does not support this type of binding. Processing continues as though the option had not been specified.

MQCNO_CLIENT_BINDING

Specify this option to make the application attempt a client connection only. This option has the following limitations:

- MQCNO_CLIENT_BINDING is rejected on z/OS with MQRC_OPTIONS_ERROR.
- MQCNO_CLIENT_BINDING is rejected with MQRC_OPTIONS_ERROR if it is specified with any MQCNO binding option other than MQCNO_STANDARD_BINDING.
- MQCNO_CLIENT_BINDING is not available for Java as it has its own mechanisms for choosing the bind type.
- **V7.5.0.7** Before IBM WebSphere MQ Version 7.5.0, Fix Pack 7, MQCNO_CLIENT_BINDING is not available for .NET as it has its own mechanisms for choosing the bind type. From Version 7.5.0, Fix Pack 7, the restriction on using .NET for MQCNO_CLIENT_BINDING is removed.
- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

MQCNO_LOCAL_BINDING

Specify this option to make the application attempt a server connection. If either MQCNO_FASTPATH_BINDING, MQCNO_ISOLATED_BINDING, or MQCNO_SHARED_BINDING is also specified, then the connection is of that type instead, and is documented in this section. Otherwise a standard server connection is attempted using the default bind type. MQCNO_LOCAL_BINDING has the following limitations:

- MQCNO_LOCAL_BINDING is ignored on z/OS.

- MQCNO_LOCAL_BINDING is rejected with MQRC_OPTIONS_ERROR if it is specified with any MQCNO reconnect option other than MQCNO_RECONNECT_AS_DEF.
- MQCNO_LOCAL_BINDING is not available for Java as it has its own mechanisms for choosing the bind type.
- **V7.5.0.7** Before IBM WebSphere MQ Version 7.5.0, Fix Pack 7, MQCNO_LOCAL_BINDING is not available for .NET as it has its own mechanisms for choosing the bind type. From Version 7.5.0, Fix Pack 7, the restriction on using .NET for MQCNO_LOCAL_BINDING is removed.
- If the MQ_CONNECT_TYPE environment variable is not set when MQCONN is called, a standard server connection using the default bind type is attempted. If the server library fails to load, a client connection is attempted.

On z/OS these options are tolerated, but only a standard bound connection is performed. MQCNO Version 3, for z/OS, allows four alternative options:

MQCNO_SERIALIZE_CONN_TAG_QSG

This allows an application to request that only one instance of an application runs at any one time in a queue-sharing group. This is achieved by registering the use of a connection tag with a value that is specified or derived by the application. The tag is a 128 byte character string specified in the Version 3 MQCNO.

MQCNO_RESTRICT_CONN_TAG_QSG

This is used where an application consists of more than one process (or a TCB), each of which can connect to a queue manager. Connection is permitted only if there is no current use of the tag, or the requesting application is within the same processing scope. This is MVS address space within the same queue-sharing group as the tag owner.

MQCNO_SERIALIZE_CONN_TAG_Q_MGR

This is similar to MQCNO_SERIALIZE_CONN_TAG_QSG, but only the local queue manager is interrogated to see if the requested tag is already in use.

MQCNO_RESTRICT_CONN_TAG_Q_MGR

This is similar to MQCNO_RESTRICT_CONN_TAG_QSG, but only the local queue manager is interrogated to see if the requested tag is already in use.

Restrictions for trusted applications

The following restrictions apply to trusted applications:

- You must explicitly disconnect trusted applications from the queue manager.
- You must stop trusted applications before ending the queue manager with the endmqm command.
- You must not use asynchronous signals and timer interrupts (such as sigkill) with MQCNO_FASTPATH_BINDING.
- On all platforms, a thread within a trusted application cannot connect to a queue manager while another thread in the same process is connected to a different queue manager.
- On WebSphere MQ on UNIX and Linux systems you must use mqm as the effective userID and groupID for all MQI calls. You can change these IDs before making a non-MQI call requiring authentication (for example, opening a file), but you *must* change it back to mqm before making the next MQI call.
- On WebSphere MQ for HP-UX, multithreaded fast-path applications are likely to need to set a larger stack size than the default. Use a size of 256 KB.
- On WebSphere MQ for Windows trusted 64-bit applications are not supported. If you try to run a trusted 64-bit application, it will be downgraded to a standard bound connection.
- On WebSphere MQ on UNIX and Linux systems trusted 32-bit applications are not supported. If you try to run a trusted 32-bit application, it will be downgraded to a standard bound connection.

Shared (thread independent) connections with MQCONN

Use this information to learn about Shared connections with MQCONN, and some usage notes to consider.

Note: Not supported on WebSphere MQ for z/OS.

On WebSphere MQ platforms other than WebSphere MQ for z/OS, a connection made with MQCONN is available only to the thread that made the connection. Options on the MQCONN call allow you to create a connection that can be shared by all the threads in a process. If your application is running in a transactional environment that requires MQI calls to be issued on the same thread, you must use the following default option:

MQCNO_HANDLE_SHARE_NONE

Creates a non-shared connection.

In most other environments, you can use one of the following thread independent, shared connection options:

MQCNO_HANDLE_SHARE_BLOCK

Creates a shared connection. On a MQCNO_HANDLE_SHARE_BLOCK connection, if the connection is currently in use by an MQI call on another thread, the MQI call waits until the current MQI call has completed.

MQCNO_HANDLE_SHARE_NO_BLOCK

Creates a shared connection. On a MQCNO_HANDLE_SHARE_NO_BLOCK connection, if the connection is currently in use by an MQI call on another thread, the MQI call fails immediately with a reason of MQRC_CALL_IN_PROGRESS.

Except for the MTS (Microsoft Transaction Server) environment, the default value is MQCNO_HANDLE_SHARE_NONE. In the MTS environment, the default value is MQCNO_HANDLE_SHARE_BLOCK.

A connection handle is returned from the MQCONN call. The handle can be used by subsequent MQI calls from any thread in the process, associating those calls with the handle returned from the MQCONN. MQI calls using a single shared handle are serialized across threads.

For example, the following sequence of activity is possible with a shared handle:

1. Thread 1 issues MQCONN and gets a shared handle *h1*
2. Thread 1 opens a queue and issues a get request using *h1*
3. Thread 2 issues a put request using *h1*
4. Thread 3 issues a put request using *h1*
5. Thread 2 issues MQDISC using *h1*

While the handle is in use by any thread, access to the connection is unavailable to other threads. In circumstances where it is acceptable that a thread waits for any previous call from another thread to complete, use MQCONN with the option MQCNO_HANDLE_SHARE_BLOCK.

However blocking can cause difficulties. Suppose that in step “2” on page 204, thread 1 issues a get request that waits for messages that might not have yet arrived (a get with wait). In this case, threads 2 and 3 are also left waiting (blocked) for as long as the get request on thread 1 takes. If you prefer that an MQI call returns with an error if another MQI call is already running on the handle, use MQCONN with the option MQCNO_HANDLE_SHARE_NO_BLOCK.

Shared connection usage notes

1. Any object handles (Hobj) created by opening an object are associated with an Hconn; so for a shared Hconn, the Hobjs are also shared and usable by any thread using the Hconn. Similarly, any unit of work started under an Hconn is associated with that Hconn; so this too is shared across threads with the shared Hconn.
2. Any thread can call MQDISC to disconnect a shared Hconn, not just the thread that called the corresponding MQCONN. The MQDISC terminates the Hconn making it unavailable to all threads.

3. A single thread can use multiple shared Hconns serially, for example use MQPUT to put one message under one shared Hconn then put another message using another shared Hconn, with each operation being under a different local unit of work.
4. Shared Hconns cannot be used within a global unit of work.

Use of MQCONN call options with MQ_CONNECT_TYPE

Use this information to understand the different MQCONN call options how they are used with MQ_CONNECT_TYPE.

On WebSphere MQ for IBM i, WebSphere MQ for Windows, and WebSphere MQ on UNIX and Linux systems, you can use the environment variable, MQ_CONNECT_TYPE in combination with the type of binding specified in the *Options* field of the MQCNO structure used on an MQCONN call.

Table 33. The MQ_CONNECT_TYPE environment variable		
MQCONN call option	MQ_CONNECT_TYPE environment variable	Result
STANDARD	UNDEFINED	STANDARD
STANDARD	STANDARD	STANDARD
STANDARD	FASTPATH	STANDARD
STANDARD	CLIENT	CLIENT
STANDARD	LOCAL	STANDARD

If MQCNO_STANDARD_BINDING is not specified, you can use MQCNO_NONE, which defaults to MQCNO_STANDARD_BINDING.

Disconnecting programs from a queue manager using MQDISC

Use this information to learn about disconnecting programs from a queue manager using MQDISC.

When a program that has connected to a queue manager using the MQCONN or MQCONN call has finished all interaction with the queue manager, it breaks the connection using the MQDISC call, except:

- On CICS Transaction Server for z/OS applications, where the call is optional unless MQCONN was used and you want to drop the connection tag before the application ends.
- On WebSphere MQ for IBM i where, when you sign off from the operating system, an implicit MQDISC call is made.

As input to the MQDISC call, you must supply the connection handle (Hconn) that was returned by MQCONN or MQCONN when you connected to the queue manager.

Except on CICS on z/OS, after MQDISC is called the connection handle (Hconn) is no longer valid, and you cannot issue any further MQI calls until you call MQCONN or MQCONN again. MQDISC does an implicit MQCLOSE for any objects that are still open using this handle.

If you use MQCONN to connect on WebSphere MQ for z/OS, MQDISC also ends the scope of the connection tag established by the MQCONN. However, in a CICS, IMS, or RRS application, if there is an active unit of recovery associated with a connection tag, the MQDISC is rejected with a reason code of MQRC_CONN_TAG_NOT_RELEASED.

Descriptions of the parameters are given in the description of the MQDISC call in [MQDISC](#).

When no MQDISC is issued

A standard, non-shared connection (Hconn) is cleaned up when the creating thread terminates. A shared connection is only implicitly backed out and disconnected when the whole process terminates. If the thread that created the shared Hconn terminates while the Hconn still exists the Hconn is still usable.

Authority checking

The MQCLOSE and MQDISC calls usually perform no authority checking.

In the normal course of events a job that has the authority to open or connect to a WebSphere MQ object closes or disconnect from that object. Even if the authority of a job that has connected to or opened a WebSphere MQ object is revoked, the MQCLOSE and MQDISC calls are accepted.

Opening and closing objects

This information provides an insight into opening and closing WebSphere MQ objects.

To perform any of the following operations, you must first *open* the relevant WebSphere MQ object:

- Put messages on a queue
- Get (browse or retrieve) messages from a queue
- Set the attributes of an object
- Inquire about the attributes of any object

Use the MQOPEN call to open the object, using the options of the call to specify what you want to do with the object. The only exception is if you want to put a single message on a queue, then close the queue immediately. In this case, you can bypass the *opening* stage by using the MQPUT1 call (see [“Putting one message on a queue using the MQPUT1 call”](#) on page 224).

Before you open an object using the MQOPEN call, you must connect your program to a queue manager. This is explained in detail, for all environments, in [“Connecting to and disconnecting from a queue manager”](#) on page 198.

There are four types of WebSphere MQ object that you can open:

- Queue
- Namelist
- Process definition
- Queue manager

You open all these objects in a similar way using the MQOPEN call. For more information about WebSphere MQ objects, see [Objects](#).

You can open the same object more than once, and each time you get a new object handle. You might want to browse messages on a queue using one handle, and remove messages from the same queue using another handle. This saves using up resources to close and reopen the same object. You can also open a queue for browsing *and* removing messages at the same time.

Moreover, you can open multiple objects with a single MQOPEN and close them using MQCLOSE. See [“Distribution lists”](#) on page 225 for information about how to do this.

When you attempt to open an object, the queue manager checks that you are authorized to open that object for the options that you specify in the MQOPEN call.

Objects are closed automatically when a program disconnects from the queue manager. In the IMS environment, disconnection is forced when a program starts processing for a new user following a GU (get unique) IMS call. On the IBM i platform, objects are closed automatically when a job ends.

It is good programming practice to close objects you have opened. Use the MQCLOSE call to do this.

Use the following links to find out more about opening and closing objects:

- [“Opening objects using the MQOPEN call”](#) on page 207
- [“Creating dynamic queues”](#) on page 214
- [“Opening remote queues”](#) on page 215
- [“Closing objects using the MQCLOSE call”](#) on page 215

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager.

Use this information to learn how to connect to and disconnect from a queue manager.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Opening objects using the MQOPEN call

Use this information to learn about opening objects using the MQOPEN call.

As input to the MQOPEN call, you must supply:

- A connection handle. For CICS applications on z/OS, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNEX call. For other programs, always use the connection handle returned by the MQCONN or MQCONNEX call.
- A description of the object that you want to open, using the object descriptor structure (MQOD).
- One or more options that control the action of the call.

The output from MQOPEN is:

- An object handle that represents your access to the object. Use this on input to any subsequent MQI calls.
- A modified object-descriptor structure, if you are creating a dynamic queue (and it is supported on your platform).
- A completion code.
- A reason code.

Scope of an object handle

The scope of an object handle (Hobj) is the same as the scope of a connection handle (Hconn).

This is covered in [“Scope of MQCONN or MQCONNEX” on page 200](#) and [“Shared \(thread independent\) connections with MQCONNEX” on page 204](#). However, there are additional considerations in some environments:

CICS

In a CICS program, you can use the handle only within the same CICS task from which you made the MQOPEN call.

IMS and z/OS batch

In the IMS and batch environments, you can use the handle within the same task, but not within any subtasks.

Descriptions of the parameters of the MQOPEN call are given in [MQOPEN](#).

The following sections describe the information that you must supply as input to MQOPEN.

Identifying objects (the MQOD structure)

Use the MQOD structure to identify the object that you want to open. This structure is an input parameter for the MQOPEN call. (The structure is modified by the queue manager when you use the MQOPEN call to create a dynamic queue.)

For full details of the MQOD structure, see [MQOD](#).

For information about using the MQOD structure for distribution lists, see [“Using the MQOD structure”](#) on page 227 under [“Distribution lists”](#) on page 225.

Name resolution

How the MQOPEN call resolves queue and queue manager names.

Note: A Queue manager alias is a remote queue definition without an RNAME field.

When you open a WebSphere MQ queue, the MQOPEN call performs a name resolution function on the queue name that you specify. This determines on which queue the queue manager performs subsequent operations. This means that when you specify the name of an alias queue or a remote queue in your object descriptor (MQOD), the call resolves the name either to a local queue or to a transmission queue. If a queue is opened for any type of input, browse, or set, it resolves to a local queue if there is one, and fails if there is not one. It resolves to a nonlocal queue only if it is opened for output only, inquire only, or output and inquire only. See [Table 34 on page 208](#) for an overview of the name resolution process. The name that you supply in *ObjectQMGrName* is resolved *before* that in *ObjectName*.

[Table 34 on page 208](#) also shows how you can use a local definition of a remote queue to define an alias for the name of a queue manager. This allows you to select which transmission queue is used when you put messages on a remote queue, so you could, for example, use a single transmission queue for messages destined for many remote queue managers.

To use the following table, first read down the two left-hand columns, under the heading **Input to MQOD**, and select the appropriate case. Then read across the corresponding row, following any instructions. Following the instructions in the **Resolved names** columns, you can either return to the **Input to MQOD** columns and insert values as directed, or you can exit the table with the results supplied. For example, you might be required to input *ObjectName*.

Table 34. Resolving queue names when using MQOPEN				
Input to MQOD		Resolved names		
<i>ObjectQMGrName</i>	<i>ObjectName</i>	<i>ObjectQMGrName</i>	<i>ObjectName</i>	Transmission queue
Blank or local queue manager	Local queue with no CLUSTER attribute	Local queue manager	Input <i>ObjectName</i>	Not applicable (local queue used)
Blank queue manager	Local queue with CLUSTER attribute	Workload management selected cluster queue manager or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER.TRANSMIT.QUEUE and local queue used SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Local queue manager	Local queue with CLUSTER attribute	Local queue manager	Input <i>ObjectName</i>	Not applicable (local queue used)

Table 34. Resolving queue names when using MQOPEN (continued)				
Input to MQOD		Resolved names		
<i>ObjectQMgrName</i>	<i>ObjectName</i>	<i>ObjectQMgrName</i>	<i>ObjectName</i>	Transmission queue
Blank or local queue manager	Model queue	Local queue manager	Generated name	Not applicable (local queue used)
Blank or local queue manager	Alias queue with or without CLUSTER attribute	Perform name resolution again with <i>ObjectQMgrName</i> unchanged, and input <i>ObjectName</i> set to the <i>BaseQName</i> in the alias queue definition object. Must not resolve to an alias locally defined where the <i>ObjectQMgrName</i> is specified, but can resolve to a clustered alias (hosted on other queue managers) where the <i>ObjectQMgrName</i> is blank.		
Local queue manager	Alias queue with CLUSTER attribute	The alias must not resolve to a cluster queue that is not locally defined, or a cluster queue that has the same <i>ObjectName</i> as the alias.		
Blank queue manager	Alias queue with CLUSTER attribute	The alias can resolve to a cluster queue with same <i>ObjectName</i> as the alias.		
Blank or local queue manager	Local definition of a remote queue	Perform name resolution again with <i>ObjectQMgrName</i> set to <i>RemoteQMgrName</i> , and <i>ObjectName</i> set to <i>RemoteQName</i> . Must not resolve remote queues		Name of <i>XmitQName</i> attribute, if non-blank; otherwise <i>RemoteQMgrName</i> in the remote queue definition object. SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Blank queue manager	No matching local object; cluster queue found	Workload management selected cluster queue manager or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER.TRANSMIT.QUEUE SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Blank or local queue manager	No matching local object; cluster queue not found		Error, queue not found	Not applicable

Table 34. Resolving queue names when using MQOPEN (continued)				
Input to MQOD		Resolved names		
<i>ObjectQMgrName</i>	<i>ObjectName</i>	<i>ObjectQMgrName</i>	<i>ObjectName</i>	Transmission queue
Name of queue manager in same queue sharing group as local queue manager	Local shared queue	Local queue manager	Input <i>ObjectName</i>	Not applicable
Name of a local transmission queue	(Not resolved)	Input <i>ObjectQMgrName</i>	Input <i>ObjectName</i>	Input <i>ObjectQMgrName</i> SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Queue manager alias definition (<i>RemoteQMgrName</i> may be the local queue manager)	(Not resolved, remote queue)	Perform name resolution again with <i>ObjectQMgrName</i> set to <i>RemoteQMgrName</i> . Must not resolve to remote queues	Input <i>ObjectName</i>	Name of <i>XmitQName</i> attribute, if non-blank; otherwise <i>RemoteQMgrName</i> in the remote queue definition object. SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Queue manager is not the name of any local object; cluster queue managers or queue manager alias found	(Not resolved)	<i>ObjectQMgrName</i> or specific cluster queue manager selected on PUT	Input <i>ObjectName</i>	SYSTEM.CLUSTER.TRANSMIT.QUEUE SYSTEM.QSG.TRANSMIT.QUEUE (see note)
Queue manager is not the name of any local object; no cluster objects found	(Not resolved)	Input <i>ObjectQMgrName</i>	Input <i>ObjectName</i>	<i>DefXmitQName</i> attribute of the queue manager where <i>DefXmitQName</i> is supported. SYSTEM.QSG.TRANSMIT.QUEUE (see note)

Notes:

1. *BaseQName* is the name of the base queue from the definition of the alias queue.
2. *RemoteQName* is the name of the remote queue from the local definition of the remote queue.
3. *RemoteQMgrName* is the name of the remote queue manager from the local definition of the remote queue.
4. *XmitQName* is the name of the transmission queue from the local definition of the remote queue.
5. When using WebSphere MQ for z/OS queue managers that are part of a queue-sharing group (QSG), the name of the QSG can be used instead of the local queue manager name in [Table 34 on page 208](#).

If the local queue manager cannot open the target queue, or put a message to the queue, the message is transferred to the specified *ObjectQMgrName* through, either intra-group queueing, or a WebSphere MQ channel.

6. In the *ObjectName* column of the table, CLUSTER refers to both the CLUSTER and CLUSNL attributes of the queue.
7. The SYSTEM.QSG.TRANSMIT.QUEUE is used if local and remote queue managers are in the same queue-sharing group; intra-group queueing is enabled.

8. If you have assigned a different cluster transmission queue to each cluster-sender channel, `SYSTEM.CLUSTER.TRANSMIT.QUEUE` might not be the name of the cluster transmission queue. For more information about multiple cluster transmission queues, see [Clustering: Planning how to configure cluster transmission queues](#).
9. In the situation where the queue manager is not the name of any local object; cluster queue managers, or queue manager alias found.

When you have supplied a queue manager name using **ObjectQMgrName**, and there are multiple cluster channels with different cluster names known by the local queue manager that would reach that destination, then any of these channels might be used to move the message, regardless of the cluster name of the destination queue.

This might be unexpected, if you were anticipating messages for that queue only to be sent through a channel that has the same cluster name as the queue.

However, the **ObjectQMgrName** takes precedence in this case, and cluster workload balancing takes into consideration all channels that might reach that queue manager, regardless of the cluster name they are in.

Opening an alias queue also opens the base queue to which the alias resolves, and opening a remote queue also opens the transmission queue. Therefore you cannot delete either the queue that you specify or the queue to which it resolves while the other one is open.

While an alias queue is unable to resolve to another locally defined alias queue (shared in a cluster or not), resolving to a remotely defined cluster alias queue is permitted and can therefore be specified as the base queue.

The resolved queue name and the resolved queue manager name are stored in the *ResolvedQName* and *ResolvedQMgrName* fields in the MQOD.

For more information about name resolution in a distributed queuing environment see [What is queue name resolution?](#).

Using the options of the MQOPEN call

In the *Options* parameter of the MQOPEN call, you must choose one or more options to control the access that you are given to the object that you are opening. With these options you can:

- Open a queue and specify that all messages put to that queue must be directed to the same instance of it
- Open a queue to allow you to put messages on it
- Open a queue to allow you to browse messages on it
- Open a queue to allow you to remove messages from it
- Open an object to allow you to inquire about and set its attributes (but you can set the attributes of queues only)
- Open a topic or topic string to publish messages to it
- Associate context information with a message
- Nominate an alternative user identifier to be used for security checks
- Control the call if the queue manager is in a quiescing state

MQOPEN option for cluster queue

The binding used for the queue handle is taken from the *DefBind* queue attribute, which can take the value `MQBND_BIND_ON_OPEN`, `MQBND_BIND_NOT_FIXED`, or `MQBND_BIND_ON_GROUP`.

To route all messages put to a queue using MQPUT to the same queue manager by the same route, use the `MQOO_BIND_ON_OPEN` option on the MQOPEN call.

To specify that a destination is to be selected at MQPUT time, that is, on a message-by-message basis, use the `MQOO_BIND_NOT_FIXED` option on the MQOPEN call.

To specify that all messages in a message groups put to a queue using MQPUT are allocated to the same destination instance, use the MQOO_BIND_ON_GROUP option on the MQOPEN call.

Either MQOO_BIND_ON_OPEN or MQOO_BIND_ON_GROUP must be specified when using message groups with clusters to ensure that all messages in the group are processed at the same destination.

If you do not specify any of these options the default, MQOO_BIND_AS_Q_DEF, is used.

If you specify the name of a queue manager in the MQOD, the queue at that queue manager is selected. If the queue manager name is blank, any instance can be selected. See [“MQOPEN and clusters”](#) on page 332 for more information.

If you open a cluster queue using a QALIAS definition, some queue attributes are defined by the alias queue, and not the base queue. Cluster attributes are among the attributes of the base queue definition that are overridden by the alias queue. For example, in the following snippet, the cluster queue is opened with MQOO_BIND_NOT_FIXED and not MQOO_BIND_ON_OPEN. The cluster queue definition is advertised throughout the cluster, the alias queue definition is local to the queue manager.

```
DEFINE QLOCAL(CLQ1) CLUSTER(MYCLUSTER) DEFBIND(OPEN) REPLACE
DEFINE QALIAS(ACLQ1) TARGET(CLQ1) DEFBIND(NOTFIXED) REPLACE
```

MQOPEN option for putting messages

To open a queue or topic to put messages on it, use the MQOO_OUTPUT option.

MQOPEN option for browsing messages

To open a queue so that you can *browse* the messages on it, use the MQOPEN call with the MQOO_BROWSE option.

This creates a *browse cursor* that the queue manager uses to identify the next message on the queue. For more information, see [“Browsing messages on a queue”](#) on page 260.

Note:

1. You cannot browse messages on a remote queue; do not open a remote queue using the MQOO_BROWSE option.
2. You cannot specify this option when opening a distribution list. For further information about distribution lists, see [“Distribution lists”](#) on page 225.
3. Use the MQOO_CO_OP in conjunction with MQOO_BROWSE if you are using cooperative browsing; see [Options](#)

MQOPEN options for removing messages

Three options control the opening of a queue to remove messages from it.

You can use only one of them in any MQOPEN call. These options define whether your program has exclusive or shared access to the queue. *Exclusive access* means that, until you close the queue, only you can remove messages from it. If another program attempts to open the queue to remove messages, its MQOPEN call fails. *Shared access* means that more than one program can remove messages from the queue.

The most advisable approach is to accept the type of access that was intended for the queue when the queue was defined. The queue definition involved the setting of the *Shareability* and the *DefInputOpenOption* attributes. To accept this access, use the MQOO_INPUT_AS_Q_DEF option. Refer to [Table 35 on page 212](#) to see how the setting of these attributes affects the type of access that you will be given when you use this option.

Table 35. How queue attributes and options of the MQOPEN call affect access to queues				
Queue attributes		Type of access with MQOPEN options		
<i>Shareability</i>	<i>DefInputOpenOption</i>	AS_Q_DEF	SHARED	EXCLUSIVE
SHAREABLE	SHARED	shared	shared	exclusive
SHAREABLE	EXCLUSIVE	exclusive	shared	exclusive

Table 35. How queue attributes and options of the MQOPEN call affect access to queues (continued)				
Queue attributes		Type of access with MQOPEN options		
NOT_SHAREABLE*	SHARED*	exclusive	exclusive	exclusive
NOT_SHAREABLE	EXCLUSIVE	exclusive	exclusive	exclusive
Note: * Although you can define a queue to have this combination of attributes, the default input open option is overridden by the shareability attribute.				

Alternatively:

- If you know that your application can work successfully even if other programs can remove messages from the queue at the same time, use the MQOO_INPUT_SHARED option. [Table 35 on page 212](#) shows how, in some cases you will be given exclusive access to the queue, even with this option.
- If you know that your application can work successfully only if other programs are prevented from removing messages from the queue at the same time, use the MQOO_INPUT_EXCLUSIVE option.

Note:

1. You cannot remove messages from a remote queue. Therefore you cannot open a remote queue using any of the MQOO_INPUT_* options.
2. You cannot specify this option when opening a distribution list. For further information, see [“Distribution lists” on page 225](#).

MQOPEN options for setting and inquiring about attributes

To open a queue so that you can set its attributes, use the MQOO_SET option.

You cannot set the attributes of any other type of object (see [“Inquiring about and setting object attributes” on page 306](#)).

To open an object so that you can inquire about its attributes, use the MQOO_INQUIRE option.

Note: You cannot specify this option when opening a distribution list.

MQOPEN options relating to message context

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

The options allow you to differentiate between context information that relates to the *user* who originated the message, and that which relates to the *application* that originated the message. Also, you can opt to set the context information when you put the message on the queue, or you can opt to have the context taken automatically from another queue handle.

Related concepts

[“Message context” on page 37](#)

Message context information allows the application that retrieves the message to find out about the originator of the message.

[“Controlling context information” on page 222](#)

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

MQOPEN option for alternative user authority

When you attempt to open an object using the MQOPEN call, the queue manager checks that you have the authority to open that object. If you are not authorized, the call fails.

However, server programs might want the queue manager to check the authorization of the user they are working for, rather than the server's own authorization. To do this, they must use the MQOO_ALTERNATE_USER_AUTHORITY option of the MQOPEN call, and specify the alternative user ID

in the *AlternateUserId* field of the MQOD structure. Typically, the server would get the user ID from the context information in the message it is processing.

MQOPEN option for queue manager quiescing

In the CICS environment on z/OS, if you use the MQOPEN call when the queue manager is in a quiescing state, the call always fails.

In other z/OS environments, IBM i, Windows systems and in UNIX and Linux systems environments, the call fails when the queue manager is quiescing only if you use the MQOO_FAIL_IF QUIESCING option of the MQOPEN call.

MQOPEN option for resolving local queue names

When you open a local, alias or model queue, the local queue is returned.

However, when you open a remote queue or cluster queue, the *ResolvedQName* and *ResolvedQMGrName* fields of the MQOD structure are filled with the names of the remote queue and remote queue manager found in the remote queue definition, or with the chosen remote cluster queue.

Use the MQOO_RESOLVE_LOCAL_Q option of the MQOPEN call to fill the *ResolvedQName* in the MQOD structure with the name of the local queue that was opened. The *ResolvedQMGrName* is similarly filled with the name of the local queue manager hosting the local queue. This field is available only with Version 3 of the MQOD structure; if the structure is less than Version 3, MQOO_RESOLVE_LOCAL_Q is ignored without an error being returned.

If you specify MQOO_RESOLVE_LOCAL_Q when opening, for example, a remote queue, *ResolvedQName* is the name of the transmission queue to which messages will be put. *ResolvedQMGrName* is the name of the local queue manager hosting the transmission queue.

Creating dynamic queues

Use a dynamic queue when you do not need the queue after your application ends.

For example, you could use a dynamic queue for your reply-to queue. You specify the name of the reply-to queue in the *ReplyToQ* field of the MQMD structure when you put a message on a queue (see [“Defining messages using the MQMD structure”](#) on page 217).

To create a dynamic queue, you use a template known as a model queue, together with the MQOPEN call. You create a model queue using the WebSphere MQ commands or the operations and control panels. The dynamic queue that you create takes the attributes of the model queue.

When you call MQOPEN, specify the name of the model queue in the *ObjectName* field of the MQOD structure. When the call completes, the *ObjectName* field is set to the name of the dynamic queue that is created. Also, the *ObjectQMGrName* field is set to the name of the local queue manager.

You can specify the name of the dynamic queue that you create in three ways:

- Give the full name that you want in the *DynamicQName* field of the MQOD structure.
- Specify a prefix (fewer than 33 characters) for the name, and allow the queue manager to generate the rest of the name. This means that the queue manager generates a unique name, but you still have some control (for example, you might want each user to use a certain prefix, or you might want to give a special security classification to queues with a certain prefix in their name). To use this method, specify an asterisk (*) for the last non-blank character of the *DynamicQName* field. Do not specify a single asterisk (*) for the dynamic queue name.
- Allow the queue manager to generate the full name. To use this method, specify an asterisk (*) in the first character position of the *DynamicQName* field.

For more information about these methods, see the description of the [DynamicQName](#) field.

There is more information on dynamic queues in [Dynamic and Model queues](#).

Opening remote queues

A remote queue is a queue that is owned by a queue manager other than the one to which the application is connected.

To open a remote queue, use the MQOPEN call as for a local queue. You can specify the name of the queue as follows:

1. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue as known to the *local* queue manager.

Note: Leave the *ObjectQMGrName* field blank in this case.

2. In the *ObjectName* field of the MQOD structure, specify the name of the remote queue, as known to the *remote* queue manager. In the *ObjectQMGrName* field, specify either:

- The name of the transmission queue that has the same name as the remote queue manager. The name and case (uppercase, lowercase or a mixture) must match *exactly*.
- The name of a queue manager alias object that resolves to the destination queue manager or the transmission queue.

This tells the queue manager the destination of the message as well as the transmission queue that it needs to be put on to get there.

3. If *DefXmitQname* is supported, in the *ObjectName* field of the MQOD structure, specify the name of the remote queue as known by the *remote* queue manager.

Note: Set the *ObjectQMGrName* field to the name of the remote queue manager (it cannot be left blank in this case).

Only local names are validated when you call MQOPEN; the last check is for the existence of the transmission queue to be used.

These methods are summarized in [Table 34 on page 208](#).

Closing objects using the MQCLOSE call

To close an object, use the MQCLOSE call.

If the object is a queue, note the following:

- You do not need to empty a temporary dynamic queue before you close it.

When you close a temporary dynamic queue, the queue is deleted, along with any messages that might still be on it. This is true even if there are uncommitted MQGET, MQPUT, or MQPUT1 calls outstanding against the queue.

- On WebSphere MQ for z/OS, if you have any MQGET requests with an MQGMO_SET_SIGNAL option outstanding for that queue, they are canceled.
- If you opened the queue using the MQOO_BROWSE option, your browse cursor is destroyed.

Closure is unrelated to sync point, so you can close queues before or after sync point.

As input to the MQCLOSE call, you must supply:

- A connection handle. Use the same connection handle used to open it, or alternatively, for CICS applications on z/OS, you can specify the constant MQHC_DEF_HCONN (which has the value zero).
- The handle of the object that you want to close. Get this from the output of the MQOPEN call.
- MQCO_NONE in the *Options* field (unless you are closing a permanent dynamic queue).
- The control option to determine whether the queue manager should delete the queue even if there are still messages on it (when closing a permanent dynamic queue).

The output from MQCLOSE is:

- A completion code
- A reason code

- The object handle, reset to the value MQHO_UNUSABLE_HOBJ

Descriptions of the parameters of the MQCLOSE call are given in [MQCLOSE](#).

Putting messages on a queue

Use this information to learn how to put messages on a queue.

Use the MQPUT call to put messages on the queue. You can use MQPUT repeatedly to put many messages on the same queue, following the initial MQOPEN call. Call MQCLOSE when you have finished putting all your messages on the queue.

If you want to put a single message on a queue and close the queue immediately afterward, you can use the MQPUT1 call. MQPUT1 performs the same functions as the following sequence of calls:

- MQOPEN
- MQPUT
- MQCLOSE

Generally however, if you have more than one message to put on the queue, it is more efficient to use the MQPUT call. This depends on the size of the message and the platform that you are working on.

Use the following links to find out more about putting messages on a queue:

- [“Putting messages on a local queue using the MQPUT call” on page 217](#)
- [“Putting messages on a remote queue” on page 221](#)
- [“Setting properties of a message” on page 222](#)
- [“Controlling context information” on page 222](#)
- [“Putting one message on a queue using the MQPUT1 call” on page 224](#)
- [“Distribution lists” on page 225](#)
- [“Some cases where the put calls fail” on page 230](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Putting messages on a local queue using the MQPUT call

Use this information to learn about putting messages on a local queue using the MQPUT call.

As input to the MQPUT call, you must supply:

- A connection handle (*Hconn*).
- A queue handle (*Hobj*).
- A description of the message that you want to put on the queue. This is in the form of a message descriptor structure (MQMD).
- Control information, in the form of a put-message options structure (MQPMO).
- The length of the data contained within the message (MQLONG).
- The message data itself.

The output from the MQPUT call is as follows:

- A reason code (MQLONG)
- A completion code (MQLONG)

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generates a unique value for the identifier of the message you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you. Reset this value before you issue another MQPUT.

There is a description of the MQPUT call in [MQPUT](#).

For more description on the information needed as input to the MQPUT call, see the following links:

- [“Specifying handles” on page 217](#)
- [“Defining messages using the MQMD structure” on page 217](#)
- [“Specifying options using the MQPMO structure” on page 218](#)
- [“The data in your message” on page 220](#)
- [“Putting messages: Using message handles” on page 221](#)

Specifying handles

For the connection handle (*Hconn*) in CICS on z/OS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or you can use the connection handle returned by the MQCONN or MQCONNEX call. For other applications, always use the connection handle returned by the MQCONN or MQCONNEX call.

Whatever environment you are working in, use the same queue handle (*Hobj*) that is returned by the MQOPEN call.

Defining messages using the MQMD structure

The message descriptor structure (MQMD) is an input/output parameter for the MQPUT and MQPUT1 calls. Use it to define the message you are putting on a queue.

If MQPRI_PRIORITY_AS_Q_DEF or MQPER_PERSISTENCE_AS_Q_DEF is specified for the message and the queue is a cluster queue, the values used are those of the queue to which the MQPUT resolves. If that queue is disabled for MQPUT, the call will fail. See [Configuring a queue manager cluster](#) for more information.

Note: Use MQPMO_NEW_MSG_ID and MQPMO_NEW_CORREL_ID before putting a new message to ensure that the *MsgId* and *CorrelId* are unique. The values in these fields are returned on a successful MQPUT.

There is an introduction to the message properties that MQMD describes in [“IBM WebSphere MQ messages”](#) on page 9, and there is a description of the structure itself in [MQMD](#).

Specifying options using the MQPMO structure

Use the MQPMO (Put Message Option) structure to pass options to the MQPUT and MQPUT1 calls.

The following sections give you help on filling in the fields of this structure. There is a description of the structure in [MQPMO](#).

The structure includes the following fields:

- *StrucId*
- *Version*
- *Options*
- *Context*
- *ResolvedQName*
- *ResolvedQMgrName*
- *RecsPresent*
- *PutMsgRecsFields*
- *ResponseRecOffset* and *ResponseRecPtr*
- *OriginalMsgHandle*
- *NewMsgHandle*
- *Action*
- *PubLevel*

The contents of these fields is as follows:

StrucId

This identifies the structure as a put-message options structure. This is a 4-character field. Always specify MQPMO_STRUC_ID.

Version

This describes the version number of the structure. The default is MQPMO_VERSION_1. If you enter MQPMO_VERSION_2, you can use distribution lists (see [“Distribution lists”](#) on page 225). If you enter MQPMO_VERSION_3, you can use message handles and message properties. If you enter MQPMO_CURRENT_VERSION, your application is set always to use the most recent level.

Options

This controls the following:

- Whether the put operation is included in a unit of work
- How much context information is associated with a message
- Where the context information is taken from
- Whether the call fails if the queue manager is in a quiescing state
- Whether grouping or segmentation is allowed
- Generation of a new message identifier and correlation identifier
- The order in which messages and segments are put on a queue
- Whether to resolve local queue names

If you leave the *Options* field set to the default value (MQPMO_NONE), the message you put has default context information associated with it.

Also, the way that the call operates with sync points is determined by the platform. The sync point control default is yes in z/OS; for other platforms, it is no.

Context

This states the name of the queue handle that you want context information to be copied from (if requested in the *Options* field).

For an introduction to message context, see “[Message context](#)” on page 37. For information about using the MQPMO structure to control the context information in a message, see “[Controlling context information](#)” on page 222.

ResolvedQName

This contains the name (after resolution of any alias name) of the queue that was opened to receive the message. This is an output field.

ResolvedQMgrName

This contains the name (after resolution of any alias name) of the queue manager that owns the queue in *ResolvedQName*. This is an output field.

The MQPMO can also accommodate fields required for distribution lists (see “[Distribution lists](#)” on page 225). If you want to use this facility, Version 2 of the MQPMO structure is used. This includes the following fields:

RecsPresent

This field contains the number of queues in the distribution list; that is, the number of Put Message Records (MQPMR) and corresponding Response Records (MQRR) present.

The value that you enter can be the same as the number of Object Records provided at MQOPEN. However, if the value is less than the number of Object Records provided on the MQOPEN call, or if you provide no Put Message Records, the values of the queues that are not defined are taken from the default values provided by the message descriptor. Also, if the value is greater than the number of Object Records provided, the excess Put Message Records are ignored.

You are recommended to do one of the following:

- If you want to receive a report or reply from each destination, enter the same value as appears in the MQOR structure and use MQPMRs containing *MsgId* fields. Either initialize these *MsgId* fields to zeros or specify MQPMO_NEW_MSG_ID.

When you have put the message to the queue, *MsgId* values that the queue manager has created become available in the MQPMRs; you can use these to identify which destination is associated with each report or reply.

- If you do not want to receive reports or replies, choose one of the following:
 1. If you want to identify destinations that fail immediately, you might still want to enter the same value in the *RecsPresent* field as appears in the MQOR structure and provide MQRRs to identify these destinations. Do not specify any MQPMRs.
 2. If you do not want to identify failed destinations, enter zero in the *RecsPresent* field and do not provide MQPMRs nor MQRRs.

Note: If you are using MQPUT1, the number of Response Record Pointers and Response Record Offsets must be zero.

For a full description of Put Message Records (MQPMR) and Response Records (MQRR), see [MQPMR](#) and [MQRR](#).

PutMsgRecFields

This indicates which fields are present in each Put Message Record (MQPMR). For a list of these fields, see “[Using the MQPMR structure](#)” on page 229.

PutMsgRecOffset and PutMsgRecPtr

Pointers (typically in C) and offsets (typically in COBOL) are used to address the Put Message Records (see “[Using the MQPMR structure](#)” on page 229 for an overview of the MQPMR structure).

Use the *PutMsgRecPtr* field to specify a pointer to the first Put Message Record, or the *PutMsgRecOffset* field to specify the offset of the first Put Message Record. This is the offset from the start of the MQPMO. Depending on the *PutMsgRecFields* field, enter a nonnull value for either *PutMsgRecOffset* or *PutMsgRecPtr*.

ResponseRecOffset and ResponseRecPtr

You also use pointers and offsets to address the Response Records (see [“Using the MQRR structure”](#) on page 228 for further information about Response Records).

Use the *ResponseRecPtr* field to specify a pointer to the first Response Record, or the *ResponseRecOffset* field to specify the offset of the first Response Record. This is the offset from the start of the MQPMO structure. Enter a nonnull value for either *ResponseRecOffset* or *ResponseRecPtr*.

Note: If you are using MQPUT1 to put messages to a distribution list, *ResponseRecPtr* must be null or zero and *ResponseRecOffset* must be zero.

Version 3 of the MQPMO structure additionally includes the following fields:

OriginalMsgHandle

The use you can make of this field depends on the value of the *Action* field. If you are putting a new message with associated message properties, set this field to the message handle you previously created and set properties on. If you are forwarding, replying to, or generating a report in response to a previously retrieved message, this field contains the message handle of that message.

NewMsgHandle

If you specify a *NewMsgHandle*, any properties associated with the handle override properties associated with the *OriginalMsgHandle*. For more information, see [Action \(MQLONG\)](#).

Action

Use this field to specify the type of put being performed. Possible values and their meanings are as follows:

MQACTP_NEW

This is a new message unrelated to any other.

MQACTP_FORWARD

This message was retrieved previously and is now being forwarded.

MQACTP_REPLY

This message is a reply to a previously retrieved message.

MQACTP_REPORT

This message is a report generated as a result of a previously retrieved message.

For more information, see [Action \(MQLONG\)](#).

PubLevel

If this message is a publication, you can set this field to determine which subscriptions receive it. Only subscriptions with a *SubLevel* less than or equal to this value will receive this publication. The default value is 9 which is the highest level and means that subscriptions with any *SubLevel* can receive this publication.

The data in your message

Give the address of the buffer that contains your data in the *Buffer* parameter of the MQPUT call. You can include anything in the data in your messages. The amount of data in the messages, however, affects the performance of the application that is processing them.

The maximum size of the data is determined by:

- The *MaxMsgLength* attribute of the queue manager
- The *MaxMsgLength* attribute of the queue on which you are putting the message
- The size of any message header added by WebSphere MQ (including the dead-letter header, MQDLH and the distribution list header, MQDLH)

The *MaxMsgLength* attribute of the queue manager holds the size of message that the queue manager can process. This has a default of 100 MB for all WebSphere MQ products at V6 or higher.

To determine the value of this attribute, use the MQINQ call on the queue manager object. For large messages, you can change this value.

The *MaxMsgLength* attribute of a queue determines the maximum size of message that you can put on the queue. If you attempt to put a message with a size larger than the value of this attribute, your MQPUT call fails. If you are putting a message on a remote queue, the maximum size of message that you can successfully put is determined by the *MaxMsgLength* attribute of the remote queue, of any intermediate transmission queues that the message is put on along the route to its destination, and of the channels used.

For an MQPUT operation, the size of the message must be smaller than or equal to the *MaxMsgLength* attribute of both the queue and the queue manager. The values of these attributes are independent, but you are recommended to set the *MaxMsgLength* of the queue to a value less than or equal to that of the queue manager.

WebSphere MQ adds header information to messages in the following circumstances:

- When you put a message on a remote queue, WebSphere MQ adds a transmission header structure (MQXQH) to the message. This structure includes the name of the destination queue and its owning queue manager.
- If WebSphere MQ cannot deliver a message to a remote queue, it attempts to put the message on the dead-letter (undelivered-message) queue. It adds an MQDLH structure to the message. This structure includes the name of the destination queue and the reason that the message was put on the dead-letter queue.
- If you want to send a message to multiple destination queues, WebSphere MQ adds an MQDH header to the message. This describes the data that is present in a message, belonging to a distribution list, on a transmission queue. Consider this when choosing an optimum value for the maximum message length.
- If the message is a segment or a message in a group, WebSphere MQ might add an MQMDE.

These structures are described in [MQDH](#) and [MQMDE](#).

If your messages are of the maximum size allowed for these queues, the addition of these headers means that the put operations fail because the messages are now too big. To reduce the possibility of the put operations failing:

- Make the size of your messages smaller than the *MaxMsgLength* attribute of the transmission and dead-letter queues. Allow at least the value of the MQ_MSG_HEADER_LENGTH constant (more for large distribution lists).
- Make sure that the *MaxMsgLength* attribute of the dead-letter queue is set to the same as the *MaxMsgLength* of the queue manager that owns the dead-letter queue.

The attributes for the queue manager and the message queuing constants are described in [Attributes for the queue manager](#).

Putting messages: Using message handles

Two message handles are available in the MQPMO structure, *OriginalMsgHandle* and *NewMsgHandle*. The relationship between these message handles is defined by the value of the MQPMO *Action* field.

For full details see [Action \(MQLONG\)](#). A message handle is not necessarily required in order to put a message. Its purpose is to associate properties with a message, so it is required only if you are using message properties.

Putting messages on a remote queue

When you want to put a message on a remote queue (that is, a queue owned by a queue manager other than the one to which your application is connected) rather than a local queue, the only extra consideration is how you specify the name of the queue when you open it. This is described in [Opening](#)

remote queues” on page 215. There is no change to how you use the MQPUT or MQPUT1 call for a local queue.

For more information on using remote and transmission queues, see [WebSphere MQ distributed-messaging techniques](#).

Setting properties of a message

Call MQSETMP for each property you want to set. When you put the message set the message handle and action fields of the MQPMO structure.

To associate properties with a message, the message must have a message handle. Create a message handle using the MQCRTMH function call. Call MQSETMP specifying this message handle for each property you want to set. A sample program, amqsstma.c, is provided to illustrate the use of MQSETMP.

If this is a new message, when you put it to a queue, using MQPUT or MQPUT1, set the OriginalMsgHandle field in the MQPMO to the value of this message handle, and set the MQPMO Action field to MQACTP_NEW (this is the default value).

If this is a message you have previously retrieved, and you are now forwarding or replying to it or sending a report in response to it, put the original message handle in the OriginalMsgHandle field of the MQPMO and the new message handle in the NewMsgHandle field. Set the Action field to MQACTP_FORWARD, MQACTP_REPLY, or MQACTP_REPORT, as appropriate.

If you have properties in an MQRFH2 header from a message you have previously retrieved, you can convert them to message handle properties using the MQBUFMH call.

If you are putting your message to a queue on a queue manager at a level earlier than WebSphere MQ Version 7.0, which cannot process message properties, you can set the PropertyControl parameter in the channel definition to specify how the properties are to be treated.

Controlling context information

When you use the MQPUT or MQPUT1 call to put a message on a queue, you can specify that the queue manager is to add some default context information to the message descriptor. Applications that have the appropriate level of authority can add extra context information. You can use the options field in the MQPMO structure to control context information.

To control context information, use the *Options* field in the MQPMO structure.

If you do not, the queue manager overwrites context information that might already be in the message descriptor with the identity and context information that it has generated for your message. This is the same as specifying the MQPMO_DEFAULT_CONTEXT option. You might want this default context information when you create a new message (for example, when processing user input from an inquiry screen).

If you want no context information associated with your message, use the MQPMO_NO_CONTEXT option. When putting a message with no context, any authority checks made by IBM WebSphere MQ are made using a blank user ID. A blank user ID cannot be assigned explicit authority to IBM WebSphere MQ resources but is treated as a member of the special group 'nobody'. For more details on the special group nobody, see [Installable services interface reference information](#).

If you want no context information associated with your message, use the MQPMO_NO_CONTEXT option.

The following sections of this topic explain the use of identity context, user context, and all context.

- [“Passing identity context” on page 223](#)
- [“Passing user context” on page 223](#)
- [“Passing all context” on page 223](#)
- [“Setting identity context” on page 223](#)
- [“Setting user context” on page 224](#)
- [“Setting all context” on page 224](#)

Passing identity context

In general, programs should pass identity context information from message to message around an application until the data reaches its final destination.

Programs should change the origin context information each time that they change the data. However, applications that want to change or set any context information must have the appropriate level of authority. The queue manager checks this authority when the applications open the queues; they must have authority to use the appropriate context options for the MQOPEN call.

If your application gets a message, processes the data from the message, then puts the changed data into another message (possibly for processing by another application), the application must pass the identity context information from the original message to the new message. You can allow the queue manager to create the origin context information.

To save the context information from the original message, use the MQOO_SAVE_ALL_CONTEXT option when you open the queue for getting the message. This is in addition to any other options you use with the MQOPEN call. Note, however, that you cannot save context information if you only browse the message.

When you create the second message:

- Open the queue using the MQOO_PASS_IDENTITY_CONTEXT option (in addition to the MQOO_OUTPUT option).
- In the *Context* field of the put-message options structure, give the handle of the queue from which you saved the context information.
- In the *Options* field of the put-message options structure, specify the MQPMO_PASS_IDENTITY_CONTEXT option.

Passing user context

You cannot choose to pass only user context. To pass user context when putting a message, specify MQPMO_PASS_ALL_CONTEXT. Any properties in the user context are passed in the same way as the origin context.

When an MQPUT or MQPUT1 takes place and the context is being passed, all properties in the user context are passed from the retrieved message to the put message. Any user context properties that the putting application has altered are put with their original values. Any user context properties that the putting application has deleted are restored in the put message. Any user context properties that the putting application has added to the message are retained.

Passing all context

If your application gets a message, and puts the message data (unchanged) into another message, the application must pass all (identity, origin, and user) context information from the original message to the new message. An example of an application that might do this is a message mover, which moves messages from one queue to another.

Follow the same procedure as for passing identity context, except that you use the MQOPEN option MQOO_PASS_ALL_CONTEXT and the put-message option MQPMO_PASS_ALL_CONTEXT.

Setting identity context

If you want to set the identity context information for a message:

- Open the queue using the MQOO_SET_IDENTITY_CONTEXT option.
- Put the message on the queue, specifying the MQPMO_SET_IDENTITY_CONTEXT option. In the message descriptor, specify whatever identity context information you require.

Note: When you set some (but not all) of the identity context fields using the MQOO_SET_IDENTITY_CONTEXT and MQPMO_SET_IDENTITY_CONTEXT options, it is important to realize that the queue manager does not set any of the other fields.

In order to modify any of the message context options, you must have the appropriate authorizations to issue the call. For example, in order to use MQOO_SET_IDENTITY_CONTEXT or MQPMO_SET_IDENTITY_CONTEXT, you must have +setid permission.

Setting user context

To set a property in the user context, set the Context field of the message property descriptor (MQPD) to MQPD_USER_CONTEXT when you make the MQSETMP call.

You do not need any special authority to set a property in the user context. User context has no MQOO_SET_* or MQPMO_SET_* context options.

Setting all context

If you want to set both the identity and the origin context information for a message:

1. Open the queue using the MQOO_SET_ALL_CONTEXT option.
2. Put the message on the queue, specifying the MQPMO_SET_ALL_CONTEXT option. In the message descriptor, specify whatever identity and origin context information you require.

Appropriate authority is needed for each type of context setting.

Related concepts

[“Message context” on page 37](#)

Message context information allows the application that retrieves the message to find out about the originator of the message.

Related reference

[“MQOPEN options relating to message context” on page 213](#)

If you want to be able to associate context information with a message when you put it on a queue, you must use one of the message context options when you open the queue.

Putting one message on a queue using the MQPUT1 call

Use the MQPUT1 call when you want to close the queue immediately after you have put a single message on it. For example, a server application is likely to use the MQPUT1 call when it is sending a reply to each of the different queues.

MQPUT1 is functionally equivalent to calling MQOPEN followed by MQPUT, followed by MQCLOSE. The only difference in the syntax for the MQPUT and MQPUT1 calls is that for MQPUT you specify an object handle, whereas for MQPUT1 you specify an object descriptor structure (MQOD) as defined in MQOPEN (see [“Identifying objects \(the MQOD structure\)” on page 208](#)). This is because you need to give information to the MQPUT1 call about the queue that it has to open, whereas when you call MQPUT, the queue must already be open.

As input to the MQPUT1 call, you must supply:

- A connection handle.
- A description of the object that you want to open. This is in the form of an object descriptor structure (MQOD).
- A description of the message that you want to put on the queue. This is in the form of a message descriptor structure (MQMD).
- Control information in the form of a put-message options structure (MQPMO).
- The length of the data contained within the message (MQLONG).
- The address of the message data.

The output from MQPUT1 is:

- A completion code
- A reason code

If the call completes successfully, it also returns your options structure and your message descriptor structure. The call modifies your options structure to show the name of the queue and the queue manager to which the message was sent. If you request that the queue manager generate a unique value for the identifier of the message that you are putting (by specifying binary zero in the *MsgId* field of the MQMD structure), the call inserts the value in the *MsgId* field before returning this structure to you.

Note: You cannot use MQPUT1 with a model queue name; however, once a model queue has been opened, you can issue an MQPUT1 to the dynamic queue.

The six input parameters for MQPUT1 are:

Hconn

This is a connection handle. For CICS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNX call. For other programs, always use the connection handle returned by the MQCONN or MQCONNX call.

ObjDesc

This is an object descriptor structure (MQOD).

In the *ObjectName* and *ObjectQMgrName* fields, give the name of the queue on which you want to put a message, and the name of the queue manager that owns this queue.

The *DynamicQName* field is ignored for the MQPUT1 call because it cannot use model queues.

Use the *AlternateUserId* field if you want to nominate an alternate user identifier that is to be used to test authority to open the queue.

MsgDesc

This is a message descriptor structure (MQMD). As with the MQPUT call, use this structure to define the message that you are putting on the queue.

PutMsgOpts

This is a put-message options structure (MQPMO). Use it as you would for the MQPUT call (see [“Specifying options using the MQPMO structure”](#) on page 218).

When the *Options* field is set to zero, the queue manager uses your own user ID when it performs tests for authority to access the queue. Also, the queue manager ignores any alternate user identifier given in the *AlternateUserId* field of the MQOD structure.

BufferLength

This is the length of your message.

Buffer

This is the buffer area that contains the text of your message.

When you use clusters, MQPUT1 operates as though MQOO_BIND_NOT_FIXED is in effect. Applications must use the resolved fields in the MQPMO structure rather than the MQOD structure to determine where the message was sent. See [Configuring a queue manager cluster](#) for more information.

There is a description of the MQPUT1 call in [MQPUT1](#).

Distribution lists

Not supported on WebSphere MQ for z/OS. Distribution lists allow you to put a message to multiple destinations in a single MQPUT or MQPUT1 call. A single MQOPEN call can open multiple queues and a single MQPUT call can then put a message to each of those queues. Some generic information from the MQI structures used for this process can be superseded by specific information relating to the individual destinations included in the distribution list.

V7.5.0.8



Attention: Distribution lists do not support the use of alias queues that point to topic objects.

From Version 7.5.0, Fix Pack 8, if an alias queue points to a topic object in a distribution list, IBM WebSphere MQ returns MQRC_ALIAS_BASE_Q_TYPE_ERROR.

When an MQOPEN call is issued, generic information is taken from the Object Descriptor (MQOD). If you specify MQOD_VERSION_2 in the *Version* field and a value greater than zero in the *RecsPresent* field, the *Hobj* can be defined as a handle of a list (of one or more queues) rather than of a queue. In this case, specific information is given through the object records (MQORs), which give details of destination (that is, *ObjectName* and *ObjectQMgrName*).

The object handle (*Hobj*) is passed to the MQPUT call, allowing you to put to a list rather than to a single queue.

When a message is put on the queues (MQPUT), generic information is taken from the Put Message Option structure (MQPMO) and the Message Descriptor (MQMD). Specific information is given in the form of Put Message Records (MQPMRs).

Response Records (MQRR) can receive a completion code and reason code specific to each destination queue.

Figure 29 on page 226 shows how distribution lists work.

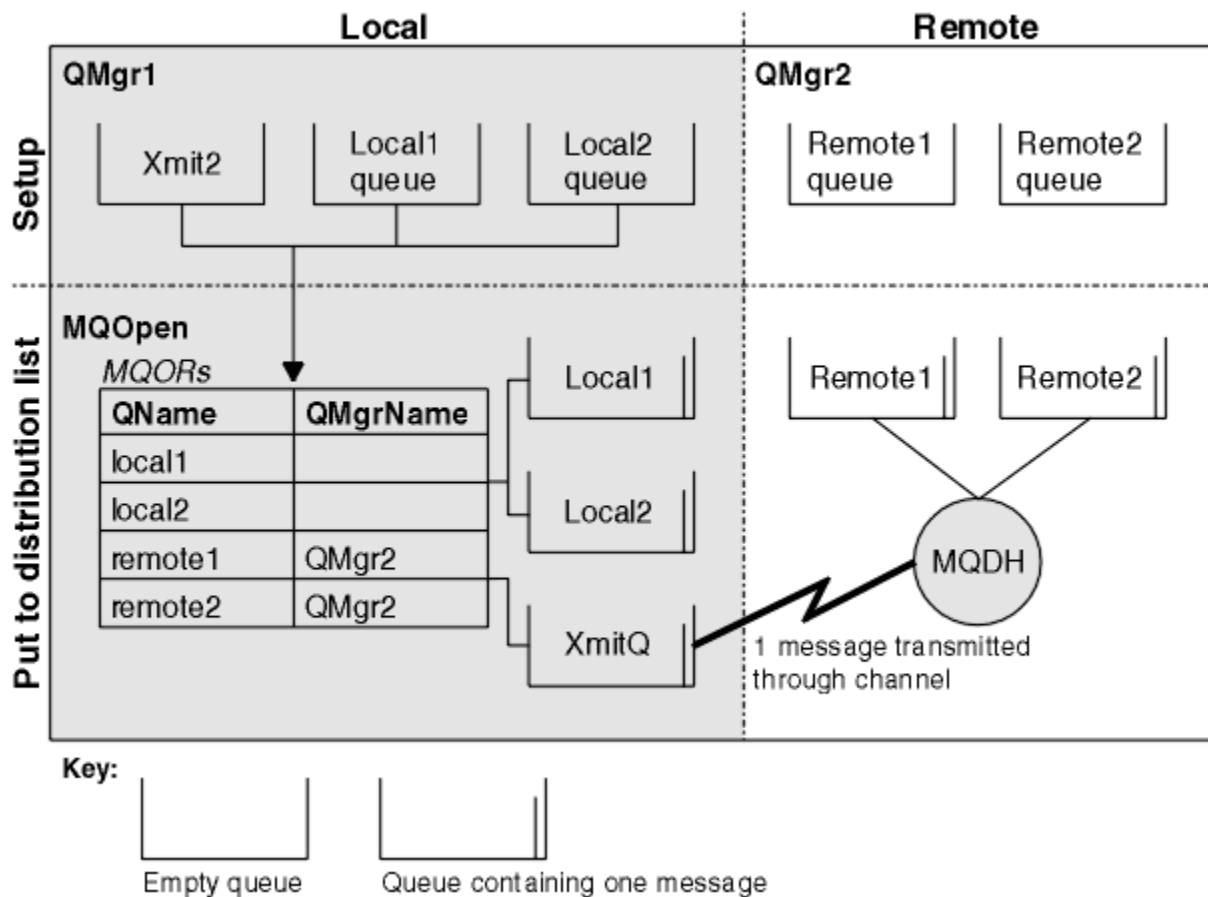


Figure 29. How distribution lists work

Opening distribution lists

Use the MQOPEN call to open a distribution list, and use the options of the call to specify what you want to do with the list.

As input to MQOPEN, you must supply:

- A connection handle (see “Putting messages on a queue” on page 216 for a description)
- Generic information in the Object Descriptor structure (MQOD)
- The name of each queue that you want to open, using the Object Record structure (MQOR)

The output from MQOPEN is:

- An object handle that represents your access to the distribution list

- A generic completion code
- A generic reason code
- Response Records (optional), containing a completion code and reason for each destination

Using the MQOD structure

Use the MQOD structure to identify the queues that you want to open.

To define a distribution list, you must specify MQOD_VERSION_2 in the *Version* field, a value greater than zero in the *RecsPresent* field, and MQOT_Q in the *ObjectType* field. See [MQOD](#) for a description of all the fields of the MQOD structure.

Using the MQOR structure

Provide an MQOR structure for each destination.

The structure contains the destination queue and queue manager names. The *ObjectName* and *ObjectQMgrName* fields in the MQOD are not used for distribution lists. There must be one or more object records. If the *ObjectQMgrName* is left blank, the local queue manager is used. See [ObjectName](#) and [ObjectQMgrName](#) for further information about these fields.

You can specify the destination queues in two ways:

- By using the offset field *ObjectRecOffset*.

In this case, the application must declare its own structure containing an MQOD structure, followed by the array of MQOR records (with as many array elements as are needed), and set *ObjectRecOffset* to the offset of the first element in the array from the start of the MQOD. Ensure that this offset is correct.

Use of built-in facilities provided by the programming language is recommended, if these are available in all the environments in which the application runs. The following code illustrates this technique for the COBOL programming language:

```
01 MY-OPEN-DATA.
02 MY-MQOD.
   COPY CMQODV.
02 MY-MQOR-TABLE OCCURS 100 TIMES.
   COPY CMQORV.
MOVE LENGTH OF MY-MQOD TO MQOD-OBJECTRECOFFSET.
```

Alternatively, use the constant MQOD_CURRENT_LENGTH if the programming language does not support the necessary built-in facilities in all the environments concerned. The following code illustrates this technique:

```
01 MY-MQ-CONSTANTS.
   COPY CMQV.
01 MY-OPEN-DATA.
02 MY-MQOD.
   COPY CMQODV.
02 MY-MQOR-TABLE OCCURS 100 TIMES.
   COPY CMQORV.
MOVE MQOD-CURRENT-LENGTH TO MQOD-OBJECTRECOFFSET.
```

However, this works correctly only if the MQOD structure and the array of MQOR records are contiguous; if the compiler inserts skip bytes between the MQOD and the MQOR array, these must be added to the value stored in *ObjectRecOffset*.

Using *ObjectRecOffset* is recommended for programming languages that do not support the pointer data type, or that implement the pointer data type in a way that is not portable to different environments (for example, the COBOL programming language).

- By using the pointer field *ObjectRecPtr*.

In this case, the application can declare the array of MQOR structures separately from the MQOD structure, and set *ObjectRecPtr* to the address of the array. The following code illustrates this technique for the C programming language:

```
MQOD MyMqod;
MQOR MyMqor[100];
MyMqod.ObjectRecPtr = MyMqor;
```

Using *ObjectRecPtr* is recommended for programming languages that support the pointer data type in a way that is portable to different environments (for example, the C programming language).

Whichever technique you choose, you must use one of *ObjectRecOffset* and *ObjectRecPtr*; the call fails with reason code MQRC_OBJECT_RECORDS_ERROR if both are zero, or both are nonzero.

Using the MQRR structure

These structures are destination-specific; each Response Record contains a *CompCode* and *Reason* field for each queue of a distribution list. You must use this structure to enable you to distinguish where any problems lie.

For example, if you receive a reason code of MQRC_MULTIPLE_REASONS and your distribution list contains five destination queues, you will not know which queues the problems apply to if you do not use this structure. However, if you have a completion code and reason code for each destination, you can locate the errors more easily.

See [MQRR](#) for further information about the MQRR structure.

Figure 30 on page 228 shows how you can open a distribution list in C.




Figure 30. Opening a distribution list in C

Figure 31 on page 228 shows how you can open a distribution list in COBOL.




Figure 31. Opening a distribution list in COBOL

Using the MQOPEN options

You can specify the following options when opening a distribution list:

- MQOO_OUTPUT
- MQOO_FAIL_IF QUIESCING (optional)
- MQOO_ALTERNATE_USER_AUTHORITY (optional)
- MQOO_*_CONTEXT (optional)

See [“Opening and closing objects” on page 206](#) for a description of these options.

Putting messages to a distribution list

To put messages to a distribution list, you can use MQPUT or MQPUT1.

As input, you must supply:

- A connection handle (see [“Putting messages on a queue”](#) on page 216 for a description).
- An object handle. If a distribution list is opened using MQOPEN, the *Hobj* allows you only to put to the list.
- A message descriptor structure (MQMD). See [MQMD](#) for a description of this structure.
- Control information in the form of a put-message option structure (MQPMO). See [“Specifying options using the MQPMO structure”](#) on page 218 for information about completing the fields of the MQPMO structure.
- Control information in the form of Put Message Records (MQPMR).
- The length of the data contained within the message (MQLONG).
- The message data itself.

The output is:

- A completion code
- A reason code
- Response Records (optional)

Using the MQPMR structure

This structure is optional and gives destination-specific information for some fields that you might want to identify differently from those already identified in the MQMD.

For a description of these fields, see [MQPMR](#).

The content of each record depends on the information given in the *PutMsgRecFields* field of the MQPMO. For example, in the sample program AMQSPTL0.C (see [“The Distribution List sample program”](#) on page 122 for a description) showing the use of distribution lists, the sample chooses to provide values for *MsgId* and *CorrelId* in the MQPMR. This section of the sample program looks like this:

```
typedef struct
{
  MQBYTE24 MsgId;
  MQBYTE24 CorrelId;
} PutMsgRec;
...
/*****
MQLONG PutMsgRecFields=MQPMRF_MSG_ID | MQPMRF_CORREL_ID;
```

This implies that *MsgId* and *CorrelId* are provided for each destination of a distribution list. The Put Message Records are provided as an array.

Figure 32 on page 229 shows how you can put a message to a distribution list in C.

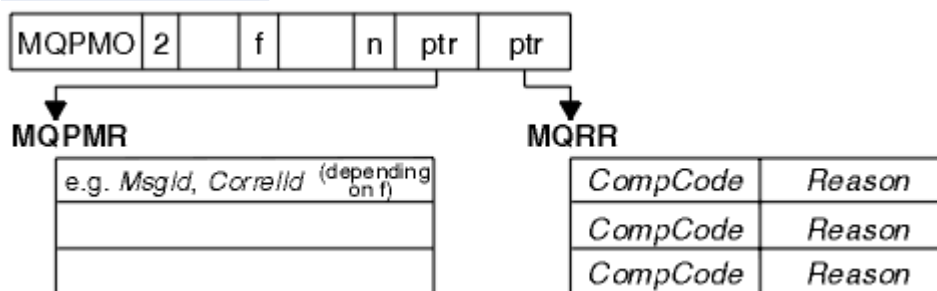


Figure 32. Putting a message to a distribution list in C

Figure 33 on page 230 shows how you can put a message to a distribution list in COBOL.




Figure 33. Putting a message to a distribution list in COBOL

Using MQPUT1

If you are using MQPUT1, consider the following points:

1. The values of the *ResponseRecOffset* and *ResponseRecPtr* fields must be null or zero.
2. The Response Records, if required, must be addressed from the MQOD.

Some cases where the put calls fail

If certain attributes of a queue are changed using the FORCE option on a command during the interval between you issuing an MQOPEN and an MQPUT call, the MQPUT call fails and returns the MQRC_OBJECT_CHANGED reason code.

The queue manager marks the object handle as being no longer valid. This also happens if the changes are made while an MQPUT1 call is being processed, or if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in [MQOPEN](#). If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to put a message again.

If put operations are inhibited for a queue on which you are attempting to put messages (or any queue to which the queue name resolves), the MQPUT or MQPUT1 call fails and returns the MQRC_PUT_INHIBITED reason code. You might be able to put a message successfully if you attempt the call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

Furthermore, if the queue that you are trying to put your message on is full, the MQPUT or MQPUT1 call fails and returns MQRC_Q_FULL.

If a dynamic queue (either temporary or permanent) has been deleted, MQPUT calls using a previously-acquired object handle fail and return the MQRC_Q_DELETED reason code. In this situation, it is good practice to close the object handle as it is no longer of any use to you.

In the case of distribution lists, multiple completion codes and reason codes can occur in a single request. These cannot be handled using only the *CompCode* and *Reason* output fields on MQOPEN and MQPUT.

When you use distribution lists to put messages to multiple destinations, the Response Records contain the specific *CompCode* and *Reason* for each destination. If you receive a completion code of MQCC_FAILED, no message is put on any destination queue successfully. If the completion code is MQCC_WARNING, the message is successfully put on one or more of the destination queues. If you receive a return code of MQRC_MULTIPLE_REASONS, the reason codes are not all the same for every destination. Therefore, it is recommended to use the MQRR structure so that you can determine which queue or queues caused an error and the reasons for each.

Getting messages from a queue

Use this information to learn about getting messages from a queue.

You can get messages from a queue in two ways:

1. You can remove a message from the queue so that other programs can no longer see it.
2. You can copy a message, leaving the original message on the queue. This is known as *browsing*. You can remove the message once you have browsed it.

In both cases, you use the MQGET call, but first your application must be connected to the queue manager, and you must use the MQOPEN call to open the queue (for input, browse, or both). These

operations are described in [“Connecting to and disconnecting from a queue manager” on page 198](#) and [“Opening and closing objects” on page 206](#).

When you have opened the queue, you can use the MQGET call repeatedly to browse or remove messages on the same queue. Call MQCLOSE when you have finished getting all the messages that you want from the queue.

Use the following links to find out more about getting messages from a queue:

- [“Getting messages from a queue using the MQGET call” on page 231](#)
- [“The order in which messages are retrieved from a queue” on page 235](#)
- [“Getting a particular message” on page 246](#)
- [“Improving performance of non-persistent messages” on page 247](#)
- [“Handling messages greater than 4 MB long” on page 251](#)
- [“Waiting for messages” on page 256](#)
-
- [“Skipping backout” on page 257](#)
- [“Application data conversion” on page 259](#)
- [“Browsing messages on a queue” on page 260](#)
- [“Some cases where the MQGET call fails” on page 266](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Getting messages from a queue using the MQGET call

The MQGET call gets a message from an open local queue. It cannot get a message from a queue on another system.

As input to the MQGET call, you must supply:

- A connection handle.
- A queue handle.
- A description of the message that you want to get from the queue. This is in the form of a message descriptor (MQMD) structure.
- Control information in the form of a Get Message Options (MQGMO) structure.

- The size of the buffer that you have assigned to hold the message (MQLONG).
- The address of the storage in which to put the message.

The output from MQGET is:

- A reason code
- A completion code
- The message in the buffer area that you specified, if the call completes successfully
- Your options structure, modified to show the name of the queue from which the message was retrieved
- Your message descriptor structure, with the contents of the fields modified to describe the message that was retrieved
- The length of the message (MQLONG)

There is a description of the MQGET call in [MQGET](#).

The following sections describe the information you must supply as input to the MQGET call.

- [“Specifying connection handles” on page 232](#)
- [“Describing messages using the MQMD structure and the MQGET call” on page 232](#)
- [“Specifying MQGET options using the MQGMO structure” on page 232](#)
- [“Specifying the size of the buffer area” on page 234](#)

Specifying connection handles

For CICS on z/OS applications, you can specify the constant MQHC_DEF_HCONN (which has the value zero), or use the connection handle returned by the MQCONN or MQCONNEX call. For other applications, always use the connection handle returned by the MQCONN or MQCONNEX call.

Use the queue handle (*Hobj*) that is returned when you call MQOPEN.

Describing messages using the MQMD structure and the MQGET call

To identify the message that you want to get from a queue, use the message descriptor structure (MQMD).

This is an input/output parameter for the MQGET call. There is an introduction to the message properties that MQMD describes in [“IBM WebSphere MQ messages” on page 9](#), and there is a description of the structure itself in [MQMD](#).

If you know which message you want to get from the queue, see [“Getting a particular message” on page 246](#).

If you do not specify a particular message, MQGET retrieves the *first* message in the queue. [“The order in which messages are retrieved from a queue” on page 235](#) describes how the priority of a message, the *MsgDeliverySequence* attribute of the queue, and the MQGMO_LOGICAL_ORDER option determine the order of the messages in the queue.

Note: If you want to use MQGET more than once (for example, to step through the messages in the queue), you must set the *MsgId* and *CorrelId* fields of this structure to null after each call. This clears these fields of the identifiers of the message that was retrieved.

However, if you want to group your messages, the *GroupId* must be the same for messages in the same group, so that the call looks for a message having the same identifiers as the previous message in order to make up the whole group.

Specifying MQGET options using the MQGMO structure

The MQGMO structure is an input/output variable for passing options to the MQGET call. The following sections help you to complete some of the fields of this structure.

There is a description of the MQGMO structure in [MQGMO](#).

StrucId

StrucId is a 4-character field used to identify the structure as a get-message options structure. Always specify MQGMO_STRUC_ID.

Version

Version describes the version number of the structure. MQGMO_VERSION_1 is the default. If you want to use the Version 2 fields or retrieve messages in logical order, specify MQGMO_VERSION_2. If you want to use the Version 3 fields or retrieve messages in logical order, specify MQGMO_VERSION_3. MQGMO_CURRENT_VERSION sets your application to use the most recent level.

Options

Within your code, you can select the options in any order; each option is represented by a bit in the *Options* field.

The *Options* field controls:

- Whether the MQGET call waits for a message to arrive on the queue before it completes (see [“Waiting for messages” on page 256](#))
- Whether the get operation is included in a unit of work.
- Whether a nonpersistent message is retrieved outside sync point, allowing fast messaging
- On WebSphere MQ for z/OS, whether the message retrieved is marked as skipping backout (see [“Skipping backout” on page 257](#))
- Whether the message is removed from the queue, or merely browsed
- Whether to select a message by using a browse cursor or by other selection criteria
- Whether the call succeeds even if the message is longer than your buffer
- On WebSphere MQ for z/OS, whether to allow the call to complete. This option also sets a signal to indicate that you want to be notified when a message arrives
- Whether the call fails if the queue manager is in a quiescing state
- On WebSphere MQ for z/OS, whether the call fails if the connection is in a quiescing state
- Whether application message data conversion is required (see [“Application data conversion” on page 259](#))
- The order in which messages and (except for WebSphere MQ for z/OS) segments are retrieved from a queue
- Except on WebSphere MQ for z/OS, whether complete, logical messages only are retrievable
- Whether messages in a group can be retrieved only when *all* messages in the group are available
- Except on WebSphere MQ for z/OS, whether segments in a logical message can be retrieved only when *all* segments in the logical message are available

If you leave the *Options* field set to the default value (MQGMO_NO_WAIT), the MQGET call operates this way:

- If there is no message matching your selection criteria on the queue, the call does not wait for a message to arrive, but completes immediately. Also, in WebSphere MQ for z/OS, the call does not set a signal requesting notification when such a message arrives.
- The way that the call operates with sync points is determined by the platform:

Platform	Under sync point control
IBM i	No
UNIX and Linux systems	No
z/OS	Yes
Windows systems	No

- On WebSphere MQ for z/OS, the message retrieved is not marked as skipping backout.

- The selected message is removed from the queue (not browsed).
- No application message data conversion is required.
- The call fails if the message is longer than your buffer.

WaitInterval

The *WaitInterval* field specifies the maximum time (in milliseconds) that the MQGET call waits for a message to arrive on the queue when you use the MQGMO_WAIT option. If no message arrives within the time specified in *WaitInterval*, the call completes and returns a reason code showing that there was no message that matched your selection criteria on the queue.

On WebSphere MQ for z/OS, if you use the MQGMO_SET_SIGNAL option, the *WaitInterval* field specifies the time for which the signal is set.

For more information about these options, see [“Waiting for messages” on page 256](#).

Signal1

Signal1 is supported on WebSphere MQ for z/OS and MQSeries for HP Integrity NonStop Server only.

If you use the MQGMO_SET_SIGNAL option to request that your application is notified when a suitable message arrives, you specify the type of signal in the *Signal1* field. In WebSphere MQ on all other platforms, the *Signal1* field is reserved and its value is not significant.

Signal2

The *Signal2* field is reserved on all platforms and its value is not significant.

ResolvedQName

ResolvedQName is an output field in which the queue manager returns the name of the queue (after resolution of any alias) from which the message was retrieved.

MatchOptions

MatchOptions controls the selection criteria for MQGET.

GroupStatus

GroupStatus indicates whether the message that you have retrieved is in a group.

SegmentStatus

SegmentStatus indicates whether the item that you have retrieved is a segment of a logical message.

Segmentation

Segmentation indicates whether segmentation is allowed for the message retrieved.

MsgToken

MsgToken uniquely identifies a message.

ReturnedLength

ReturnedLength is an output field in which the queue manager returns the length of message data returned (in bytes).

MsgHandle

The handle to a message that is to be populated with the properties of the message being retrieved from the queue. The handle has previously been created by an MQCRTMH call. Any properties already associated with the handle are cleared before retrieving a message.

Specifying the size of the buffer area

In the *BufferLength* parameter of the MQGET call, specify the size of the buffer area to hold the message data that you retrieve. You decide how large this should be in three ways:

1. You might already know what length of messages to expect from this program. If so, specify a buffer of this size.

However, you can use the MQGMO_ACCEPT_TRUNCATED_MSG option in the MQGMO structure if you want the MQGET call to complete even if the message is too large for the buffer. In this case:

- The buffer is filled with as much of the message as it can hold
- The call returns a warning completion code
- The message is removed from the queue (discarding the remainder of the message), or the browse cursor is advanced (if you are browsing the queue)
- The real length of the message is returned in *DataLength*

Without this option, the call still completes with a warning, but it does not remove the message from the queue (or advance the browse cursor).

2. Estimate a size for the buffer (or even specify a size of zero bytes) and *do not* use the MQGMO_ACCEPT_TRUNCATED_MSG option. If the MQGET call fails (for example, because the buffer is too small), the length of the message is returned in the *DataLength* parameter of the call. (The buffer is still filled with as much of the message as it can hold, but the processing of the call is not completed.) Store the *MsgId* of this message, then repeat the MQGET call, specifying a buffer area of the correct size, and the *MsgId* that you noted from the first call.

If your program is serving a queue that is also being served by other programs, one of those other programs might remove the message that you want before your program can issue another MQGET call. Your program could waste time searching for a message that no longer exists. To avoid this, first browse the queue until you find the message that you want, specifying a *BufferLength* of zero and using the MQGMO_ACCEPT_TRUNCATED_MSG option. This positions the browse cursor under the message that you want. You can then retrieve the message by calling MQGET again, specifying the MQGMO_MSG_UNDER_CURSOR option. If another program removes the message between your browse and removal calls, your second MQGET fails immediately (without searching the whole queue), because there is no message under your browse cursor.

3. The *MaxMsgLength queue* attribute determines the maximum length of messages accepted for that queue; the *MaxMsgLength queue manager* attribute determines the maximum length of messages accepted for that queue manager. If you do not know what length of message to expect, you can inquire about the *MaxMsgLength* attribute (using the MQINQ call), then specify a buffer of this size.

Try to make the buffer size as close as possible to the actual message size to avoid reduced performance.

For further information about the *MaxMsgLength* attribute, see [“Increasing the maximum message length” on page 251](#).

The order in which messages are retrieved from a queue

You can control the order in which you retrieve messages from a queue. This section looks at the options.

Priority

A program can assign a priority to a message when it puts the message on a queue (see [“Message priorities” on page 17](#)). Messages of equal priority are stored in a queue in order of arrival, not the order in which they are committed.

The queue manager maintains queues either in strict FIFO (first in, first out) sequence, or in FIFO within priority sequence. This depends on the setting of the *MsgDeliverySequence* attribute of the queue. When a message arrives on a queue, it is inserted immediately following the last message that has the same priority.

Programs can either get the first message from a queue, or they can get a particular message from a queue, ignoring the priority of those messages. For example, a program might want to process the reply to a particular message that it sent earlier. For more information, see [“Getting a particular message” on page 246](#).

If an application puts a sequence of messages on a queue, another application can retrieve those messages in the same order that they were put, provided:

- The messages all have the same priority
- The messages were all put within the same unit of work, or all put outside a unit of work

- The queue is local to the putting application

If these conditions are not met, and the applications depend on the messages being retrieved in a certain order, the applications must either include sequencing information in the message data, or establish a means of acknowledging receipt of a message before the next one is sent.

Logical and physical ordering

Messages on queues can occur (within each priority level) in *physical* or *logical* order.

Physical order is the order in which messages arrive on a queue. Logical order is when all of the messages and segments within a group are in their logical sequence, next to each other, in the position determined by the physical position of the first item belonging to the group.

For a description of groups, messages, and segments, see [“Message groups” on page 34](#). These physical and logical orders can differ because:

- Groups can arrive at a destination at similar times from different applications, therefore losing any distinct physical order.
- Even within a single group, messages can get out of order because of rerouting or delay of some of the messages in the group.

For example, the logical order might look like [Figure 34 on page 236](#):




Figure 34. Logical order on a queue

These messages would occur in the following logical order on a queue:

1. Message A (not in a group)
2. Logical message 1 of group Y

3. Logical message 2 of group Y
4. Segment 1 of (last) logical message 3 of group Y
5. (Last) segment 2 of (last) logical message 3 of group Y
6. Logical message 1 of group Z
7. (Last) logical message 2 of group Z
8. Message B (not in a group)

The physical order, however, might be entirely different. The physical position of the *first* item within each group determines the logical position of the whole group. For example, if groups Y and Z arrived at similar times, and message 2 of group Z overtook message 1 of the same group, the physical order would look like Figure Figure 35 on page 237:

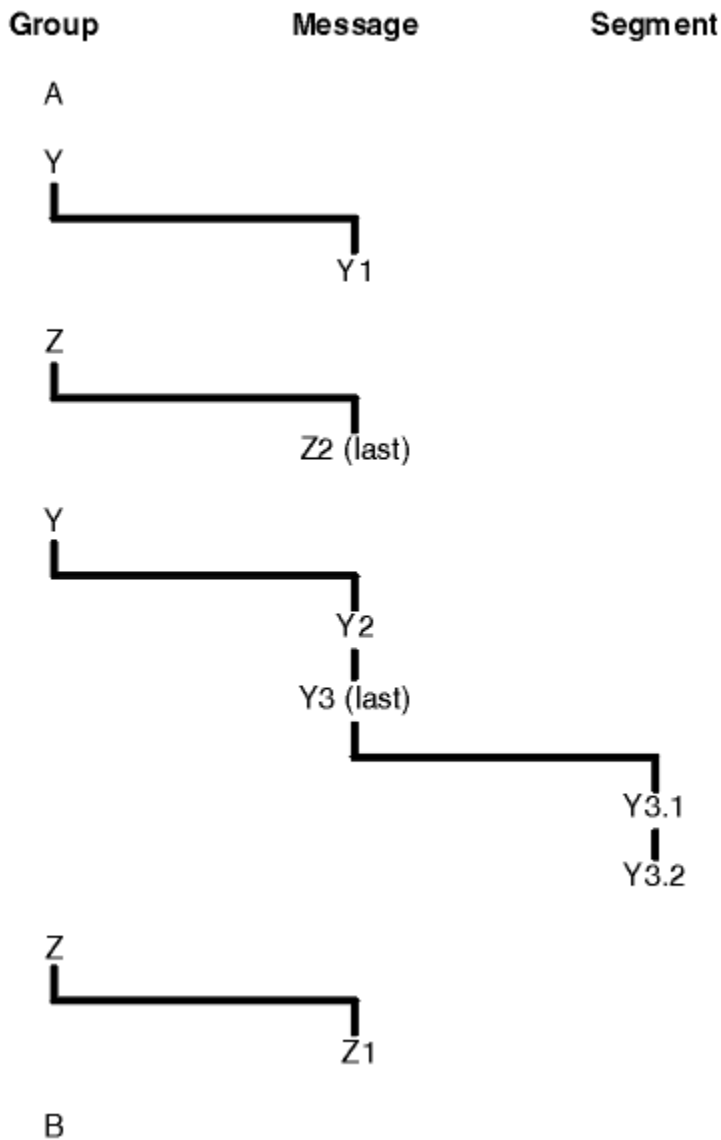


Figure 35. Physical order on a queue

These messages occur in the following physical order on the queue:

1. Message A (not in a group)
2. Logical message 1 of group Y
3. Logical message 2 of group Z
4. Logical message 2 of group Y

5. Segment 1 of (last) logical message 3 of group Y
6. (Last) segment 2 of (last) logical message 3 of group Y
7. Logical message 1 of group Z
8. Message B (not in a group)

Note: On IBM WebSphere MQ for z/OS, the physical order of messages on the queue is not guaranteed if the queue is indexed by GROUPID.

When getting messages, you can specify MQGMO_LOGICAL_ORDER to retrieve messages in logical order rather than physical order.

If you issue an MQGET call with MQGMO_BROWSE_FIRST and MQGMO_LOGICAL_ORDER, subsequent MQGET calls with MQGMO_BROWSE_NEXT must also specify MQGMO_LOGICAL_ORDER. Conversely, if the MQGET with MQGMO_BROWSE_FIRST does not specify MQGMO_LOGICAL_ORDER, neither must the following MQGETs with MQGMO_BROWSE_NEXT.

The group and segment information that the queue manager retains for MQGET calls that browse messages on the queue is separate from the group and segment information that the queue manager retains for MQGET calls that remove messages from the queue. When you specify MQGMO_BROWSE_FIRST, the queue manager ignores the group and segment information for browsing, and scans the queue as though there were no current group and no current logical message.

Note: Do not use an MQGET call to browse *beyond the end* of a message group (or logical message not in a group) without specifying MQGMO_LOGICAL_ORDER. For example, if the last message in the group *precedes* the first message in the group on the queue, using MQGMO_BROWSE_NEXT to browse beyond the end of the group, specifying MQMO_MATCH_MSG_SEQ_NUMBER with *MsgSeqNumber* set to 1 (to find the first message of the next group) returns again the first message in the group already browsed. This could happen immediately, or a number of MQGET calls later (if there are intervening groups).

Avoid the possibility of an infinite loop by opening the queue *twice* for browse:

- Use the first handle to browse only the first message in each group.
- Use the second handle to browse only the messages within a specific group.
- Use the MQMO_* options to move the second browse cursor to the position of the first browse cursor, before browsing the messages in the group.
- Do not use the MQGMO_BROWSE_NEXT browse beyond the end of a group.

For further information about this, see [MQGET](#), [MQMD](#), and [Rules for validating MQI options](#).

For most applications you will probably choose either logical or physical ordering when browsing. However, if you want to switch between these modes, remember that when you first issue a browse with MQGMO_LOGICAL_ORDER, your position within the logical sequence is established.

If the first item within the group is not present at this time, the group that you are in is not considered to be part of the logical sequence.

Once the browse cursor is within a group, it can continue within the same group, even if the first message is removed. Initially though, you can never move into a group using MQGMO_LOGICAL_ORDER where the first item is not present.

MQPMO_LOGICAL_ORDER

The MQPMO option tells the queue manager how the application puts messages in groups and segments of logical messages. It can be specified only on the MQPUT call; it is not valid on the MQPUT1 call.

If MQPMO_LOGICAL_ORDER is specified, it indicates that the application uses successive MQPUT calls to:

1. Put the segments in each logical message in the order of increasing segment offset, starting from 0, with no gaps.
2. Put all the segments in one logical message before putting the segments in the next logical message.

3. Put the logical messages in each message group in the order of increasing message sequence number, starting from 1, with no gaps. IBM WebSphere MQ increments the message sequence number automatically.
4. Put all the logical messages in one message group before putting logical messages in the next message group.

Because the application has told the queue manager how it puts messages in groups and segments of logical messages, the application does not have to maintain and update the group and segment information about each MQPUT call, because the queue manager maintains and updates this information. Specifically, it means that the application does not need to set the *GroupId*, *MsgSeqNumber*, and *Offset* fields in MQMD, because the queue manager sets these fields to the appropriate values. The application must only set the *MsgFlags* field in MQMD, to indicate when messages belong to groups or are segments of logical messages, and to indicate the last message in a group or last segment of a logical message.

After a message group or logical message has been started, subsequent MQPUT calls must specify the appropriate MQMF_* flags in *MsgFlags* in MQMD. If the application tries to put a message that is not in a group when there is an unterminated message group, or put a message that is not a segment when there is an unterminated logical message, the call fails with reason code MQRC_INCOMPLETE_GROUP or MQRC_INCOMPLETE_MSG, as appropriate. However, the queue manager retains the information about the current message group or current logical message, and the application can terminate them by sending a message (possibly with no application message data) specifying MQMF_LAST_MSG_IN_GROUP or MQMF_LAST_SEGMENT as appropriate, before reissuing the MQPUT call to put the message that is not in the group or not a segment.

Figure 35 on page 237 shows the combinations of options and flags that are valid, and the values of the *GroupId*, *MsgSeqNumber*, and *Offset* fields that the queue manager uses in each case. Combinations of options and flags that are not shown in the table are not valid. The columns in the table have the following meanings; Either means Yes or No:

LOG ORD

Whether the MQPMO_LOGICAL_ORDER option is specified on the call.

MIG

Whether the MQMF_MSG_IN_GROUP or MQMF_LAST_MSG_IN_GROUP option is specified on the call.

SEG

Whether the MQMF_SEGMENT or MQMF_LAST_SEGMENT option is specified on the call.

SEG OK

Whether the MQMF_SEGMENTATION_ALLOWED option is specified on the call.

Cur grp

Whether a current message group exists before the call.

Cur log msg

Whether a current logical message exists before the call.

Other columns

Show the values that the queue manager uses. Previous denotes the value used for the field in the previous message for the queue handle.

Table 36. MQPUT options relating to messages in groups and segments of logical messages

Options you specify				Group and log-msg status before call		Values the queue manager uses		
LOG ORD	MIG	SEG	SEG OK	Cur grp	Cur log msg	GroupId	MsgSeqNumber	Offset
Yes	No	No	No	No	No	MQGI_NONE	1	0
Yes	No	No	Yes	No	No	New group id	1	0
Yes	No	Yes	Either	No	No	New group id	1	0
Yes	No	Yes	Either	No	Yes	Previous group id	1	Previous offset + previous segment length
Yes	Yes	Either	Either	No	No	New group id	1	0
Yes	Yes	Either	Either	Yes	No	Previous group id	Previous sequence number + 1	0
Yes	Yes	Yes	Either	Yes	Yes	Previous group id	Previous sequence number	Previous offset + previous segment length
No	No	No	No	Either	Either	MQGI_NONE	1	0
No	No	No	Yes	Either	Either	New group id if MQGI_NONE, else value in field	1	0
No	No	Yes	Either	Either	Either	New group id if MQGI_NONE, else value in field	1	Value in field
No	Yes	No	Either	Either	Either	New group id if MQGI_NONE, else value in field	Value in field	0
No	Yes	Yes	Either	Either	Either	New group id if MQGI_NONE, else value in field	Value in field	Value in field

Note:

- MQPMO_LOGICAL_ORDER is not valid on the MQPUT1 call.
- For the *MsgId* field, the queue manager generates a new message identifier if MQPMO_NEW_MSG_ID or MQMI_NONE is specified, and uses the value in the field otherwise.
- For the *CorrelId* field, the queue manager generates a new correlation identifier if MQPMO_NEW_CORREL_ID is specified, and uses the value in the field otherwise.

When you specify MQPMO_LOGICAL_ORDER, the queue manager requires that all messages in a group and segments in a logical message are put with the same value in the *Persistence* field in MQMD, that is, all must be persistent, or all must be nonpersistent. If this condition is not satisfied, the MQPUT call fails with reason code MQRC_INCONSISTENT_PERSISTENCE.

The MQPMO_LOGICAL_ORDER option affects units of work as follows:

- If the first physical message in a group or logical message is put within a unit of work, all the other physical messages in the group or logical message must be put within a unit of work, if the same queue handle is used. However, they do not need to be put within the same unit of work, allowing a message group or logical message that consists of many physical messages to be split across two or more consecutive units of work for the queue handle.
- If the first physical message in a group or logical message is not put within a unit of work, none of the other physical messages in the group or logical message can be put within a unit of work, if the same queue handle is used.

If these conditions are not satisfied, the MQPUT call fails with reason code MQRC_INCONSISTENT_UOW.

When MQPMO_LOGICAL_ORDER is specified, the MQMD supplied on the MQPUT call must not be less than MQMD_VERSION_2. If this condition is not satisfied, the call fails with reason code MQRC_WRONG_MD_VERSION.

If MQPMO_LOGICAL_ORDER is not specified, messages in groups and segments of logical messages can be put in any order, and it is not necessary to put complete message groups or complete logical messages. It is the responsibility of the application to ensure that the *GroupId*, *MsgSeqNumber*, *Offset*, and *MsgFlags* fields have appropriate values.

Use this technique to restart a message group or logical message in the middle, after a system failure has occurred. When the system restarts, the application can set the *GroupId*, *MsgSeqNumber*, *Offset*, *MsgFlags*, and *Persistence* fields to the appropriate values, and then issue the MQPUT call with MQPMO_SYNCPOINT or MQPMO_NO_SYNCPOINT set as required, but without specifying MQPMO_LOGICAL_ORDER. If this call is successful, the queue manager retains the group and segment information, and subsequent MQPUT calls using that queue handle can specify MQPMO_LOGICAL_ORDER as normal.

The group and segment information that the queue manager retains for the MQPUT call is separate from the group and segment information that it retains for the MQGET call.

For any given queue handle, the application can mix MQPUT calls which specify MQPMO_LOGICAL_ORDER with MQPUT calls that do not, but note the following points:

- If MQPMO_LOGICAL_ORDER is not specified, each successful MQPUT call causes the queue manager to set the group and segment information for the queue handle to the values specified by the application, replacing the existing group and segment information retained by the queue manager for the queue handle.
- If MQPMO_LOGICAL_ORDER is not specified, the call does not fail if there is a current message group or logical message; the call might succeed with an MQCC_WARNING completion code. [Table 37 on page 241](#) shows the various cases that can arise. In these cases, if the completion code is not MQCC_OK, the reason code is one of the following (as appropriate):
 - MQRC_INCOMPLETE_GROUP
 - MQRC_INCOMPLETE_MSG
 - MQRC_INCONSISTENT_PERSISTENCE
 - MQRC_INCONSISTENT_UOW

Note: The queue manager does not check the group and segment information for the MQPUT1 call.

Table 37. Outcome when MQPUT or MQCLOSE call is not consistent with group and segment information		
Current call is	Previous call was MQPUT with MQPMO_LOGICAL_ORDER	Previous call was MQPUT without MQPMO_LOGICAL_ORDER
MQPUT with MQPMO_LOGICAL_ORDER	MQCC_FAILED	MQCC_FAILED
MQPUT without MQPMO_LOGICAL_ORDER	MQCC_WARNING	MQCC_OK

Table 37. Outcome when MQPUT or MQCLOSE call is not consistent with group and segment information (continued)

Current call is	Previous call was MQPUT with MQPMO_LOGICAL_ORDER	Previous call was MQPUT without MQPMO_LOGICAL_ORDER
MQCLOSE with an unterminated group or logical message	MQCC_WARNING	MQCC_OK

For applications that put messages and segments in logical order, specify MQPMO_LOGICAL_ORDER, as it is the simplest option to use. This option relieves the application of the need to manage the group and segment information, because the queue manager manages that information. However, specialized applications might need more control than that provided by the MQPMO_LOGICAL_ORDER option, which can be achieved by not specifying that option; if you do so, you must ensure that the *GroupId*, *MsgSeqNumber*, *Offset*, and *MsgFlags* fields in MQMD are set correctly, before each MQPUT or MQPUT1 call.

For example, an application that wants to forward physical messages that it receives, without regard for whether those messages are in groups or segments of logical messages, must not specify MQPMO_LOGICAL_ORDER, for two reasons:

- If the messages are retrieved and put in order, specifying MQPMO_LOGICAL_ORDER assigns a new group identifier to the messages, which might make it difficult or impossible for the originator of the messages to correlate any reply or report messages that result from the message group.
- In a complex network with multiple paths between sending and receiving queue managers, the physical messages might arrive out of order. By not specifying MQPMO_LOGICAL_ORDER and MQGMO_LOGICAL_ORDER on the MQGET call, the forwarding application can retrieve and forward each physical message as soon as it arrives, without waiting for the next one in logical order to arrive.

Applications that generate report messages for messages in groups or segments of logical messages must also not specify MQPMO_LOGICAL_ORDER when putting the report message.

MQPMO_LOGICAL_ORDER can be specified with any of the other MQPMO_* options.

Putting Logically Ordered Groups to a Clustered Queue (MQOO_BIND_ON_GROUP)

The MQOO_BIND_ON_OPEN option ensures that all messages from this application, and therefore all groups, are routed to a single instance. This has the drawback that the application traffic is not load balanced across multiple instances of a cluster queue. In order to enable workload balancing while keeping groups of messages intact, you must set the following options:

- The MQPUT call must specify MQPMO_LOGICAL_ORDER
- The MQOPEN call must specify one of the following two options:
 - MQOO_BIND_ON_GROUP
 - MQOO_BIND_AS_Q_DEF, and the queue definition must specify DEFBIND(GROUP)

Workload balancing is then driven *between groups* of messages without requiring an MQCLOSE and MQOPEN of the queue. *Between groups* means that MQMF_MSG_IN_GROUP is set in the MQMD(v2) or MQMDE, and there is no partially complete group in progress. When a group is in progress, the resolved queue manager and queue name in the object handle are reused.

If the previous message was MQPMO_LOGICAL_ORDER and/or MQMF_MSG_IN_GROUP was set but the current message is not part of the group, then the PUT call fails with MQRC_INCOMPLETE_GROUP.

If an individual MQPUT does not specify MQPMO_LOGICAL_ORDER, and no current group is active, then workload balancing is driven for that message (as if the MQOPEN call has specified MQOO_BIND_NOT_FIXED).

No reallocation takes place for messages bound to a destination using MQOO_BIND_ON_GROUP. For more information on reallocation, see [“Message groups” on page 34](#).

Grouping logical messages

There are two main reasons for using logical messages in a group:

- You might need to process the messages in a particular order.
- You might need to process each message in a group in a related way.

In either case, retrieve the entire group with the same getting application instance.

For example, assume that the group consists of four logical messages. The putting application looks like this:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP

MQCMIT
```

The getting application specifies the MQGMO_ALL_MSGS_AVAILABLE option for the first message in the group. This ensures that processing does not start until all the messages within the group have arrived. The MQGMO_ALL_MSGS_AVAILABLE option is ignored for subsequent messages within the group.

When the first logical message of the group is retrieved, you can use MQGMO_LOGICAL_ORDER to ensure that the remaining logical messages of the group are retrieved in order.

So, the getting application looks like this:

```
/* Wait for the first message in a group, or a message not in a group */
GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
              | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_LOGICAL_ORDER
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Process each remaining message in the group */
  ...
MQCMIT
```

For further examples of grouping messages, see [“Application segmentation of logical messages” on page 254](#) and [“Putting and getting a group that spans units of work” on page 243](#).

For information on allowing an application to request that a group of messages are all allocated to the same destination instance for cluster queues, see [DefBind](#).

Putting and getting a group that spans units of work

In the previous case, messages or segments cannot start to leave the node (if its destination is remote) or start to be retrieved until the whole group has been put and the unit of work is committed. This might not be what you want if it takes a long time to put the whole group, or if queue space is limited on the node. To overcome this, put the group in several units of work.

If the group is put within multiple units of work, it is possible for some of the group to commit even when the putting application fails. The application must therefore save status information, committed with each unit of work, which it can use after a restart to resume an incomplete group. The simplest place to record this information is in a STATUS queue. If a complete group has been successfully put, the STATUS queue is empty.

If segmentation is involved, the logic is similar. In this case, the StatusInfo must include the *Offset* .

Here is an example of putting the group in several units of work:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

/* First UOW */
```

```

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Next and subsequent UOWs */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
StatusInfo = GroupId,MsgSeqNumber from MQMD
MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT

/* Last UOW */
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP
MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
MQCMIT

```

If all the units of work have been committed, the entire group has been put successfully, and the STATUS queue is empty. If not, the group must be resumed at the point indicated by the status information. MQPMO_LOGICAL_ORDER cannot be used for the first put, but can thereafter.

Restart processing looks like this:

```

MQGET (StatusInfo from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
    /* Proceed to normal processing */
    ...
else
    /* Group was terminated prematurely */
    Set GroupId, MsgSeqNumber in MQMD to values from Status message
    PMO.Options = MQPMO_SYNCPOINT
    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

    /* Now normal processing is resumed.
       Assume this is not the last message */
    PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT
    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
    MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP
    StatusInfo = GroupId,MsgSeqNumber from MQMD
    MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
    MQCMIT

```

From the getting application, you might want to start processing the messages in a group before the whole group has arrived. This improves response times on the messages within the group, and also means that storage is not required for the entire group. In order to realize the benefits, use several units of work for each group of messages. For recovery reasons, you must retrieve each message within a unit of work.

As with the corresponding putting application, this requires status information to be recorded somewhere automatically as each unit of work is committed. Again, the simplest place to record this information is on a STATUS queue. If a complete group has been successfully processed, the STATUS queue is empty.

Note: For intermediate units of work, you can avoid the MQGET calls from the STATUS queue by specifying that each MQPUT to the status queue is a segment of a message (that is, by setting the MQMF_SEGMENT flag), instead of putting a complete new message for each unit of work. In the last unit of work, a final segment is put to the status queue specifying MQMF_LAST_SEGMENT, and then the status information is cleared with an MQGET specifying MQGMO_COMPLETE_MSG.

During restart processing, instead of using a single MQGET to get a possible status message, browse the status queue with MQGMO_LOGICAL_ORDER until you reach the last segment (that is, until no further segments are returned). In the first unit of work after restart, also specify the offset explicitly when putting the status segment.

In the following example, we consider only messages within a group, assuming that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. MQGMO_COMPLETE_MSG is therefore specified on each MQGET. The same principles apply if segmentation is involved (in this case, the StatusInfo must include the *Offset*).

For simplicity, we assume that a maximum of 4 messages are retrieved within a single UOW:

```

msgs = 0      /* Counts messages retrieved within UOW */
/* Should be no status message at this point */

/* Retrieve remaining messages in the group */
do while ( GroupStatus == MQGS_MSG_IN_GROUP )

    /* Process up to 4 messages in the group */
    GMO.Options = MQGMO_SYNCPOINT | MQGMO_WAIT
                | MQGMO_LOGICAL_ORDER
    do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
        MQGET
        msgs = msgs + 1
        /* Process this message */
        ...
    /* end while

    /* Have retrieved last message or 4 messages */
    /* Update status message if not last in group */
    MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
    if ( GroupStatus == MQGS_MSG_IN_GROUP )
        StatusInfo = GroupId,MsgSeqNumber from MQMD
        MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
    MQCMIT
    msgs = 0
/* end while

if ( msgs > 0 )
    /* Come here if there was only 1 message in the group */
    MQCMIT

```

If all the units of work have been committed, the entire group has been retrieved successfully, and the STATUS queue is empty. If not, the group must be resumed at the point indicated by the status information. MQGMO_LOGICAL_ORDER cannot be used for the first retrieve, but can thereafter.

Restart processing looks like this:

```

MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if (Reason == MQRC_NO_MSG_AVAILABLE)
    /* Proceed to normal processing */
    ...
else
    /* Group was terminated prematurely */
    /* The next message on the group must be retrieved by matching
       the sequence number and group id with those retrieved from the
       status information. */
    GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
    MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID | MQMO_MATCH_MSG_SEQ_NUMBER,
        MQMD.GroupId      = value from Status message,
        MQMD.MsgSeqNumber = value from Status message plus 1
    msgs = 1
    /* Process this message */
    ...

    /* Now normal processing is resumed */
    /* Retrieve remaining messages in the group */
    do while ( GroupStatus == MQGS_MSG_IN_GROUP )

        /* Process up to 4 messages in the group */
        GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT | MQGMO_WAIT
                    | MQGMO_LOGICAL_ORDER
        do while ( (GroupStatus == MQGS_MSG_IN_GROUP) && (msgs < 4) )
            MQGET
            msgs = msgs + 1
            /* Process this message */
            ...

        /* Have retrieved last message or 4 messages */
        /* Update status message if not last in group */

```

```

MQGET (from STATUS queue) GMO.Options = MQGMO_SYNCPOINT
if ( GroupStatus == MQGS_MSG_IN_GROUP )
    StatusInfo = GroupId,MsgSeqNumber from MQMD
    MQPUT (StatusInfo to STATUS queue) PMO.Options = MQPMO_SYNCPOINT
MQCMIT
msgs = 0

```

Getting a particular message

There are a number of ways of getting a particular message from a queue. These are: selecting on the *MsgId* and *CorrelId*, selecting on the *GroupId*, *MsgSeqNumber* and *Offset*, and selecting on the *MsgToken*. You can also use a selection string when you open the queue.

To get a particular message from a queue, use the *MsgId* and *CorrelId* fields of the MQMD structure. However, applications can explicitly set these fields, so the values that you specify might not identify a unique message. [Table 38 on page 246](#) shows which message is retrieved for the possible settings of these fields. These fields are ignored on input if you specify MQGMO_MSG_UNDER_CURSOR in the *GetMsgOpts* parameter of the MQGET call.

Table 38. Using message and correlation identifiers		
To retrieve ...	<i>MsgId</i>	<i>CorrelId</i>
First message in the queue	MQMI_NONE	MQCI_NONE
First message that matches <i>MsgId</i>	Nonzero	MQCI_NONE
First message that matches <i>CorrelId</i>	MQMI_NONE	Nonzero
First message that matches both <i>MsgId</i> and <i>CorrelId</i>	Nonzero	Nonzero

In each case, *first* means the first message that satisfies the selection criteria (unless MQGMO_BROWSE_NEXT is specified, when it means the *next* message in the sequence satisfying the selection criteria).

On return, the MQGET call sets the *MsgId* and *CorrelId* fields to the message and correlation identifiers of the message returned, if any.

If you set the *Version* field of the MQMD structure to 2, you can use the *GroupId*, *MsgSeqNumber*, and *Offset* fields. [Table 39 on page 246](#) shows which message is retrieved for the possible settings of these fields.

Table 39. Using the group identifier	
To retrieve ...	Match options
First message in the queue	MQMO_NONE
First message that matches <i>MsgId</i>	MQMO_MATCH_MSG_ID
First message that matches <i>CorrelId</i>	MQMO_MATCH_CORREL_ID
First message that matches <i>GroupId</i>	MQMO_MATCH_GROUP_ID
First message that matches <i>MsgSeqNumber</i>	MQMO_MATCH_MSG_SEQ_NUMBER
First message that matches <i>MsgToken</i>	MQMO_MATCH_MSG_TOKEN
First message that matches <i>Offset</i>	MQMO_MATCH_OFFSET

Table 39. Using the group identifier (continued)

To retrieve ...	Match options
Notes: <ol style="list-style-type: none"> 1. MQMO_MATCH_XXX implies that the XXX field in the MQMD structure is set to the value to be matched. 2. The MQMO flags can be used in combination. For example, MQMO_MATCH_GROUP_ID, MQMO_MATCH_MSG_SEQ_NUMBER, and MQMO_MATCH_OFFSET can be used together to give the segment identified by the <i>GroupId</i>, <i>MsgSeqNumber</i>, and <i>Offset</i> fields. 3. If you specify MQGMO_LOGICAL_ORDER, the message that you are trying to retrieve is affected because the option depends on state information controlled for the queue handle. For information about this, see “Logical and physical ordering” on page 236 and Options. 	

The MQGET call usually retrieves the first message from a queue. If you specify a particular message when you use the MQGET call, the queue manager must search the queue until it finds that message. This can affect the performance of your application.

If you are using Version 2 or later of the MQGMO structure and do not specify the MQMO_MATCH_MSG_ID or MQMO_MATCH_CORREL_ID flags, you do not need to reset the *MsgId* or *CorrelId* fields between MQGETs.

You can get a specific message from a queue by specifying its *MsgToken* and the *MatchOption* MQMO_MATCH_MSG_TOKEN in the MQGMO structure. The *MsgToken* is returned by the MQPUT call that originally put that message on the queue, or by previous MQGET operations and remains constant unless the queue manager is restarted.

If you are interested in only a subset of messages on the queue, you can specify which messages you want to process by using a selection string with the MQOPEN or MQSUB call. MQGET then retrieves the next message that satisfies that selection string. For more information about selection strings, see [“Selectors”](#) on page 22.

Improving performance of non-persistent messages

When a client requires a message from a server, it sends a request to the server. It sends a separate request for each of the messages it consumes. To improve the performance of a client consuming non persistent messages by avoiding having to send these request messages, a client can be configured to use *read ahead*. Read ahead allows messages to be sent to a client without an application having to request them.

When read ahead is enabled, messages are sent to a memory buffer on the client called the *read ahead buffer*. The client will have a read ahead buffer for each queue it has open with read ahead enabled. The messages in the read ahead buffer are not persisted. The client periodically updates the server with information about the amount of data it has consumed.

When you call MQOPEN with MQOO_READ_AHEAD, the WebSphere MQ client only enables readahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded WebSphere MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

Using read ahead can improve performance when consuming non persistent messages from a client application. This performance improvement is available to both MQI and JMS applications. Client applications using MQGET or asynchronous consumption will benefit from the performance improvements when consuming non-persistent messages.

Not all client application designs are suited to using read ahead as not all options are supported for use with read ahead and some options are required to be consistent between MQGET calls when read ahead is enabled. If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer will remain stranded in the client read ahead buffer.

If a backlog of stranded messages with the previous selection criteria are no longer required, a configurable purge interval can be set on the client to automatically purge these messages from the client. The purge interval is one of a group of read ahead tuning options determined by the client. It is possible to tune these options to meet your requirements.

If a client application is restarted, messages in the read ahead buffer can be lost. Conversely, a message that has been moved to a read ahead buffer could then be deleted from the underlying queue; this does not result in it being removed from the buffer, so an MQGET call using read ahead can return a message that no longer exists.

Read ahead is only performed for client bindings. The attribute is ignored for all other bindings.

Read ahead has no effect on triggering. No trigger message is generated when a message is read ahead by the client. Read ahead does not generate accounting and statistics information when it is enabled.

Using read ahead with publish subscribe messaging

When a subscribing application specifies a destination queue to which publications are sent, the DEFREADA value of the specified queue is used as the default read ahead value.

When a subscribing application requests that WebSphere MQ manages the destination to which publications are sent, a managed queue is created as a dynamic queue based upon a predefined model queue. It is the DEFREADA value of the model queue that is used as the default read ahead value. The default model queues SYSTEM.DURABLE.PUBLICATIONS.MODEL or SYSTEM.NONDURABLE.PUBLICATIONS.MODEL are used unless a model queue is defined for this or a parent topic.

Related concepts

[“Tuning performance for nonpersistent messages on AIX” on page 250](#)

If you are using AIX V5.3 or later, consider setting your tuning parameter to use full performance for nonpersistent messages.

Related tasks

[“Enabling and disabling read ahead” on page 249](#)

By default read ahead is disabled. You can enable read ahead at queue or application level.

Related reference

[“MQGET options and read ahead” on page 248](#)

Not all MQGET options are supported when read ahead is enabled; some options are required to be consistent between MQGET calls.

MQGET options and read ahead

Not all MQGET options are supported when read ahead is enabled; some options are required to be consistent between MQGET calls.

When you call MQOPEN with MQOO_READ_AHEAD, the WebSphere MQ client only enables readahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded WebSphere MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

The following table indicates which options are supported for use with read ahead and whether they can be altered between MQGET calls.

Table 40. MQGET options and read ahead			
	Permitted when read ahead is enabled and can be altered between MQGET calls ⁵	Permitted when read ahead is enabled but cannot be altered between MQGET calls ⁴	MQGET Options that are not permitted when read ahead is enabled ²
MQGET MQMD values	MsgId ³ CorrelId ³	Encoding CodedCharSetId	
MQGET MQGMO Options	<ul style="list-style-type: none"> • MQGMO_NO_WAIT • MQGMO_BROWSE_MESSAGE_UNDER_CURSOR • MQGMO_BROWSE_FIRST • MQGMO_BROWSE_NEXT • MQGMO_FAIL_IF QUIESCING 	<ul style="list-style-type: none"> • MQGMO_SYNCPOINT_IF_PERSISTENT • MQGMO_NO_SYNCPOINT • MQGMO_ACCEPT_TRUNCATED_MSG • MQGMO_CONVERT 	<ul style="list-style-type: none"> • MQGMO_SET_SIGNAL • MQGMO_SYNCPOINT • MQGMO_MARK_SKIP_BACKOUT • MQGMO_MSG_UNDER_CURSOR⁴ • MQGMO_LOCK • MQGMO_UNLOCK • MQGMO_LOGICAL_ORDER • MQGMO_COMPLETE_MSG • MQGMO_ALL_MSGS_AVAILABLE • MQGMO_ALL_SEGMENTS_AVAILABLE

Notes:

1. If these options are altered between MQGET calls, an MQRC_OPTIONS_CHANGED reason code is returned.
2. If these options are specified on the first MQGET call then read ahead is disabled. If these options are specified on a subsequent MQGET call a reason code MQRC_OPTIONS_ERROR is returned.
3. If a client application alters MsgId and CorrelId values between MQGET calls, messages with the previous values might already have been sent to the client and will remain in the client read ahead buffer until consumed (or automatically purged).
4. MQGMO_MSG_UNDER_CURSOR is not possible with read ahead. Read ahead is disabled when both MQOO_BROWSE and one of the MQOO_INPUT_SHARED or MQOO_INPUT_EXCLUSIVE options are specified when opening the queue.
5. When read ahead is enabled, the first MQGET determines whether messages are to be browsed or got from a queue. If the client application then uses MQGET with changed options, such as attempting to browse following an initial get, or attempting to get following an initial browse, an MQRC_OPTIONS_CHANGED reason code is returned.

If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer that match the initial selection criteria are not consumed by the client application, and remain stranded in the client read ahead buffer. In situations where the client read ahead buffer contains many stranded messages, the benefits associated with read ahead are lost and a separate request to the server is required for each message consumed. To determine whether read ahead is being used efficiently you can use the connection status parameter, READA.

Read ahead can be inhibited when requested by an application due to incompatible options specified on the first MQGET call. In this situation the connection status shows read ahead as being inhibited.

If, because of these restrictions on MQGET, you decide that a client application design is not suited to read ahead, specify the MQOPEN option MQOO_READ_AHEAD_NO. Alternatively set the default read ahead value of the queue being opened altered to either NO or DISABLED.

Enabling and disabling read ahead

By default read ahead is disabled. You can enable read ahead at queue or application level.

About this task

When you call MQOPEN with MQOO_READ_AHEAD, the WebSphere MQ client only enables readahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded WebSphere MQ MQI client libraries.

- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

To enable read ahead:

- To configure read ahead at the queue level set the queue attribute, DEFREADA to YES.
- To configure read ahead at the application level:
 - to use read ahead wherever possible use the MQOO_READ_AHEAD option on the MQOPEN function call. It is not possible for the client application to use read ahead if the DEFREADA queue attribute has been set to DISABLED.
 - to use read ahead only when read ahead is enabled on a queue, use the MQOO_READ_AHEAD_AS_Q_DEF option on the MQOPEN function call.

If a client application design is not suited to read ahead you can disable it:

- at the queue level by setting the queue attribute, DEFREADA to NO if you do not want read ahead to be used unless it is requested by a client application, or DISABLED if you do not want read ahead to be used regardless of whether read ahead is required by a client application.
- at the application level by using the MQOO_NO_READ_AHEAD option on the MQOPEN function call.

Two MQCLOSE options allow you to configure what happens to any messages that are being stored in the read ahead buffer if the queue is closed.

- Use MQCO_IMMEDIATE to discard messages in the read ahead buffer.
- Use MQCO_QUIESCE to ensure that messages in the read ahead buffer are consumed by the application before the queue is closed. When MQCLOSE with the MQCO_QUIESCE is issued and there are messages remaining on the read ahead buffer, MQRC_READ_AHEAD_MSGS returns with MQCC_WARNING.

Tuning performance for nonpersistent messages on AIX

If you are using AIX V5.3 or later, consider setting your tuning parameter to use full performance for nonpersistent messages.

To set the tuning parameter so that it takes effect immediately, issue the following command as a root user:

```
/usr/sbin/ios -o j2_nPagesPerWriteBehindCluster=0
```

To set the tuning parameter so that it takes effect immediately and persists over reboots, issue the following command as a root user:

```
/usr/sbin/ios -p -o j2_nPagesPerWriteBehindCluster=0
```

Normally, nonpersistent messages are kept only in memory, but there are circumstances where AIX can schedule nonpersistent messages to be written to disk. Messages scheduled to be written to disk are unavailable for MQGET until the disk write completes. The suggested tuning command varies this threshold; instead of scheduling messages to be written to disk when 16 kilobytes of data are queued, the write-to-disk occurs only when real storage on the machine becomes close to full. This is a global alteration and might affect other software components.

On AIX, when using multithreaded applications and especially when running on machines with multiple processors, we strongly recommend setting AIXTHREAD_SCOPE=S in the mqm id .profile or setting AIXTHREAD_SCOPE=S in the environment before starting the application, for better performance and more solid scheduling. For example:

```
export AIXTHREAD_SCOPE=S
```

Setting AIXTHREAD_SCOPE=S means that user threads created with default attributes are placed into system-wide contention scope. If a user thread is created with system-wide contention scope, it is bound

to a kernel thread and it is scheduled by the kernel. The underlying kernel thread is not shared with any other user thread.

File descriptors

When running a multi-threaded process such as the agent process, you might reach the soft limit for file descriptors. This limit gives you the IBM WebSphere MQ reason code MQRC_UNEXPECTED_ERROR (2195) and, if there are enough file descriptors, an IBM WebSphere MQ FFST™ file.

To avoid this problem, you can increase the process limit for the number of file descriptors. To do so, alter the `nfiles` attribute in `/etc/security/limits` to 10,000 for the `mqm` user ID or in the default stanza.

System Resource Limits

Set the system resource limit for data segment and stack segment to unlimited using the following commands in a command prompt:

```
ulimit -d unlimited
ulimit -s unlimited
```

Handling messages greater than 4 MB long

Messages can be too large for the application, queue, or queue manager. Depending on the environment, WebSphere MQ provides a number of ways of dealing with messages that are longer than 4 MB.

You can increase the *MaxMsgLength* attribute up to 100 MB on all WebSphere MQ systems at V6 or later. Set this value to reflect the size of the messages using the queue. On WebSphere MQ systems other than WebSphere MQ for z/OS, you can also:

1. Use segmented messages. (Messages can be segmented by either the application or the queue manager.)
2. Use reference messages.

Each of these approaches is described in the remainder of this section.

Increasing the maximum message length

The *MaxMsgLength* queue manager attribute defines the maximum length of a message that can be handled by a queue manager. Similarly, the *MaxMsgLength* queue attribute is the maximum length of a message that can be handled by a queue. The *default* maximum message length supported depends on the environment in which you are working.

If you are handling large messages, you can alter these attributes independently. You can set the queue manager attribute value in the range 32768 bytes through 100 MB; you can set the queue attribute value in the range 0 through 100 MB.

After changing one or both of the *MaxMsgLength* attributes, restart your applications and channels to ensure that the changes take effect.

When these changes are made, the message length must be less than or equal to both the queue and the queue manager *MaxMsgLength* attributes. However, existing messages might be longer than either attribute.

If the message is too big for the queue, MQRC_MSG_TOO_BIG_FOR_Q is returned. Similarly, if the message is too big for the queue manager, MQRC_MSG_TOO_BIG_FOR_Q_MGR is returned.

This method of handling large messages is easy and convenient. However, consider the following factors before using it:

- Uniformity among queue managers is reduced. The maximum size of message data is determined by the *MaxMsgLength* for each queue (including transmission queues) on which the message will be put. This

value is often defaulted to the queue manager's *MaxMsgLength*, especially for transmission queues. This makes it difficult to predict whether a message is too large when it is to travel to a remote queue manager.

- Usage of system resources is increased. For example, applications need larger buffers, and on some platforms, there might be increased usage of shared storage. Queue storage should be affected only if actually required for larger messages.
- Channel batching is affected. A large message still counts as just one message towards the batch count but needs longer to transmit, thereby increasing response times for other messages.

Message segmentation

Use this information to learn about segmenting messages.

Note: Not supported in IBM WebSphere MQ for z/OS or by applications using IBM WebSphere MQ classes for JMS.

Increasing the maximum message length as explained in topic [“Increasing the maximum message length”](#) on page 251 has some negative implications. Also, it can still result in the message being too large for the queue or queue manager. In these cases, you can segment a message. For information about segments, see [“Message groups”](#) on page 34.

The next sections look at common uses for segmenting messages. For putting and destructively getting, it is assumed that the MQPUT or MQGET calls *always* operate within a unit of work. Always consider using this technique to reduce the possibility of incomplete groups being present in the network. Single-phase commit by the queue manager is assumed, but other coordination techniques are equally valid.

Also, in the getting applications, it is assumed that if multiple servers are processing the same queue, each server executes similar code, so that one server never fails to find a message or segment that it expects to be there (because it had specified MQGMO_ALL_MSGS_AVAILABLE or MQGMO_ALL_SEGMENTS_AVAILABLE earlier).

Putting and getting a segmented message that spans units of work

You can put and get a segmented message that spans a unit of work in a similar way to [“Putting and getting a group that spans units of work”](#) on page 243.

You cannot, however, put or get segmented messages in a global unit of work.

Segmentation and reassembly by queue manager

This is the simplest scenario, in which one application puts a message to be retrieved by another. The message might be large: not too large for either the putting or the getting application to handle in a single buffer, but too large for the queue manager or a queue on which the message is to be put.

The only changes necessary for these applications are for the putting application to authorize the queue manager to perform segmentation if necessary:

```
PMO.Options = (existing options)
MD.MsgFlags = MQMF_SEGMENTATION_ALLOWED
memcpy(MD.GroupId, MQGI_NONE, MQ_GROUP_ID_LENGTH)
MQPUT
```

and for the getting application to ask the queue manager to reassemble the message if it has been segmented:

```
GMO.Options = MQGMO_COMPLETE_MSG | (existing options)
MQGET
```

In this simplest scenario, the application must reset the GroupId field to MQGI_NONE before the MQPUT call, so that the queue manager can generate a unique group identifier for each message. If this is not done, unrelated messages can have the same group identifier, which might subsequently lead to incorrect processing.

The application buffer must be large enough to contain the reassembled message (unless you include the MQGMO_ACCEPT_TRUNCATED_MSG option).

If the MAXMSGLEN attribute of a queue is to be modified to accommodate message segmentation, then consider:

- The minimum message segment supported on a local queue is 16 bytes.
- For a transmission queue, MAXMSGLEN must also include the space required for headers. Consider using a value at least 4000 bytes larger than the maximum expected length of user data in any message segment that could be put on a transmission queue.

If data conversion is necessary, the getting application might have to do it by specifying MQGMO_CONVERT. This should be straightforward because the data conversion exit is presented with the complete message. Do not attempt to convert data in a sender channel if the message is segmented, and the format of the data is such that the data-conversion exit cannot carry out the conversion on incomplete data.

Application segmentation

Application segmentation is used when queue-manager segmentation is not adequate, or when applications require data conversion with specific segment boundaries.

Application segmentation is used for two main reasons:

1. Queue-manager segmentation alone is not adequate because the message is too large to be handled in a single buffer by the applications.
2. Data conversion must be performed by sender channels, and the format is such that the putting application must stipulate where the segment boundaries are to be in order for conversion of an individual segment to be possible.

However, if data conversion is not an issue, or if the getting application always uses MQGMO_COMPLETE_MSG, queue-manager segmentation can also be allowed by specifying MQMF_SEGMENTATION_ALLOWED. In our example, the application segments the message into four segments:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_SEGMENT

MQCMIT
```

If you do not use MQPMO_LOGICAL_ORDER, the application must set the *Offset* and the length of each segment. In this case, logical state is not maintained automatically.

The getting application cannot guarantee to have a buffer large enough to hold any reassembled message. It must therefore be prepared to process segments individually.

For messages that are segmented, this application does not want to start processing one segment until all the segments that constitute the logical message are present. MQGMO_ALL_SEGMENTS_AVAILABLE is therefore specified for the first segment. If you specify MQGMO_LOGICAL_ORDER and there is a current logical message, MQGMO_ALL_SEGMENTS_AVAILABLE is ignored.

After the first segment of a logical message has been retrieved, use MQGMO_LOGICAL_ORDER to ensure that the remaining segments of the logical message are retrieved in order.

No consideration is given to messages within different groups. If such messages occur, they are processed in the order in which the first segment of each message occurs on the queue.

```
GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_SEGMENTS_AVAILABLE | MQGMO_WAIT
do while ( SegmentStatus == MQSS_SEGMENT )
  MQGET
  /* Process each remaining segment of the logical message */
```

```
...
MQCMIT
```

Application segmentation of logical messages

The messages must be maintained in logical order in a group, and some or all of them might be so large that they require application segmentation.

In our example, a group of four logical messages is to be put. All but the third message are large, and require segmentation, which is performed by the putting application:

```
PMO.Options = MQPMO_LOGICAL_ORDER | MQPMO_SYNCPOINT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP      | MQMF_LAST_SEGMENT

MQPUT MD.MsgFlags = MQMF_MSG_IN_GROUP

MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_SEGMENT
MQPUT MD.MsgFlags = MQMF_LAST_MSG_IN_GROUP | MQMF_LAST_SEGMENT

MQCMIT
```

In the getting application, MQGMO_ALL_MSGS_AVAILABLE is specified on the first MQGET. This means that no messages or segments of a group are retrieved until the entire group is available. When the first physical message of a group has been retrieved, MQGMO_LOGICAL_ORDER is used to ensure that the segments and messages of the group are retrieved in order:

```
GMO.Options = MQGMO_SYNCPOINT | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_WAIT

do while ( (GroupStatus != MQGS_LAST_MSG_IN_GROUP) ||
           (SegmentStatus != MQGS_LAST_SEGMENT) )
    MQGET
    /* Process a segment or complete logical message. Use the GroupStatus
       and SegmentStatus information to see what has been returned */
    ...
MQCMIT
```

Note: If you specify MQGMO_LOGICAL_ORDER and there is a current group, MQGMO_ALL_MSGS_AVAILABLE is ignored.

Reference messages

Use this information to learn more about reference messages.

Note: Not supported in WebSphere MQ for z/OS.

This method allows a large object to be transferred from one node to another without storing the object on WebSphere MQ queues at either the source or the destination nodes. This is of particular benefit when the data exists in another form, for example, for mail applications.

To do this, you specify a message exit at both ends of a channel. For information about how to do this, see [“Channel message exit programs” on page 393](#).

WebSphere MQ defines the format of a reference message header (MQRMH). See [MQRMH](#) for a description of this. This is recognized with a defined format name and might be followed by actual data.

To initiate transfer of a large object, an application can put a message consisting of a reference message header with no data following it. As this message leaves the node, the message exit retrieves the object in an appropriate way and appends it to the reference message. It then returns the message (now larger than before) to the sending Message Channel Agent for transmission to the receiving MCA.

Another message exit is configured at the receiving MCA. When this message exit receives one of these messages, it creates the object using the object data that was appended and passes on the reference message *without* it. The reference message can now be received by an application and this application knows that the object (or at least the portion of it represented by this reference message) has been created at this node.

The maximum amount of object data that a sending message exit can append to the reference message is limited by the negotiated maximum message length for the channel. The exit can return only a single message to the MCA for each message that it is passed, so the putting application can put several messages to cause one object to be transferred. Each message must identify the *logical* length and offset of the object that is to be appended to it. However, in cases where it is not possible to know the total size of the object or the maximum size allowed by the channel, design the sending message exit so that the putting application just puts a single message, and the exit itself puts the next message on the transmission queue when it has appended as much data as it can to the message it has been passed.

Before using this method of dealing with large messages, consider the following points:

- The MCA and the message exit run under a WebSphere MQ user ID. The message exit (and therefore, the user ID) needs to access the object to either retrieve it at the sending end or create it at the receiving end; this might only be feasible in cases where the object is universally accessible. This raises a security issue.
- If the reference message with bulk data appended to it must travel through several queue managers before reaching its destination, the bulk data *is* present on WebSphere MQ queues at the intervening nodes. However, no special support or exits need to be provided in these cases.
- Designing your message exit is made difficult if rerouting or dead-letter queuing is allowed. In these cases, the portions of the object might arrive out of order.
- When a reference message arrives at its destination, the receiving message exit creates the object. However, this is not synchronized with the MCA's unit of work, so if the batch is backed out, another reference message containing this same portion of the object will arrive in a later batch, and the message exit might attempt to re-create the same portion of the object. If the object is, for example, a series of database updates, this might be unacceptable. If so, the message exit must keep a log of which updates have been applied; this might require the use of a WebSphere MQ queue.
- Depending on the characteristics of the object type, the message exits and applications might need to cooperate in maintaining use counts, so that the object can be deleted when it is no longer needed. An instance identifier might also be required; a field is provided for this in the reference message header (see [MQRMH](#)).
- If a reference message is put as a distribution list, the object must be retrievable for each resulting distribution list or individual destination at that node. You might need to maintain use counts. Also consider the possibility that a node might be the final node for some of the destinations in the list, but an intermediate node for others.
- Bulk data is not typically converted. This is because conversion takes place *before* the message exit is invoked. For this reason, conversion must not be requested on the originating sender channel. If the reference message passes through an intermediate node, the bulk data is converted when sent from the intermediate node, if requested.
- Reference messages cannot be segmented.

Using the MQRMH and MQMD structures

See [MQRMH](#) and [MQMD](#) for a description of the fields in the reference message header and the message descriptor.

In the MQMD structure, set the *Format* field to MQFMT_REF_MSG_HEADER. The MQHREF format, when requested on MQGET, is converted automatically by WebSphere MQ along with any bulk data that follows.

Here is an example of the use of the *DataLogicalOffset* and *DataLogicalLength* fields of the MQRMH:

A putting application might put a reference message with:

- No physical data
- *DataLogicalLength* = 0 (this message represents the entire object)
- *DataLogicalOffset* = 0.

Assuming that the object is 70 000 bytes long, the sending message exit sends the first 40 000 bytes along the channel in a reference message containing:

- 40 000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 40000
- *DataLogicalOffset* = 0 (from the start of the object).

It then places another message on the transmission queue containing:

- No physical data
- *DataLogicalLength* = 0 (to the end of the object). You could specify a value of 30 000 here.
- *DataLogicalOffset* = 40000 (starting from this point).

When this message exit is seen by the sending message exit, the remaining 30,000 bytes of data are appended, and the fields are set to:

- 30,000 bytes of physical data following the MQRMH
- *DataLogicalLength* = 30000
- *DataLogicalOffset* = 40000 (starting from this point).

The MQRMHF_LAST flag is also set.

For a description of the sample programs provided for the use of reference messages, see [“Sample programs for distributed platforms” on page 92](#).

Waiting for messages

If you want a program to wait until a message arrives on a queue, specify the MQGMO_WAIT option in the *Options* field of the MQGMO structure.

Use the *WaitInterval* field of the MQGMO structure to specify the maximum time (in milliseconds) that you want an MQGET call to wait for a message to arrive on a queue.

If the message does not arrive within this time, the MQGET call completes with the MQRC_NO_MSG_AVAILABLE reason code.

You can specify an unlimited wait interval using the constant MQWI_UNLIMITED in the *WaitInterval* field. However, events outside your control could cause your program to wait for a long time, so use this constant with caution. IMS applications must not specify an unlimited wait interval because this would prevent the IMS system terminating. (When IMS terminates, it requires all dependent regions to end.) Instead, IMS applications can specify a finite wait interval; then, if the call completes without retrieving a message after that interval, issue another MQGET call with the wait option.

Note: If more than one program is waiting on the same shared queue to *remove* a message, only one program is activated by a message arriving. However, if more than one program is waiting to browse a message, all the programs can be activated. For more information, see the description of the *Options* field of the MQGMO structure in [MQGMO](#).

If the state of the queue or the queue manager changes before the wait interval expires, the following actions occur:

- If the queue manager enters the quiescing state, and you used the MQGMO_FAIL_IF QUIESCING option, the wait is canceled and the MQGET call completes with the MQRC_Q_MGR QUIESCING reason code. Without this option, the call remains waiting.
- If the queue manager is forced to stop, or is canceled, the MQGET call completes with either the MQRC_Q_MGR STOPPING or the MQRC_CONNECTION_BROKEN reason code.

- If the attributes of the queue (or a queue to which the queue name resolves) are changed so that get requests are now inhibited, the wait is canceled and the MQGET call completes with the MQRC_GET_INHIBITED reason code.
- If the attributes of the queue (or a queue to which the queue name resolves) are changed in such a way that the FORCE option is required, the wait is canceled and the MQGET call completes with the MQRC_OBJECT_CHANGED reason code.

For more information about the circumstances in which these actions occur, see [MQGMO](#).

Skipping backout

You can prevent an application program from entering an *MQGET-error-backout* loop by specifying the **MQGMO_MARK_SKIP_BACKOUT** option on the MQGET call.

Note: Supported only on WebSphere MQ for z/OS.

As part of a unit of work, an application program can issue one or more MQGET calls to get messages from a queue. If the application program detects an error, it can back out the unit of work. This restores all the resources updated during that unit of work to the state that they were in before the unit of work started, and reinstates the messages retrieved by the MQGET calls.

Once reinstated, these messages are available to subsequent MQGET calls issued by the application program. In many cases, this does not cause a problem for the application program. However, in cases where the error leading to the backout cannot be circumvented, having the message reinstated on the queue can cause the application program to enter an *MQGET-error-backout* loop.

To avoid this problem, specify the MQGMO_MARK_SKIP_BACKOUT option on the MQGET call. This marks the MQGET request as not being involved in application-initiated backout; that is, it must not be backed out. Use of this option means that when a backout occurs, updates to other resources are backed out as required, but the marked message is treated as if it had been retrieved under a new unit of work.

The application program must issue a WebSphere MQ call either to commit the new unit of work, or to back out the new unit of work. For example, the program can perform exception handling, such as informing the originator that the message has been discarded, and commit the unit of work so removing the message from the queue. If the new unit of work is backed out (for any reason) the message is reinstated on the queue.

Within a unit of work, there can be only one MQGET request marked as skipping backout; however, there can be several other messages that are not marked as skipping backout. Once a message has been marked as skipping backout, any further MQGET calls within the unit of work that specify MQGMO_MARK_SKIP_BACKOUT fail with reason code MQRC_SECOND_MARK_NOT_ALLOWED.

Note:

1. The marked message skips backout only if the unit of work containing it is terminated by an application request to back it out. If the unit of work is backed out for any other reason, the message is backed out onto the queue in the same way that it would be if it was not marked to skip backout.
2. Skip backout is not supported within DB2 stored procedures participating in units of work controlled by RRS. For example, an MQGET call with the MQGMO_MARK_SKIP_BACKOUT option will fail with the reason code MQRC_OPTION_ENVIRONMENT_ERROR.

[Figure 36 on page 258](#) illustrates a typical sequence of steps that an application program might contain when an MQGET request is required to skip backout.




Figure 36. Skipping backout using MQGMO_MARK_SKIP_BACKOUT

The steps in [Figure 36 on page 258](#) are:

Step 1

Initial processing occurs within the transaction, including an MQOPEN call to open the queue (specifying one of the MQOO_INPUT_* options in order to get messages from the queue in Step 2).

Step 2

MQGET is called, with MQGMO_SYNCPOINT and MQGMO_MARK_SKIP_BACKOUT. MQGMO_SYNCPOINT is required because MQGET must be within a unit of work for MQGMO_MARK_SKIP_BACKOUT to be effective. In [Figure 36 on page 258](#) this unit of work is referred to as UOW1.

Step 3

Other resource updates are made as part of UOW1. These can include further MQGET calls (issued without MQGMO_MARK_SKIP_BACKOUT).

Step 4

All updates from Steps 2 and 3 complete as required. The application program commits the updates and UOW1 ends. The message retrieved in Step 2 is removed from the queue.

Step 5

Some of the updates from Steps 2 and 3 do not complete as required. The application program requests that the updates made during these steps are backed out.

Step 6

The updates made in Step 3 are backed out.

Step 7

The MQGET request made in Step 2 skips backout and becomes part of a new unit of work, UOW2.

Step 8

UOW2 performs exception handling in response to UOW1 being backed out. (For example, an MQPUT call to another queue, indicating that a problem occurred that caused UOW1 to be backed out.)

Step 9

Step 8 completes as required, the application program commits the activity, and UOW2 ends. As the MQGET request is part of UOW2 (see Step 7), this commit causes the message to be removed from the queue.

Step 10

Step 8 does not complete as required and the application program backs out UOW2. Because the get message request is part of UOW2 (see Step 7), it too is backed out and reinstated on the queue. It is now available to further MQGET calls issued by this or another application program (in the same way as any other message on the queue).

Application data conversion

When necessary, MCAs convert the message descriptor and header data into the required character set and encoding. Either end of the link (that is, the local MCA or the remote MCA) can do the conversion.

When an application puts messages on a queue, the local queue manager adds control information to the message descriptors to facilitate the control of the messages when they are processed by queue managers and MCAs. Depending on the environment, the message header data fields are created in the character set and encoding of the local system.

When you move messages between systems, you sometimes need to convert the application data into the character set and encoding required by the receiving system. This can be done either from within application programs on the receiving system or by the MCAs on the sending system. If data conversion is supported on the receiving system, use application programs to convert the application data, rather than depending on the conversion having already occurred at the sending system.

Application data is converted within an application program when you specify the MQGMO_CONVERT option in the *Options* field of the MQGMO structure passed to an MQGET call, and *all* the following are true:

- The *CodedCharSetId* or *Encoding* fields set in the MQMD structure associated with the message on the queue differ from the *CodedCharSetId* or *Encoding* fields set in the MQMD structure specified on the MQGET call.
- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE.
- The *BufferLength* specified on the MQGET call is not zero.
- The message data length is not zero.
- The queue manager supports conversion between the *CodedCharSetId* and *Encoding* fields specified in the MQMD structures associated with the message and the MQGET call. See [CodedCharSetId](#) and [Encoding](#) for details of the coded character set identifiers and machine encodings supported.
- The queue manager supports conversion of the message format. If the *Format* field of the MQMD structure associated with the message is one of the built-in formats, the queue manager can convert the message. If the *Format* is not one of the built-in formats, you need to write a data-conversion exit to convert the message.

If the sending MCA is to convert the data, specify the CONVERT(YES) keyword on the definition of each sender or server channel for which conversion is required. If the data conversion fails, the message is sent to the DLQ at the sending queue manager and the *Feedback* field of the MQDLH structure indicates the reason. If the message cannot be put on the DLQ, the channel closes and the unconverted message remains on the transmission queue. Data conversion within applications rather than at sending MCAs avoids this situation.

As a rule, data in the message that is described as *character* data by the built-in format or data-conversion exit is converted from the coded character set used by the message to that requested, and *numeric* fields are converted to the encoding requested.

For further details of the conversion processing conventions used when converting the built-in formats, and for information about writing your own data-conversion exits, see [“Writing data-conversion exits”](#) on page 397. See also [National languages and Machine encodings](#) for information about the language support tables and about the supported machine encodings.

Conversion of EBCDIC newline characters

If you need to ensure that the data that you send from an EBCDIC platform to an ASCII one is identical to the data that you receive back again, you must control the conversion of EBCDIC newline characters.

You can do this using a platform-dependent switch that forces WebSphere MQ to use the unmodified conversion tables, but you must be aware of the inconsistent behavior that might result.

The problem arises because the EBCDIC newline character is not converted consistently across platforms or conversion tables. As a result, if the data is displayed on an ASCII platform, the formatting might be incorrect. This would make it difficult, for example, to administer an IBM i system remotely from an ASCII platform using RUNMQSC.

See [Data conversion](#) for further information about converting EBCDIC-format data to ASCII format.

Browsing messages on a queue

Use this information to find out about browsing messages on a queue using the MQGET call.

To use the MQGET call to browse the messages on a queue:

1. Call MQOPEN to open the queue for browsing, specifying the MQOO_BROWSE option.
2. To browse the first message on the queue, call MQGET with the MQGMO_BROWSE_FIRST option. To find the message that you want, call MQGET repeatedly with the MQGMO_BROWSE_NEXT option to step through many messages.

You *must* set the *MsgId* and *CorrelId* fields of the MQMD structure to null after each MQGET call in order to see all messages.

3. Call MQCLOSE to close the queue.

The browse cursor

When you open (MQOPEN) a queue for browsing, the call establishes a browse cursor for use with MQGET calls that use one of the browse options. You can think of the browse cursor as a logical pointer that is positioned before the first message on the queue.

You can have more than one browse cursor active (from a single program) by issuing several MQOPEN requests for the same queue.

When you call MQGET for browsing, use one of the following options in your MQGMO structure:

MQGMO_BROWSE_FIRST

Gets a copy of the first message that satisfies the conditions specified in your MQMD structure.

MQGMO_BROWSE_NEXT

Gets a copy of the next message that satisfies the conditions specified in your MQMD structure.

MQGMO_BROWSE_MSG_UNDER_CURSOR

Gets a copy of the message currently pointed to by the cursor, that is, the one that was last retrieved using either the MQGMO_BROWSE_FIRST or the MQGMO_BROWSE_NEXT option.

In all cases, the message remains on the queue.

When you open a queue, the browse cursor is positioned logically just before the first message on the queue. This means that if you make your MQGET call immediately after your MQOPEN call, you can use the MQGMO_BROWSE_NEXT option to browse the first message; you do not have to use the MQGMO_BROWSE_FIRST option.

The order in which messages are copied from the queue is determined by the *MsgDeliverySequence* attribute of the queue. (For more information, see [“The order in which messages are retrieved from a queue”](#) on page 235.)

- [“Queues in FIFO \(first in, first out\) sequence”](#) on page 261
- [“Queues in priority sequence”](#) on page 261
- [“Uncommitted messages”](#) on page 261
- [“Change to queue sequence”](#) on page 261
- [“Using the queue's index”](#) on page 261

Queues in FIFO (first in, first out) sequence

The first message in a queue in this sequence is the message that has been on the queue the longest.

Use MQGMO_BROWSE_NEXT to read the messages sequentially in the queue. You will see any messages put to the queue while you are browsing, as a queue in this sequence has messages placed at the end. When the cursor recognizes that it has reached the end of the queue, the browse cursor stays where it is and returns with MQRC_NO_MSG_AVAILABLE. You can then either leave it there waiting for further messages or reset it to the beginning of the queue with a MQGMO_BROWSE_FIRST call.

Queues in priority sequence

The first message in a queue in this sequence is the message that has been on the queue the longest and that has the highest priority at the time that the MQOPEN call is issued.

Use MQGMO_BROWSE_NEXT to read the messages in the queue.

The browse cursor points to the next message, working from the priority of the first message to finish with the message at the lowest priority. It browses any messages put to the queue during this time as long as they are of priority equal to, or lower than, the message identified by the current browse cursor.

Any messages put to the queue of higher priority can be browsed only by:

- Opening the queue for browse again, at which point a new browse cursor is established
- Using the MQGMO_BROWSE_FIRST option

Uncommitted messages

An uncommitted message is never visible to a browse; the browse cursor skips past it.

Messages within a unit-of-work cannot be browsed until the unit-of-work is committed. Messages do not change their position on the queue when committed, so skipped, uncommitted messages will not be seen, even when they *are* committed, unless you use the MQGMO_BROWSE_FIRST option and work through the queue again.

Change to queue sequence

If the message delivery sequence is changed from priority to FIFO while there are messages on the queue, the order of the messages that are already queued is not changed. Messages added to the queue later, take the default priority of the queue.

Using the queue's index

When you browse an indexed queue that contains only messages of a single priority (either persistent or nonpersistent or both), the queue manager uses the index to browse when certain forms of browse are used.

Note: Supported only on WebSphere MQ for z/OS.

Any of the following forms of browse are used when an indexed queue contains only messages of single priority:

1. If the queue is indexed by MSGID, browse requests that pass a MSGID in the MQMD structure are processed using the index to find the target message.
2. If the queue is indexed by CORRELID, browse requests that pass a CORRELID in the MQMD structure are processed using the index to find the target message.
3. If the queue is indexed by GROUPID, browse requests that pass a GROUPID in the MQMD structure are processed using the index to find the target message.

If the browse request does not pass a MSGID, CORRELID, or GROUPID in the MQMD structure, the queue is indexed, and a message is returned, the index entry for the message must be found, and information within it used to update the browse cursor. If you use a wide selection of index values, this does not add significant extra processing to the browse request.

Browsing messages when the message length is unknown

To browse a message when you do not know the size of the message, and you do not want to use the *MsgId*, *CorrelId*, or *GroupId* fields to locate the message, you can use the MQGMO_BROWSE_MSG_UNDER_CURSOR option:

1. Issue an MQGET with:
 - Either the MQGMO_BROWSE_FIRST or MQGMO_BROWSE_NEXT option
 - The MQGMO_ACCEPT_TRUNCATED_MSG option
 - Buffer length zero

Note: If another program is likely to get the same message, consider using the MQGMO_LOCK option as well. MQRC_TRUNCATED_MSG_ACCEPTED should be returned.

2. Use the returned *DataLength* to allocate the storage needed.
3. Issue an MQGET with the MQGMO_BROWSE_MSG_UNDER_CURSOR.

The message pointed to is the last one that was retrieved; the browse cursor will not have moved. You can choose either to lock the message using the MQGMO_LOCK option, or to unlock a locked message using MQGMO_UNLOCK option.

The call fails if no MQGET with either the MQGMO_BROWSE_FIRST or MQGMO_BROWSE_NEXT options has been issued successfully since the queue was opened.

Removing a message that you have browsed

You can remove from the queue a message that you have already browsed provided that you have opened the queue for removing messages as well as for browsing. (You must specify one of the MQOO_INPUT_* options, as well as the MQOO_BROWSE option, on your MQOPEN call.)

To remove the message, call MQGET again, but in the *Options* field of the MQGMO structure, specify MQGMO_MSG_UNDER_CURSOR. In this case, the MQGET call ignores the *MsgId*, *CorrelId*, and *GroupId* fields of the MQMD structure.

In the time between your browsing and removal steps, another program might have removed messages from the queue, including the message under your browse cursor. In this case, your MQGET call returns a reason code to say that the message is not available.

Browsing messages in logical order

“Logical and physical ordering” on page 236 explains the difference between the logical and physical order of messages on a queue. This distinction is particularly important when browsing a queue, because, in general, messages are not being deleted and browse operations do not necessarily start at the beginning of the queue.

If an application browses through the various messages of one group (using logical order), it is important that logical order should be followed to reach the start of the next group, because the last message of one group might occur physically *after* the first message of the next group. The MQGMO_LOGICAL_ORDER option ensures that logical order is followed when scanning a queue.

Use MQGMO_ALL_MSGS_AVAILABLE (or MQGMO_ALL_SEGMENTS_AVAILABLE) with care for browse operations. Consider the case of logical messages with MQGMO_ALL_MSGS_AVAILABLE. The effect of this is that a logical message is available only if all the remaining messages in the group are also present. If they are not, the message is passed over. This can mean that when the missing messages arrive subsequently, they are not noticed by a browse-next operation.

For example, if the following logical messages are present,

```
Logical message 1 (not last) of group 123
Logical message 1 (not last) of group 456
Logical message 2 (last)      of group 456
```

and a browse function is issued with MQGMO_ALL_MSGS_AVAILABLE, the first logical message of group 456 is returned, leaving the browse cursor on this logical message. If the second (last) message of group 123 now arrives:

```
Logical message 1 (not last) of group 123
Logical message 2 (last)     of group 123
Logical message 1 (not last) of group 456 <=== browse cursor
Logical message 2 (last)     of group 456
```

and the same browse-next function is issued, it is not noticed that group 123 is now complete, because the first message of this group is *before* the browse cursor.

In some cases (for example, if messages are retrieved destructively when the group is present in its entirety), you can use MQGMO_ALL_MSGS_AVAILABLE together with MQGMO_BROWSE_FIRST. Otherwise, you must repeat the browse scan to take note of newly-arrived messages that have been missed; just issuing MQGMO_WAIT together with MQGMO_BROWSE_NEXT and MQGMO_ALL_MSGS_AVAILABLE does not take account of them. (This also happens to higher-priority messages that might arrive after scanning the messages is complete.)

The next sections look at browsing examples that deal with unsegmented messages; segmented messages follow similar principles.

Browsing messages in groups

In this example, the application browses through each message on the queue, in logical order.

Messages on the queue might be grouped. For grouped messages, the application does not want to start processing any group until all the messages within it have arrived. MQGMO_ALL_MSGS_AVAILABLE is therefore specified for the first message in the group; for subsequent messages in the group, this option is unnecessary.

MQGMO_WAIT is used in this example. However, although the wait can be satisfied if a new group arrives, for the reasons in [“Browsing messages in logical order”](#) on page 262, it is not satisfied if the browse cursor has already passed the first logical message in a group, and the remaining messages now arrive. Nevertheless, waiting for a suitable interval ensures that the application does not constantly loop while waiting for new messages or segments.

MQGMO_LOGICAL_ORDER is used throughout, to ensure that the scan is in logical order. This contrasts with the destructive MQGET example, where because each group is being removed, MQGMO_LOGICAL_ORDER is not used when looking for the first (or only) message in a group.

It is assumed that the application's buffer is always large enough to hold the entire message, whether or not the message has been segmented. MQGMO_COMPLETE_MSG is therefore specified on each MQGET.

The following gives an example of browsing logical messages in a group:

```
/* Browse the first message in a group, or a message not in a group */
GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
             | MQGMO_ALL_MSGS_AVAILABLE | MQGMO_WAIT
MQGET GMO.MatchOptions = MQMO_MATCH_MSG_SEQ_NUMBER, MD.MsgSeqNumber = 1
/* Examine first or only message */
...
```



```

GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Examine each remaining message in the group */
  ...

```

The group is repeated until MQRC_NO_MSG_AVAILABLE is returned.

Browsing and retrieving destructively

In this example, the application browses each of the logical messages within a group, before deciding whether to retrieve that group destructively.

The first part of this example is similar to the previous one. However, in this case, having browsed an entire group, we decide to go back and retrieve it destructively.

As each group is removed in this example, MQGMO_LOGICAL_ORDER is not used when looking for the first or only message in a group.

The following gives an example of browsing and then retrieving destructively:

```

GMO.Options = MQGMO_BROWSE_NEXT | MQGMO_COMPLETE_MSG | MQGMO_LOGICAL_ORDER
              | MQGMO_ALL_MESSAGES_AVAILABLE | MQGMO_WAIT
do while ( GroupStatus == MQGS_MSG_IN_GROUP )
  MQGET
  /* Examine each remaining message in the group (or as many as
     necessary to decide whether to get it destructively) */
  ...

  if ( we want to retrieve the group destructively )

    if ( GroupStatus == ' ' )
      /* We retrieved an ungrouped message */
      GMO.Options = MQGMO_MSG_UNDER_CURSOR | MQGMO_SYNCPOINT
      MQGET GMO.MatchOptions = 0
      /* Process the message */
      ...

    else
      /* We retrieved one or more messages in a group. The browse cursor */
      /* will not normally be still on the first in the group, so we have */
      /* to match on the GroupId and MsgSeqNumber = 1. */
      /* Another way, which works for both grouped and ungrouped messages, */
      /* would be to remember the MsgId of the first message when it was */
      /* browsed, and match on that. */
      GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
      MQGET GMO.MatchOptions = MQMO_MATCH_GROUP_ID
                        | MQMO_MATCH_MSG_SEQ_NUMBER,
              (MQMD.GroupId = value already in the MD)
              MQMD.MsgSeqNumber = 1
      /* Process first or only message */
      ...

      GMO.Options = MQGMO_COMPLETE_MSG | MQGMO_SYNCPOINT
                    | MQGMO_LOGICAL_ORDER
      do while ( GroupStatus == MQGS_MSG_IN_GROUP )
        MQGET
        /* Process each remaining message in the group */
        ...

```

Avoiding repeated delivery of browsed messages

By using certain open options and get-message options, you can mark messages as having been browsed so that they are not retrieved again by the current or other cooperating applications. Messages can be unmarked explicitly or automatically to make them available again for browsing.

If you browse messages on a queue, you might retrieve them in a different order to the order in which you would retrieve them if you got them destructively. In particular, you can browse the same message multiple times, which is not possible if it is removed from the queue. To avoid this you can *mark* messages as they are browsed, and avoid retrieving marked messages. This is sometimes referred to as *browse with mark*. To mark browsed messages, use the get message option MQGMO_MARK_BROWSE_HANDLE, and to retrieve only messages that are not marked, use MQGMO_UNMARKED_BROWSE_MSG. If you use the combination of options MQGMO_BROWSE_FIRST, MQGMO_UNMARKED_BROWSE_MSG, and

MQGMO_MARK_BROWSE_HANDLE, and issue repeated MQGETs, you will retrieve each message on the queue in turn. This prevents repeated delivery of messages even though MQGMO_BROWSE_FIRST is used to ensure that messages are not skipped. This combination of options can be represented by the single constant MQGMO_BROWSE_HANDLE. When there are no messages on the queue that have not been browsed, MQRC_NO_MSG_AVAILABLE is returned.

If multiple applications are browsing the same queue, they can open the queue with the options MQOO_CO_OP and MQOO_BROWSE. The object handle returned by each MQOPEN is considered to be part of a cooperating group. Any message returned by an MQGET call specifying the option MQGMO_MARK_BROWSE_CO_OP is considered to be marked for this cooperating set of handles.

If a message has been marked for some time, it can be automatically unmarked by the queue manager and made available for browsing again. The queue manager attribute MsgMarkBrowseInterval gives the time in milliseconds for which a message is to remain marked for the cooperating set of handles. A MsgMarkBrowseInterval of -1 means that messages are never automatically unmarked.

When the single process or set of cooperative processes marking messages stop, any marked messages become unmarked.

Examples of cooperative browsing

You might run multiple copies of a dispatcher application to browse messages on a queue and initiate a consumer based on the content of each message. In each dispatcher, open the queue with MQOO_CO_OP. This indicates that the dispatchers are cooperating and will be aware of each other's marked messages. Each dispatcher then makes repeated MQGET calls, specifying the options MQGMO_BROWSE_FIRST, MQGMO_UNMARKED_BROWSE_MSG, and MQGMO_MARK_BROWSE_CO_OP (you can use the single constant MQGMO_BROWSE_CO_OP to represent this combination of options). Each dispatcher application then retrieves only those messages that have not already been marked by other cooperating dispatchers. The dispatcher initializes a consumer and passes the MsgToken returned by the MQGET to the consumer, which destructively gets the message from the queue. If the consumer backs out the MQGET of the message, then the message is available for one of the browsers to re-dispatch, because it is no longer marked. If the consumer does not do an MQGET on the message, then after the MsgMarkBrowseInterval has passed, the queue manager unmarks the message for the cooperating set of handles, and it can be re-dispatched.

Rather than multiple copies of the same dispatcher application, you might have a number of different dispatcher applications browsing the queue, each suitable for processing a subset of the messages on the queue. In each dispatcher, open the queue with MQOO_CO_OP. This indicates that the dispatchers are cooperating and will be aware of each other's marked messages.

- If the order of message processing for a single dispatcher is important, each dispatcher makes repeated MQGET calls, specifying the options MQGMO_BROWSE_FIRST, MQGMO_UNMARKED_BROWSE_MSG, and MQGMO_MARK_BROWSE_HANDLE (or MQGMO_BROWSE_HANDLE). If the browsed message is suitable for this dispatcher to process, it then makes an MQGET call specifying MQGMO_MATCH_MSG_TOKEN, MQGMO_MARK_BROWSE_CO_OP, and the MsgToken returned by the previous MQGET call. If the call succeeds, the dispatcher initializes the consumer, passing the MsgToken to it.
- If the order of message processing is not important and the dispatcher is expected to process most of the messages it encounters, use the options MQGMO_BROWSE_FIRST, MQGMO_UNMARKED_BROWSE_MSG, and MQGMO_MARK_BROWSE_CO_OP (or MQGMO_BROWSE_CO_OP). If the dispatcher browses a message it cannot process, it unmarks the message by calling MQGET with the option MQGMO_MATCH_MSG_TOKEN, MQGMO_UNMARK_BROWSE_CO_OP, and the MsgToken returned previously.

Some cases where the MQGET call fails

If certain attributes of a queue are changed using the FORCE option on a command between issuing an MQOPEN and an MQGET call, the MQGET call fails and returns the MQRC_OBJECT_CHANGED reason code.

The queue manager marks the object handle as being no longer valid. This also happens if the changes apply to any queue to which the queue name resolves. The attributes that affect the handle in this way are listed in the description of the MQOPEN call in [MQOPEN](#). If your call returns the MQRC_OBJECT_CHANGED reason code, close the queue, reopen it, then try to get a message again.

If get operations are inhibited for a queue from which you are attempting to get messages (or any queue to which the queue name resolves), the MQGET call fails and returns the MQRC_GET_INHIBITED reason code. This happens even if you are using the MQGET call for browsing. You might be able to get a message successfully if you attempt the MQGET call at a later time, if the design of the application is such that other programs change the attributes of queues regularly.

If a dynamic queue (either temporary or permanent) has been deleted, MQGET calls using a previously-acquired object handle fail and return the MQRC_Q_DELETED reason code.

Writing publish/subscribe applications

Start writing publish/subscribe WebSphere MQ applications.

For an overview of publish/subscribe concepts, see [Introduction to WebSphere MQ publish/subscribe messaging](#).

See the following topics for information on writing different types of publish/subscribe applications:

- [“Writing publisher applications” on page 267](#)
- [“Writing subscriber applications” on page 273](#)
- [“Publish/subscribe lifecycles” on page 289](#)
- [“Publish/subscribe message properties” on page 294](#)
- [“Message ordering” on page 295](#)
- [“Intercepting publications” on page 296](#)
- [“Publishing options” on page 304](#)
- [“Subscription options” on page 304](#)

Related concepts

[“Application development concepts” on page 8](#)

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Sample WebSphere MQ programs” on page 91](#)

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

[“Writing a queuing application” on page 186](#)

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Writing client applications” on page 335](#)

What you need to know to write client applications on WebSphere MQ.

[“Using Web services in WebSphere MQ” on page 904](#)

You can develop IBM WebSphere MQ applications for Web services using the IBM WebSphere MQ transport for SOAP or the IBM WebSphere MQ bridge for HTTP.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

[“Handling program errors” on page 523](#)

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Writing publisher applications

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Writing a simple WebSphere MQ publisher application is just like writing a WebSphere MQ point to point application that puts messages to a queue ([Table 41 on page 267](#)). The difference is you MQPUT messages to a topic, not to a queue.

<i>Table 41. Point to point versus publish/subscribe WebSphere MQ program pattern.</i>		
Step	Point to point MQ Call	Publish MQ Call
Connect to a queue manager	MQCONN	MQCONN
Open queue	MQOPEN	
Open topic		MQOPEN
Put message(s)	MQPUT	MQPUT
Close topic		MQCLOSE
Close queue	MQCLOSE	
Disconnect from queue manager	MQDISC	MQDISC

To make that concrete, there are two examples of applications to publish stock prices. In the first example ([“Example 1: Publisher to a fixed topic” on page 268](#)), that is modeled closely on putting messages to a queue, the administrator creates a topic definition in a similar way to creating a queue. The programmer codes MQPUT to write messages to the topic instead of writing them to a queue. In the second example ([“Example 2: Publisher to a variable topic” on page 270](#)), the pattern of interaction of the program with WebSphere MQ is similar. The difference is the programmer provides the topic to which the message is written, rather than the administrator. In practice this typically means that the topic string is content defined, or provided by another source, such as human input through a browser.

Related concepts

[“Writing subscriber applications” on page 273](#)

Get started with writing subscriber applications by studying three examples: a WebSphere MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

Related reference

[DEFINE TOPIC](#)

[DISPLAY TOPIC](#)

[DISPLAY TPSTATUS](#)

Example 1: Publisher to a fixed topic

A WebSphere MQ program to illustrate publishing to an administratively defined topic.

Note: The compact coding style is intended for readability not production use.

See the output in [Figure 38 on page 268](#)

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
    char    topicNameDefault[]    = "IBMSTOCKPRICE";
    char    publicationDefault[]  = "129";
    MQCHAR48 qmName = "";

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ  Hobj  = MQHO_NONE;           /* object handle sub queue */
    MQLONG  CompCode = MQCC_OK;          /* completion code */
    MQLONG  Reason = MQRC_NONE;          /* reason code */
    MQOD    td = {MQOD_DEFAULT};         /* Object descriptor */
    MQMD    md = {MQMD_DEFAULT};         /* Message Descriptor */
    MQPMO    pmo = {MQPMO_DEFAULT};      /* put message options */
    MQCHAR    resTopicStr[151];          /* Returned value of topic string */
    char *    topicName = topicNameDefault;
    char *    publication = publicationDefault;
    memset    (resTopicStr, 0 , sizeof(resTopicStr));

    switch(argc){
        default:
            publication = argv[2];
        case(2):
            topicName = argv[1];
        case(1):
            printf("Optional parameters: TopicObject Publication\n");
    }
    do {
        MQCONN(qmName, &Hconn, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        td.ObjectType = MQOT_TOPIC;      /* Object is a topic */
        td.Version = MQOD_VERSION_4;     /* Descriptor needs to be V4 */
        strncpy(td.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
        td.ResObjectString.VSPtr = resTopicStr;
        td.ResObjectString.VSBufSize = sizeof(resTopicStr)-1;
        MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF_QUIESCING, &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        pmo.Options = MQPMO_FAIL_IF_QUIESCING | MQPMO_RETAIN;
        MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode,
        &Reason);
        if (CompCode != MQCC_OK) break;
        MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQDISC(&Hconn, &CompCode, &Reason);
    } while (0);
    if (CompCode == MQCC_OK)
        printf("Published \"%s\" using topic \"%s\" to topic string \"%s\"\n",
            publication, td.ObjectName, resTopicStr);
    printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
```

Figure 37. Simple WebSphere MQ publisher to a fixed topic.

```
X:\Publish1\Debug>PublishStock
Optional parameters: TopicObject Publication
Published "129" using topic "IBMSTOCKPRICE" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

X:\Publish1\Debug>PublishStock IBMSTOCKPRICE 155
Optional parameters: TopicObject Publication
Published "155" using topic "IBMSTOCKPRICE" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0
```

Figure 38. Sample output from first publisher example

The following selected lines of code illustrate aspects of writing a publisher application for WebSphere MQ.

```
char topicNameDefault[] = "IBMSTOCKPRICE";
```

A default topic name is defined in the program. You can override it by providing the name of a different topic object as the first argument to the program.

```
MQCHAR resTopicStr[151];
```

`resTopicStr` is pointed at by `td.ResObjectString.VSPtr` and is used by `MQOPEN` to return the resolved topic string. Make the length of `resTopicStr` one larger than the length passed in `td.ResObjectString.VSBufSize` to give space for null termination.

```
memset (resTopicStr, 0, sizeof(resTopicStr));
```

Initialize `resTopicStr` to nulls to ensure the resolved topic string returned in an `MQCHARV` is null terminated.

```
td.ObjectType = MQOT_TOPIC
```

There is a new type of object for publish/subscribe: the *topic object*.

```
td.Version = MQOD_VERSION_4;
```

To use the new type of object, you must use at least *version 4* of the object descriptor.

```
strncpy(td.ObjectName, topicName, MQ_OBJECT_NAME_LENGTH);
```

The `topicName` is the name of a topic object, sometimes called an administrative topic object. In the example the topic object needs to be created beforehand, using WebSphere MQ Explorer or this MQSC command,

```
DEFINE TOPIC(IBMSTOCKPRICE) TOPICSTR(NYSE/IBM/PRICE) REPLACE;
```

```
td.ResObjectString.VSPtr = resTopicStr;
```

The resolved topic string is echoed in the final `printf` in the program. Set up the `MQCHARV ResObjectString` structure for WebSphere MQ to return the resolved string back to the program.

```
MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF_QUIESCING, &Hobj, &CompCode, &Reason);
```

Open the topic for output; just like opening a queue for output.

```
pmo.Options = MQPMO_FAIL_IF_QUIESCING | MQPMO_RETAIN;
```

You want new subscribers to be able receive the publication, and by specifying `MQPMO_RETAIN` in the publisher, when you start a subscriber it receives the latest publication, published before the subscriber started, as its first matching publication. The alternative is to provide subscribers with publications published only after the subscriber started. Additionally a subscriber has the option to decline to receive a retained publication by specifying `MQSO_NEW_PUBLICATIONS_ONLY` in its subscription.

```
MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode, &Reason);
```

Add 1 to the length of the string passed to `MQPUT` to pass the null termination character to WebSphere MQ as part of the message buffer.

What does the first example demonstrate? The example imitates as closely as possible the tried and tested traditional pattern for writing point to point WebSphere MQ programs. An important feature of the WebSphere MQ programming pattern is that the programmer is not concerned where messages are sent. The task of the programmer is to connect to a queue manager, and pass it the messages that are to be distributed to recipients. In the point-to-point paradigm, the programmer opens a queue (probably an alias queue) that the administrator has configured. The alias queue routes messages to a target queue, either on the local queue manager, or to a remote queue manager. While the messages are waiting to be delivered, they are stored on queues somewhere between the source and the destination.

In the publish/subscribe pattern, instead of opening a queue, the programmer opens a topic. In our example, the topic is associated with a topic string by an administrator. The queue manager forwards the publication, using queues, to local or remote subscribers that have subscriptions that match the topic string of the publication. If publications are retained the queue manager keeps the latest copy of the publication, even if it has no subscribers now. The retained publication is available to forward to

future subscribers. The publisher application plays no part in selecting or routing the publication to a destination; its task is to create and put publications to the topics defined by the administrator.

This fixed topic example is atypical of many publish/subscribe applications: it is static. It requires an administrator to define the topic strings and change the topics that are published on. Commonly publish/subscribe applications need to know some or all the topic tree. Perhaps topics change frequently, or perhaps although the topics do not change much, the number of topic combinations is large and it is too onerous for an administrator to define a topic node for every topic string that might need to be published on. Perhaps topic strings are not known in advance of publication; a publisher application might use information from the publication content to specify a topic string, or it might have information about topic strings to publish on from another source, such as human input from a browser. To cater for more dynamic styles of publishing, the next example shows how to create topics dynamically, as part of the publisher application.

Topics couple publishers and subscribers together. Designing the rules, or architecture, for naming topics, and organizing them in topic trees is an important step in developing a publish/subscribe solution. Look carefully at the extent to which organization of the topic tree binds of publisher and subscriber programs together, and binds them to the content of the topic tree. Ask yourself the question whether changes in the topic tree affect publisher and subscriber applications, and how you can minimize the effect. Built into the architecture of the WebSphere MQ publish/subscribe model is the notion of an administrative topic object that provides the root part, or root subtree, of a topic. The topic object gives you the option of defining the root part of the topic tree administratively that simplifies application programming and operations, and consequently improves maintainability. For example, if you are deploying multiple publish/subscribe applications that have isolated topic trees, then by administratively defining the root part of the topic tree, you can guarantee the isolation of topic trees, even if there is no consistency in the topic naming conventions adopted by the different applications.

In practice, publisher applications cover a spectrum from solely using fixed topics, as in this example, and variable topics, as in the next. [“Example 2: Publisher to a variable topic” on page 270](#) also demonstrates combining the use of topics and topic strings.

Related concepts

[“Example 2: Publisher to a variable topic” on page 270](#)

A Websphere MQ program to illustrate publishing to a programmatically defined topic.

[“Writing subscriber applications” on page 273](#)

Get started with writing subscriber applications by studying three examples: a WebSphere MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queuing, and finally an example that uses both queuing and subscriptions.

Example 2: Publisher to a variable topic

A Websphere MQ program to illustrate publishing to a programmatically defined topic.

Note: The compact coding style is intended for readability not production use.

See the output in [Figure 40](#) on page 272.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
    char    topicNameDefault[]    = "STOCKS";
    char    topicStringDefault[]  = "IBM/PRICE";
    char    publicationDefault[]  = "130";
    MQCHAR48 qmName = "";

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ  Hobj  = MQHO_NONE;           /* object handle sub queue */
    MQLONG  CompCode = MQCC_OK;           /* completion code */
    MQLONG  Reason  = MQRC_NONE;          /* reason code */
    MQOD    td = {MQOD_DEFAULT};          /* Object descriptor */
    MQMD    md = {MQMD_DEFAULT};          /* Message Descriptor */
    MQPMO    pmo = {MQPMO_DEFAULT};       /* put message options */
    MQCHAR    resTopicStr[151];           /* Returned value of topic string */
    char *    topicName = topicNameDefault;
    char *    topicString = topicStringDefault;
    char *    publication = publicationDefault;
    memset    (resTopicStr, 0 , sizeof(resTopicStr));

    switch(argc){
        /* Replace defaults with args if provided */
        default:
            publication = argv[3];
        case(3):
            topicString = argv[2];
        case(2):
            if (strcmp(argv[1],"/")) /* "/" invalid = No topic object */
                topicName = argv[1];
            else
                *topicName = '\0';
        case(1):
            printf("Provide parameters: TopicObject TopicString Publication\n");
    }

    printf("Publish \"%s\" to topic \"%-48s\" and topic string \"%s\"\\n", publication,
topicName, topicString);
    do {
        MQCONN(qmName, &Hconn, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        td.ObjectType = MQOT_TOPIC; /* Object is a topic */
        td.Version = MQOD_VERSION_4; /* Descriptor needs to be V4 */
        strncpy(td.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
        td.ObjectString.VSPtr = topicString;
        td.ObjectString.VSLength = (MQLONG)strlen(topicString);
        td.ResObjectString.VSPtr = resTopicStr;
        td.ResObjectString.VSBufSize = sizeof(resTopicStr)-1;
        MQOPEN(Hconn, &td, MQOO_OUTPUT | MQOO_FAIL_IF_QUIESCING, &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        pmo.Options = MQPMO_FAIL_IF_QUIESCING | MQPMO_RETAIN;
        MQPUT(Hconn, Hobj, &md, &pmo, (MQLONG)strlen(publication)+1, publication, &CompCode,
&Reason);
        if (CompCode != MQCC_OK) break;
        MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQDISC(&Hconn, &CompCode, &Reason);
    } while (0);
    if (CompCode == MQCC_OK)
        printf("Published \"%s\" to topic string \"%s\"\\n", publication, resTopicStr);
    printf("Completion code %d and Return code %d\\n", CompCode, Reason);
}
```

Figure 39. Simple WebSphere MQ publisher to a variable topic.

```

X:\Publish2\Debug>PublishStock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

X:\Publish2\Debug>PublishStock / NYSE/IBM/PRICE 131
Provide parameters: TopicObject TopicString Publication
Publish "131" to topic "" and topic string "NYSE/IBM/PRICE"
Published "131" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

```

Figure 40. Sample output from second publisher example

There are a few points to note about this example.

char topicNameDefault[] = "STOCKS";

The default topic name STOCKS defines part of the topic string. You can override this topic name by providing it as the first argument to the program, or eliminate the use of the topic name by supplying / as the first parameter.

char topicString[101] = "IBM/PRICE";

IBM/PRICE is the default topic string. You can override this topic string by providing it as the second argument to the program.

The queue manager combines the topic string provided by the STOCKS topic object, "NYSE", with the topic string provided by the program "IBM/PRICE" and inserts a "/" between the two topic strings. The result is the resolved topic string "NYSE/IBM/PRICE". The resulting topic string is the same as the one defined in the IBMSTOCKPRICE topic object, and has precisely the same effect.

The administrative topic object associated with the resolved topic string is not necessarily the same topic object as passed to MQOPEN by the publisher. WebSphere MQ uses the tree implicit in the resolved topic string to work out which administrative topic object defines the attributes associated with the publication.

Suppose there are two topic objects A and B, and A defines topic "a", and B defines topic "a/b" (Figure 41 on page 273). If the publisher program refers to topic object A and provides topic string "b", resolving the topic to the topic string "a/b", then the publication inherits its properties from topic object B because the topic matches the topic string "a/b" defined for B.

if (strcmp(argv[1],"/"))

argv[1] is the optionally provided topicName. "/" is invalid as a topic name; here it signifies that there is no topic name, and the topic string is provided entirely by the program. The output in Figure 40 on page 272 shows the whole topic string being supplied dynamically by the program.

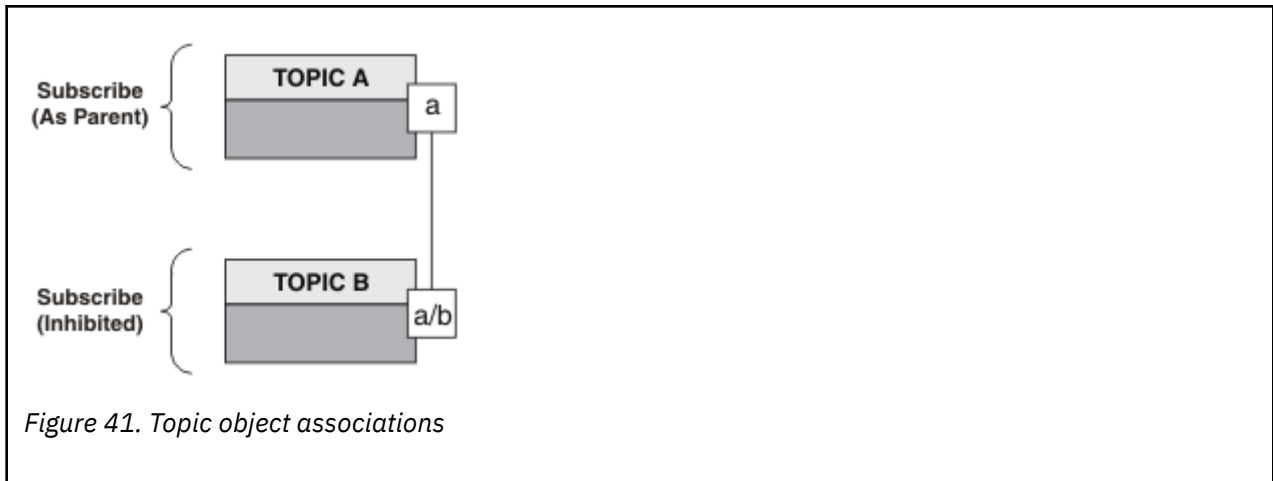
strncpy(td.ObjectName, topicName, MQ_OBJECT_NAME_LENGTH);

For the default case, the optional topicName needs to be created beforehand, using WebSphere MQ Explorer or this MQSC command:

```
DEFINE TOPIC(STOCKS) TOPICSTR(NYSE) REPLACE;
```

td.ObjectString.VSPtr = topicString;

The topic string is a MQCHARV field in the topic descriptor



What does the second example demonstrate? Although the code is very similar to the first example - effectively there are only two lines difference - the result is a significantly different program to the first. The programmer controls the destinations to which publications are sent. In conjunction with minimal administrator input used to design subscriber applications, no topics or queues need to be predefined to route publications from publishers to subscribers.

In the point-to-point messaging paradigm, queues have to be defined before messages are able to flow. For publish/subscribe, they do not, although WebSphere MQ implements publish/subscribe using its underlying queuing system; the benefits of guaranteed delivery, transactionality and loose coupling associated with messaging and queueing are inherited by publish/subscribe applications.

A designer has to decide whether publisher, and subscriber, programs are to be aware of the underlying topic tree or not, and also whether subscriber programs are aware of queueing or not. Study the subscriber example applications next. They are designed to be used with the publisher examples, typically publishing and subscribing to NYSE/IBM/PRICE.

Related concepts

[“Example 1: Publisher to a fixed topic” on page 268](#)

A WebSphere MQ program to illustrate publishing to an administratively defined topic.

[“Writing subscriber applications” on page 273](#)

Get started with writing subscriber applications by studying three examples: a WebSphere MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queueing, and finally an example that uses both queueing and subscriptions.

Writing subscriber applications

Get started with writing subscriber applications by studying three examples: a WebSphere MQ application consuming messages from a queue, an application that creates a subscription and requires no knowledge of queueing, and finally an example that uses both queueing and subscriptions.

In [Table 42 on page 274](#) the three styles of consumer or subscriber are listed, together with the sequences of WebSphere MQ function calls that characterize them.

1. The first style, MQ Publication Consumer, is identical to a point to point MQ program that only does MQGET. The application has no knowledge that it is consuming publications - it is simply reading messages from a queue. The subscription that causes publications to get routed to the queue is created administratively using WebSphere MQ Explorer or a command.
2. The second style is the preferred pattern for most subscriber applications. The subscriber application creates the subscription, and then gets publications. The queue management is all performed by the queue manager.
3. In the third style, the subscriber application elects to open and close the underlying queue that is used for publications as well as issue subscriptions to fill the queue with publications.

One way to understand these styles is to study the example C programs listed in [Table 42 on page 274](#) for each of the styles. The examples are designed to be run in conjunction with the publisher example found in [“Writing publisher applications” on page 267](#).

<i>Table 42. Point to point vs. subscribe WebSphere MQ program patterns.</i>				
Step	MQ message consumer	“Example 1: MQ Publication consumer” on page 274	“Example 2: Managed MQ subscriber” on page 276	“Example 3: Unmanaged MQ subscriber” on page 281
Connect to a queue manager	MQCONN	MQCONN	MQCONN	MQCONN
Open queue	MQOPEN	MQOPEN		MQOPEN
Subscribe			MQSUB	MQSUB
Get message(s)	MQGET	MQGET	MQGET	MQGET
Close queue	MQCLOSE	MQCLOSE	(MQCLOSE)	MQCLOSE
Close subscription			MQCLOSE	MQCLOSE
Disconnect from queue manager	MQDISC	MQDISC	MQDISC	MQDISC

Using MQCLOSE is always optional, either to release resources, pass MQCLOSE options, or just for symmetry with MQOPEN. Since you are unlikely to need to specify the MQCLOSE options when the subscription queue is closed in the Managed MQ subscriber case, and the symmetry argument is not relevant, the subscription queue is not explicitly closed in [Example 2: Managed MQ subscriber](#).

Another way to understand publish/subscribe application patterns is to look at the interactions between the different entities involved. Lifeline, or UML sequence diagrams are a good way to study interactions. Three lifeline examples are described in [“Publish/subscribe lifecycles” on page 289](#).

Example 1: MQ Publication consumer

The MQ Publication consumer is an IBM WebSphere MQ message consumer that does not subscribe to topics itself.

To create the subscription and publication queue for this example run the following commands, or define the objects using WebSphere MQ Explorer.

```
DEFINE QLOCAL(STOCKTICKER) REPLACE;
DEFINE SUB(IBMSTOCKPRICESUB) DEST(STOCKTICKER) TOPICOBJ(IBMSTOCKPRICE) REPLACE;
```

The IBMSTOCKPRICESUB subscription references the IBMSTOCK topic object created for the publisher example and the local queue STOCKTICKER. The topic object IBMSTOCK defines the topic string that is used in the subscription, NYSE/IBM/PRICE. Note that the topic object and the queue used to receive publications need to be defined before the subscription is created.

There are a number of valuable facets to the MQ publication consumer pattern:

1. Multiprocessing: sharing out of the work of reading publications. The publications all go onto the single queue associated with the subscription topic. Multiple consumers can open the queue using MQOO_INPUT_SHARED.
2. Centrally managed subscriptions. Applications do not construct their own subscription topics or subscriptions; the administrator is responsible for where publications are sent.
3. Subscription concentration: multiple different subscriptions can be sent to a single queue.
4. Subscription durability: the queue receives all publications whether or not consumers are active.
5. Migration and coexistence: the consumer code works equally well for a point-to-point and a publish/subscribe scenario.

The subscription creates a relationship between the topic string NYSE/IBM/PRICE and the queue STOCKTICKER. Publications, including any currently retained publication, are forwarded to STOCKTICKER from the moment the subscription is created.

An administratively created subscription can be managed or unmanaged. A managed subscription takes effect as soon as it has been created, just like an unmanaged subscription. Not all the pattern facets are available to a managed subscription. See [“Example 3: Unmanaged MQ subscriber”](#) on page 281

Note: The compact coding style is intended for readability not production use.

The results are shown in [Figure 43](#) on page 276.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
int main(int argc, char **argv)
{
    MQCHAR    publicationBuffer[101];
    MQCHAR48  subscriptionQueueDefault = "STOCKTICKER";
    MQCHAR48  qmName = ""; /* Use default queue manager */

    MQHCONN   Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ    Hobj = MQHO_NONE; /* object handle sub queue */
    MQLONG    CompCode = MQCC_OK; /* completion code */
    MQLONG    Reason = MQRC_NONE; /* reason code */
    MQLONG    messlen = 0;
    MQOD      od = {MQOD_DEFAULT}; /* Unmanaged subscription queue */
    MQMD      md = {MQMD_DEFAULT}; /* Message Descriptor */
    MQGMO     gmo = {MQGMO_DEFAULT}; /* Get message options */
    char *    publication=publicationBuffer;
    char *    subscriptionQueue = subscriptionQueueDefault;

    switch(argc){ /* Replace defaults with args if provided */
    default:
        subscriptionQueue = argv[1]
    case(1):
        printf("Optional parameter: subscriptionQueue\n");
    }

    do {
        MQCONN(qmName, &Hconn, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        strncpy(od.ObjectName, subscriptionQueue, MQ_Q_NAME_LENGTH);
        MQOPEN(Hconn, &od, MQOO_INPUT_AS_Q_DEF | MQOO_FAIL_IF_QUIESCING, &Hobj, &CompCode,
&Reason);
        if (CompCode != MQCC_OK) break;
        gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
        gmo.WaitInterval = 10000;
        printf("Waiting %d seconds for publications from %s\n", gmo.WaitInterval/1000,
subscriptionQueue);
        do {
            memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
            memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
            md.Encoding = MQENC_NATIVE;
            md.CodedCharSetId = MQCCSI_Q_MGR;
            memset(publication, 0, sizeof(publicationBuffer));
            MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1, publication, &messlen,
&CompCode, &Reason);
            if (Reason == MQRC_NONE)
                printf("Received publication \"%s\"\n", publication);
        }
        while (CompCode == MQCC_OK);
        if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
        MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
        MQDISC(&Hconn, &CompCode, &Reason);
    } while (0);
    printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
```

Figure 42. MQ publication consumer.

```
X:\Subscribe1\Debug>Subscribe1
Optional parameter: subscriptionQueue
Waiting 10 seconds for publications from STOCKTICKER
Received publication "129"
Completion code 0 and Return code 0
```

Figure 43. Output from MQ publication consumer

There are a couple of standard WebSphere MQ C language programming tips to be aware of:

memset(publication, 0, sizeof(publicationBuffer));

Ensure the message has a trailing null for easy formatting using printf. The publisher example includes the trailing null in the message buffer passed to MQPUT by adding 1 to strlen(publication). Setting MQCHAR buffers to null is good programming style for IBM WebSphere MQ C programs that use the buffers to store strings, ensuring a null follows an array of characters that does not completely fill the buffer.

MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1, publication, &messlen, &CompCode, &Reason);

Reserve one null at the end of the message buffer to ensure the returned message has trailing null in case "if (messlen == strlen(publication));" is true. This tip complements the preceding one, and ensures that there is at least one null in publicationBuffer that is not overwritten by the contents of publication.

Related concepts

[“Example 2: Managed MQ subscriber” on page 276](#)

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

[“Example 3: Unmanaged MQ subscriber” on page 281](#)

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

[“Writing publisher applications” on page 267](#)

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Example 2: Managed MQ subscriber

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

This simplest kind of managed subscriber typically uses a *non-durable* subscription. The example focuses on a non-durable subscription. The subscription lasts only as long as the lifetime of the subscription handle from MQSUB. Any publications that match the topic string during the lifetime of the subscription are sent to the subscription queue (and possibly a retained publication if the flag MQSO_NEW_PUBLICATIONS_ONLY is not set or defaulted, an earlier publication matching the topic string was retained, and the publication was persistent or the queue manager has not terminated, since the publication was created).

You can also use a *durable* subscription with this pattern. Typically if a managed durable subscription is used it is done for reliability reasons, rather than to establish a subscription that, without any errors occurring, would outlive the subscriber. For more information about different lifecycles associated with managed, unmanaged, durable and non-durable subscriptions see the related topics section.

Durable subscriptions are often associated with persistent publications, and non-durable subscriptions with non-persistent publications, but there is no necessary relationship between subscription durability and publication persistence. All four combinations of persistence and durability are possible.

For the managed non-durable case considered, the queue manager creates a subscription queue that is purged and deleted when the queue is closed. The publications are removed from the queue when the non-durable subscription is closed.

The valuable facets of the managed non-durable pattern exemplified by this code are as follows:

1. On demand subscription: the subscription topic string is dynamic. It is provided by the application when it runs.
2. Self managing queue: the subscription queue is self defining and managing.
3. Self managing subscription lifecycle: *non-durable* subscriptions only exist for the duration of the subscriber application.
 - If you define a *durable* managed subscription, then it results in a permanent subscription queue and publications continue to be stored on it with no subscriber programs being active. The queue manager deletes the queue (and clears any unretrieved publications from it) only after the application or administrator has chosen to delete the subscription. The subscription can be deleted using an administrative command, or by closing the subscription with the MQCO_REMOVE_SUB option.
 - Consider setting SubExpiry for durable subscriptions so that publications cease to be sent to the queue and the subscriber can consume any remaining publications before removing the subscription and causing the queue manager to delete the queue and any remaining publications on it.
4. Flexible topic string deployment: Subscription topic management is simplified by defining the root part of the subscription using an administratively defined topic. The root part of the topic tree is then hidden from the application. By hiding the root part an application can be deployed without the application inadvertently creating a topic tree that overlaps with another topic tree created by another instance, or another application.
5. Administered topics: by using a topic string in which the first part matches an administratively defined topic object, publications are managed according to the attributes of the topic object.
 - For example, if the first part of the topic string matches the topic string associated with a clustered topic object, then the subscription can receive publications from other members of the cluster
 - The selective matching of administratively defined topic objects and programmatically defined subscriptions enables you to combine the benefits of both. The administrator provides attributes for topics, and the programmer dynamically defines "sub-topics" without being concerned about the management of topics.
 - It is the resultant topic string which is used to match the topic object that provides the attributes associated with the topic, and not necessarily the topic object named in sd.Objectname, although they typically turn out to be one and the same. See [“Example 2: Publisher to a variable topic” on page 270](#).

By making the subscription durable in the example, publications continue to be sent to the subscription queue after the subscriber has closed the subscription with the MQCO_KEEP_SUB option. The queue continues to receive publications when the subscriber is not active. You can override this behavior by creating the subscription with the MQSO_PUBLICATIONS_ON_REQUEST option and using MQSUBRQ to request the retained publication.

The subscription can be resumed later by opening the subscription with the MQCO_RESUME option.

You can use the queue handle, Hobj, returned by MQSUB in a number of ways. The queue handle is used in the example to inquire on the name of the subscription queue. Managed queues are opened using the default model queues SYSTEM.NDURABLE.MODEL.QUEUE or SYSTEM.DURABLE.MODEL.QUEUE. You can override the defaults by providing your own durable and non-durable model queues on a topic by topic basis as properties of the topic object associated with the subscription.

Regardless of the attributes inherited from the model queues, you cannot reuse a managed queue handle to create an additional subscription. Nor can you obtain another handle for the managed queue by opening the managed queue a second time using the returned queue name. The queue behaves as if it has been opened for exclusive input.

Unmanaged queues are more flexible than managed queues. You can, for example share unmanaged queues, or define multiple subscriptions on the one queue. The next example, [“Example 3: Unmanaged MQ subscriber” on page 281](#), demonstrates how to combine subscriptions with an unmanaged subscription queue.

Note: The compact coding style is intended for readability not production use.

The results are shown in [Figure 46 on page 279](#).

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>

void inquireQname(MQHCONN HConn, MQHOBJ Hobj, MQCHAR48 qName);

int main(int argc, char **argv)
{
    MQCHAR48 topicNameDefault = "STOCKS";
    char topicStringDefault[] = "IBM/PRICE";
    MQCHAR48 qmName = ""; /* Use default queue manager */
    MQCHAR48 qName = ""; /* Allocate to query queue name */
    char publicationBuffer[101]; /* Allocate to receive messages */
    char resTopicStrBuffer[151]; /* Allocate to resolve topic string */

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ Hobj = MQHO_NONE; /* publication queue handle */
    MQHOBJ Hsub = MQSO_NONE; /* subscription handle */
    MQLONG CompCode = MQCC_OK; /* completion code */
    MQLONG Reason = MQRC_NONE; /* reason code */
    MQLONG messlen = 0;
    MQSD sd = {MQSD_DEFAULT}; /* Subscription Descriptor */
    MQMD md = {MQMD_DEFAULT}; /* Message Descriptor */
    MQGMO gmo = {MQGMO_DEFAULT}; /* get message options */

    char * topicName = topicNameDefault;
    char * topicString = topicStringDefault;
    char * publication = publicationBuffer;
    char * resTopicStr = resTopicStrBuffer;
    memset(resTopicStr, 0, sizeof(resTopicStrBuffer));

    switch(argc){ /* Replace defaults with args if provided */
    default:
        topicString = argv[2];
    case(2):
        if (strcmp(argv[1],"/")) /* "/" invalid = No topic object */
            topicName = argv[1];
        else
            *topicName = '\0';
    case(1):
        printf("Optional parameters: topicName, topicString\nValues \"%s\" \"%s\"\n",
            topicName, topicString);
    }
}
```

Figure 44. Managed MQ subscriber - part 1: declarations and parameter handling.

There are some additional comments to make about the declarations in this example.

MQHOBJ Hobj = MQHO_NONE;

You cannot explicitly open a non-durable managed subscription queue to receive publications, but you do need to allocate storage for the object handle the queue manager returns when it opens the queue for you. It is important to initialize the handle to MQHO_OBJECT. This indicates to the queue manager that it needs to return a queue handle to the subscription queue.

MQSD sd = {MQSD_DEFAULT};

The new subscription descriptor, used in MQSUB.

MQCHAR48 qName;

Although the example doesn't require knowledge of the subscription queue, the example does inquire the name of the subscription queue - the MQINQ binding is a little awkward in the C language, so you might find this part of the example useful to study.

```

do {
    MQCONN(qmName, &Hconn, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    strncpy(sd.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
    sd.ObjectString.VSPtr = topicString;
    sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    sd.Options = MQSO_CREATE | MQSO_MANAGED | MQSO_NON_DURABLE | MQSO_FAIL_IF QUIESCING ;
    sd.ResObjectString.VSPtr = resTopicStr;
    sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;
    MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
    gmo.WaitInterval = 10000;
    inquireQname(Hconn, Hobj, qName);
    printf("Waiting %d seconds for publications matching \"%s\" from \"%-0.48s\\\"\\n",
        gmo.WaitInterval/1000, resTopicStr, qName);
    do {
        memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
        memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
        md.Encoding = MQENC_NATIVE;
        md.CodedCharSetId = MQCCSI_Q_MGR;
        memset(publicationBuffer, 0, sizeof(publicationBuffer));
        MQGET(Hconn, Hobj, &md, &gmo, sizeof(publicationBuffer)-1,
            publication, &messlen, &CompCode, &Reason);
        if (Reason == MQRC_NONE)
            printf("Received publication \"%s\\\"\\n", publication);
    }
    while (CompCode == MQCC_OK);
    if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
    MQCLOSE(Hconn, &Hsub, MQCO_REMOVE_SUB, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQDISC(&Hconn, &CompCode, &Reason);
} while (0);
printf("Completion code %d and Return code %d\\n", CompCode, Reason);
return;
}

void inquireQname(MQHCONN Hconn, MQHOBJ Hobj, MQCHAR48 qName) {
#define _selectors 1
#define _intAttrs 1

    MQLONG select[_selectors] = {MQCA_Q_NAME}; /* Array of attribute selectors */
    MQLONG intAttrs[_intAttrs]; /* Array of integer attributes */
    MQLONG CompCode, Reason;
    MQINQ(Hconn, Hobj, _selectors, select, _intAttrs, intAttrs, MQ_Q_NAME_LENGTH, qName,
        &CompCode, &Reason);
    if (CompCode != MQCC_OK) {
        printf("MQINQ failed with Condition code %d and Reason %d\\n", CompCode, Reason);
        strcpy(qName, "unknown queue");
    }
    return;
}

```

Figure 45. Managed MQ subscriber - part 2: code body.

```

W:\Subscribe2\Debug>solution2
Optional parameters: topicName, topicString
Values "STOCKS" "IBM/PRICE"
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from
"SYSTEM.MANAGED.NDURABLE.48A0AC7403300020"
Received publication "150"
Completion code 0 and Return code 0

W:\Subscribe2\Debug>solution2 / NYSE/IBM/PRICE
Optional parameters: topicName, topicString
Values "" "NYSE/IBM/PRICE"
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from
"SYSTEM.MANAGED.NDURABLE.48A0AC7403310020"
Received publication "150"
Completion code 0 and Return code 0

```

Figure 46. Output from managed MQ subscriber

There are some additional comments to make about the code in this example.

strncpy(sd.ObjectName, topicName, MQ_Q_NAME_LENGTH);

If topicName is null or blank (*default value*), the topic name is not used to compute the resolved topic string.

sd.ObjectString.VSPtr = topicString;

Rather than solely use a predefined topic object, in this example the programmer provides a topic object and a topic string, that are combined by MQSUB. Notice the topic string is a MQCHARV structure.

sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;

An alternative to setting the length of a MQCHARV field.

**sd.Options = MQSO_CREATE | MQSO_MANAGED | MQSO_NON_DURABLE |
MQSO_FAIL_IF QUIESCING;**

After defining the topic string, the sd.Options flags need the most careful attention. There are many options, the example specifies only the most commonly used ones; the others are left to default.

1. As the subscription is *non-durable*, that is, it has a lifetime of the open subscription in the application, set the MQSO_CREATE flag. You can also set the (*default*) MQSO_NON_DURABLE flag for readability.
2. Complementing MQSO_CREATE is MQSO_RESUME. Both flags can be set together; the queue manager either creates a new subscription or resumes an existing subscription, whichever is appropriate. However, if you do specify MQSO_RESUME you must also initialize the MQCHARV structure for sd.SubName, even if there is no subscription to resume. Failure to initialize SubName results in a return code of 2440: MQRC_SUB_NAME_ERROR from MQSUB.

Note: MQSO_RESUME is always ignored for a non-durable managed subscription: but specifying it without initializing the MQCHARV structure for sd.SubName does cause the error.

3. In addition there is a third flag affecting how the subscription is opened, MQSO_ALTER. Given the right permissions, the properties of a resumed subscription are changed to match other attributes specified in MQSUB.

Note: At least one of the MQSO_CREATE, MQSO_RESUME and MQSO_ALTER flags must be specified. See Options (MQLONG). There are examples of using all three flags in [“Example 3: Unmanaged MQ subscriber”](#) on page 281.

4. Set MQSO_MANAGED for the queue manager to manage the subscription for you automatically.

sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;

Optionally, omit setting the length of MQCHARV for null terminated strings and use the null terminator flag instead.

sd.ResObjectString.VSPtr = resTopicStr;

The resulting topic string is echoed in first printf in the program. Set up MQCHARV ResObjectString for WebSphere MQ to return the resolved string back to the program.

Note: resTopicStringBuffer is initialized to nulls in memset(resTopicStr, 0, sizeof(resTopicStrBuffer)). Returned topic strings do not end with a trailing null.

sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;

Set the buffer size of the sd.ResObjectString to one less than its actual size. This prevents overwriting the null terminator that is provided, in case the resolved topic string fills the entire buffer.

Note: No error is returned if the topic string is longer than sizeof(resTopicStrBuffer)-1. Even if VSLength > VSBufSiz the length returned in sd.ResObjectString.VSLength is the length of the complete string and not necessarily the length of the returned string. Test sd.ResObjectString.VSLength < sd.ResObjectString.VSBufSiz to confirm the topic string is complete.

MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);

The MQSUB function creates a subscription. If it is non-durable you are probably not interested in its name, though you can inspect its status in WebSphere MQ Explorer. You can provide the sd.SubName parameter as input, so you know what name to look for; you obviously have to avoid name clashes with other subscriptions.

MQCLOSE(Hconn, &Hsub, MQCO_REMOVE_SUB, &CompCode, &Reason);

Closing both the subscription and the subscription queue is optional. In the example the subscription is closed, but not the queue. The MQCLOSE MQCO_REMOVE_SUB option is the default in this case anyway as the subscription is non-durable. Using MQCO_KEEP_SUB is an error.

Note: the subscription *queue* is not closed by MQSUB, and its handle, Hobj, remains valid until the queue is closed by MQCLOSE or MQDISC. If the application terminates prematurely, the queue and subscription are cleaned up by the queue manager sometime after application termination.

Related concepts

[“Example 1: MQ Publication consumer” on page 274](#)

The MQ Publication consumer is an IBM WebSphere MQ message consumer that does not subscribe to topics itself.

[“Example 3: Unmanaged MQ subscriber” on page 281](#)

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

[“Writing publisher applications” on page 267](#)

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Example 3: Unmanaged MQ subscriber

The unmanaged subscriber is an important class of subscriber application. With it, you combine the benefits of publish/subscribe with *control* of queuing and consumption of publications. The example demonstrates different ways of combining subscriptions and queues.

The unmanaged pattern is more commonly associated with *durable* subscriptions than *non-durable*. Typically the lifecycle of a subscription created by an unmanaged subscriber is independent of the lifecycle of the subscribing application itself. By making the subscription durable the subscription receives publications even when no subscribing application is active.

You can create durable *managed* subscriptions to achieve the same result, but some applications require more flexibility and control over queues and messages than is possible with a managed subscription. For a durable managed subscription, the queue manager creates a permanent queue for the publications that match the subscription topic. It deletes the queue and associated publications when the subscription is deleted.

Typically durable *managed* subscriptions are used if the lifecycle of the application and the subscription is essentially the same, but hard to guarantee. By making the subscription durable, and having the publisher create persistent publications, there are no lost messages should the queue manager or subscriber terminate prematurely and need to be recovered.

The queue manager implicitly opens the durable managed subscription queue for a subscriber in such a way that shared processing of the queue is not possible. In addition, you cannot create more than one subscription for each managed queue and you may find the queues harder to manage because you have less control over the names of the queues. For these reasons, consider whether the *unmanaged* MQ subscriber is a better fit for applications requiring durable subscriptions than the *managed* MQ subscriber.

The code in [Figure 49 on page 286](#) demonstrates an unmanaged durable subscription pattern. For illustration the code also creates unmanaged, non-durable subscriptions. This example illustrates the following pattern facets:

- On demand subscriptions: the subscription topic strings are dynamic. They are provided by the application when it runs.
- Simplified subscription topic management: subscription topic management is simplified by defining the root part of the subscription topic string using an administratively defined topic. This hides the root part of the topic tree from the application. By hiding the root part a subscriber can be deployed to different topic trees.

- Flexible subscription management: you can define a subscription either administratively, or create it on-demand in a subscriber program. There is no difference between administratively and programmatically created subscriptions, except an attribute that shows how the subscription was created. There is a third type of subscription that is created automatically by the queue manager for distribution of subscriptions. All subscriptions are displayed in the WebSphere MQ Explorer.
- Flexible association of subscriptions with queues: a predefined local queue is associated with a subscription by the MQSUB function. There are different ways to use MQSUB to associate subscriptions with queues:
 - Associate a subscription with a queue having *no* existing subscriptions, MQSO_CREATE + (Hobj from MQOPEN).
 - Associate a *new* subscription with a queue having existing subscriptions, MQSO_CREATE + (Hobj from MQOPEN).
 - Move an existing subscription to a different queue, MQSO_ALTER + (Hobj from MQOPEN).
 - Resume an existing subscription associated with an existing queue, MQSO_RESUME + (Hobj = MQHO_NONE), or MQSO_RESUME + (Hobj = from MQOPEN of queue with existing subscription).
 - By combining MQSO_CREATE | MQSO_RESUME | MQSO_ALTER in different combinations, you can cater for different input states of the subscription and the queue without having to code multiple versions of MQSUB with different sd.Options values.
 - Alternatively, by coding a specific choice of MQSO_CREATE | MQSO_RESUME | MQSO_ALTER the queue manager returns an error ([Table 43 on page 283](#)) if the states of the subscription and queue provided as input to MQSUB are inconsistent with the value of sd.Options. [Figure 55 on page 289](#) shows the results of issuing MQSUB for Subscription X with different individual settings of the sd.Options flag, and passing it three different object handles.

Explore different inputs to the example program in [Figure 48 on page 285](#) to become familiar with these different kinds of errors. One common error, RC = 2440, that is not included in the cases listed in the table, is a subscription name error. It is commonly caused by passing a null or invalid subscription name with MQSO_RESUME or MQSO_ALTER.

- Multiprocessing: You can share among many consumers the work of reading publications. The publications all go onto the single queue associated with the subscription topic. Consumers have a choice of opening the queue directly using MQOPEN or resuming the subscription using MQSUB.
- Subscription concentration: multiple subscriptions can be created on the same queue. Be cautious with this capability as it can lead to "overlapping" subscriptions, and receiving the same publication multiple times. The MQSO_GROUP_SUB option eliminates duplicate publications caused by overlapping subscriptions.
- Subscriber and consumer separation: As well as the three consumer models illustrated in the examples, another model is to separate the consumer from the subscriber. It is a variation of the unmanaged MQ Subscriber, but rather than issue the MQOPEN and MQSUB in the same program, one program subscribes to publications, and another program consumes them. For example, the subscriber might be part of a publish/subscribe cluster and the consumer attached to a queue manager outside the queue manager cluster. The consumer receives publications through standard distributed queuing by defining the subscription queue as a remote queue definition.

Understanding the behavior of MQSO_CREATE | MQSO_RESUME | MQSO_ALTER is important, especially if you plan to simplify your code by using combinations of these options. Study the table [Table 43 on page 283](#) that shows the results of passing different queue handles to MQSUB, and the results of running the example program shown in [Figure 50 on page 287](#) to [Figure 55 on page 289](#).

The scenario used to construct the table has one subscription X and two queues, A and B. The subscription name parameter sd.SubName is set to X, the name of a subscription attached to queue A. Queue B has no subscription attached to it.

In [Table 43 on page 283](#), MQSUB is passed subscription X and the queue handle to queue A. The results from subscription options are as follows:

- MQSO_CREATE fails because the queue handle corresponds to the queue A which already has a subscription to X. Contrast this behavior to the successful call. That call succeeds because queue B does not have a subscription to X attached to it.
- MQSO_RESUME succeeds because the queue handle corresponds to the queue A which already has a subscription to X. In contrast, the call fails where the subscription X does not exist on queue A.
- MQSO_ALTER behaves in a similar way to MQSO_RESUME with respect to opening the subscription and queue. However if the attributes contained within the subscription descriptor passed to MQSUB differ from the attributes of the subscription, MQSO_RESUME fails, whereas MQSO_ALTER succeeds as long as the program instance has permission to alter the attributes. Note that you can never change the topic string in a subscription; but rather than return an error, MQSUB ignores the topic name and topic string values in the subscription descriptor and uses the values in the existing subscription.

Next, look at Table 43 on page 283 where MQSUB is passed subscription X and the queue handle to queue B. The results from subscription options are as follows:

- MQSO_CREATE succeeds and creates subscription X on queue B because this is a new subscription on queue B.
- MQSO_RESUME fails. MQSUB looks for subscription X on queue B and does not find it, but rather than returning *RC = 2428 - subscription X does not exist*, it returns *RC = 2019 - Subscription queue does not match queue object handle*. The behavior of the third option MQSO_ALTER suggests the reason for this unexpected error. MQSUB expects the queue handle to point to a queue with a subscription. It checks this first before checking whether the subscription named in sd.SubName exists.
- MQSO_ALTER succeeds, and moves the subscription from queue A to queue B.

A case that is not shown in the table is if the subscription name of the subscription on queue A does not match the subscription name in sd.SubName. That call fails with a *RC = 2428 - subscription X does not exist on Queue A*.

Table 43. Errors from MQSUB with different queue handles and subscription combinations		
	Queue A Subscription X Queue B No subscription	Queue A No subscription Queue B No subscription
Hobj for Queue A passed to MQSUB	MQSO_CREATE RC = 2432 - Subscription X already exists on Queue A MQSO_RESUME Resumes subscription X on Queue A MQSO_ALTER Resumes subscription X on Queue A and makes permitted alterations	MQSO_CREATE Creates subscription X on Queue A MQSO_RESUME RC = 2428 - Subscription X does not exist on Queue A MQSO_ALTER RC = 2428 - Subscription X does not exist on Queue A
Hobj for Queue B passed to MQSUB	MQSO_CREATE Creates new subscription X on Queue B MQSO_RESUME RC = 2019 - Subscription queue does not match queue object handle MQSO_ALTER Move subscription X from Queue A to Queue B	MQSO_CREATE Creates new subscription X on Queue B MQSO_RESUME RC = 2428 - subscription X does not exist on Queue B MQSO_ALTER RC = 2428 - subscription X does not exist on Queue B

Table 43. Errors from MQSUB with different queue handles and subscription combinations (continued)

	Queue A Subscription X Queue B No subscription	Queue A No subscription Queue B No subscription
MQHO_NONE passed to MQSUB	MQSO_CREATE RC = 2019 - Bad object handle: set MQSO_MANAGED flag to create a managed subscription and create a managed queue MQSO_RESUME Resumes subscription X on Queue A and returns Hobj to Queue A MQSO_ALTER Resumes subscription X on Queue A, returns Hobj to Queue A and makes permitted alterations	MQSO_CREATE RC = 2019 - Bad object handle: set MQSO_MANAGED flag to create a managed subscription and create a managed queue MQSO_RESUME RC = 2428 - No subscription X MQSO_ALTER RC = 2019 - Bad object handle: No queue A or B

Note: The compact coding style is intended for readability not production use.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>

void inquireQname(MQHCONN HConn, MQHOBJ Hobj, MQCHAR48 qName);

int main(int argc, char **argv)
{
    MQCHAR48 topicNameDefault      = "STOCKS";
    char      topicStringDefault[]  = "IBM/PRICE";
    char      subscriptionNameDefault[] = "IBMSTOCKPRICESUB";
    char      subscriptionQueueDefault[] = "STOCKTICKER";
    char      publicationBuffer[101]; /* Allocate to receive messages */
    char      resTopicStrBuffer[151]; /* Allocate to resolve topic string */
    MQCHAR48 qmName = "";             /* Default queue manager */
    MQCHAR48 qName = "";             /* Allocate storage for MQINQ */

    MQHCONN Hconn = MQHC_UNUSABLE_HCONN; /* connection handle */
    MQHOBJ Hobj = MQHO_NONE;             /* subscription queue handle */
    MQHOBJ Hsub = MQSO_NONE;             /* subscription handle */
    MQLONG CompCode = MQCC_OK;           /* completion code */
    MQLONG Reason = MQRC_NONE;           /* reason code */
    MQLONG messlen = 0;
    MQOD od = {MQOD_DEFAULT};           /* Unmanaged subscription queue */
    MQSD sd = {MQSD_DEFAULT};           /* Subscription Descriptor */
    MQMD md = {MQMD_DEFAULT};           /* Message Descriptor */
    MQGMO gmo = {MQGMO_DEFAULT};        /* get message options */
    MQLONG sdOptions = MQSO_CREATE | MQSO_RESUME | MQSO_DURABLE |
    MQSO_FAIL_IF QUIESCING;

    char * topicName = topicNameDefault;
    char * topicString = topicStringDefault;
    char * subscriptionName = subscriptionNameDefault;
    char * subscriptionQueue = subscriptionQueueDefault;
    char * publication = publicationBuffer;
    char * resTopicStr = resTopicStrBuffer;
    memset(resTopicStrBuffer, 0, sizeof(resTopicStrBuffer));
}
```

Figure 47. Unmanaged MQ subscriber - part 1: declarations.

```

        switch(argv){
default:
    switch((argv[5][0])) {
case('A'): sdOptions = MQSO_ALTER | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
break;
case('C'): sdOptions = MQSO_CREATE | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
break;
case('R'): sdOptions = MQSO_RESUME | MQSO_DURABLE | MQSO_FAIL_IF QUIESCING;
break;
default:
;
}
case(5):
if (strcmp(argv[4],"/")) /* "/" invalid = No subscription */
subscriptionQueue = argv[4];
else {
*subscriptionQueue = '\0';
if (argc > 5) {
if (argv[5][0] == 'C') {
sdOptions = sdOptions + MQSO_MANAGED;
}
}
else
sdOptions = sdOptions + MQSO_MANAGED;
}
case(4):
if (strcmp(argv[3],"/")) /* "/" invalid = No subscription */
subscriptionName = argv[3];
else {
*subscriptionName = '\0';
sdOptions = sdOptions - MQSO_DURABLE;
}
case(3):
if (strcmp(argv[2],"/")) /* "/" invalid = No topic string */
topicString = argv[2];
else
*topicString = '\0';
case(2):
if (strcmp(argv[1],"/")) /* "/" invalid = No topic object */
topicName = argv[1];
else
*topicName = '\0';
case(1):
sd.Options = sdOptions;
printf("Optional parameters: "
printf("topicName, topicString, subscriptionName, subscriptionQueue, A(lter)|C(reate)|
R(esume)\n");
printf("Values \"%-48s\" \"%s\" \"%s\" \"%-48s\" sd.Options=%d\n",
topicName, topicString, subscriptionName, subscriptionQueue, sd.Options);
}

```

Figure 48. Unmanaged MQ subscriber - part 2: parameter handling.

Additional comments about the parameter handling in this example are as follows:

switch((argv[5][0]))

You have the choice of entering Alter | Create | Resume in parameter 5, to test the effect of overriding part of the MQSUB option setting used by default in the example. The default setting used by the example is MQSO_CREATE | MQSO_RESUME | MQSO_DURABLE.

Note: Setting MQSO_ALTER or MQSO_RESUME without setting MQSO_DURABLE is an error, and sd.SubName must be set and refer to a subscription that can be resumed or altered.

***subscriptionQueue = '\0';**

sdOptions = sdOptions + MQSO_MANAGED;

If the default subscription queue, STOCKTICKER is replaced by a null string then as long as MQSO_CREATE is set, the example sets the MQSO_MANAGED flag and creates a dynamic subscription queue. If Alter or Resume are set in the fifth parameter the behavior of the example will depend on the value of subscriptionName.

***subscriptionName = '\0';**

sdOptions = sdOptions - MQSO_DURABLE;

If the default subscription, IBMSTOCKPRICESUB, is replaced by a null string then the example removes the MQSO_DURABLE flag. If you run the example providing the default values for the other parameters an additional temporary subscription destined to STOCKTICKER is created and receives duplicate publications. Next time you run the example, without any parameters, you receive just one publication again.

```
do {
    MQCONN(qmName, &Hconn, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    if (strlen(subscriptionQueue)) {
        strncpy(od.ObjectName, subscriptionQueue, MQ_Q_NAME_LENGTH);
        MQOPEN(Hconn, &od, MQOO_INPUT_AS_Q_DEF | MQOO_FAIL_IF_QUIESCING | MQOO_INQUIRE,
            &Hobj, &CompCode, &Reason);
        if (CompCode != MQCC_OK) break;
    }
    strncpy(sd.ObjectName, topicName, MQ_TOPIC_NAME_LENGTH);
    sd.ObjectString.VSPtr = topicString;
    sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    sd.SubName.VSPtr = subscriptionName;
    sd.SubName.VSLength = MQVS_NULL_TERMINATED;
    sd.ResObjectString.VSPtr = resTopicStr;
    sd.ResObjectString.VSBufSize = sizeof(resTopicStrBuffer)-1;
    MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    gmo.Options = MQGMO_WAIT | MQGMO_NO_SYNCPOINT | MQGMO_CONVERT;
    gmo.WaitInterval = 10000;
    gmo.MatchOptions = MQMO_MATCH_CORREL_ID;
    memcpy(md.CorrelId, sd.SubCorrelId, MQ_CORREL_ID_LENGTH);
    inquireQname(Hconn, Hobj, qName);
    printf("Waiting %d seconds for publications matching \"%s\" from %-0.48s\n",
        gmo.WaitInterval/1000, resTopicStr, qName);
    do {
        memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
        memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
        md.Encoding = MQENC_NATIVE;
        md.CodedCharSetId = MQCCSI_Q_MGR;
        MQGET(Hconn, Hobj, &md, &gmo, sizeof(publication), publication, &messlen,
            &CompCode, &Reason);
        if (Reason == MQRC_NONE)
            printf("Received publication \"%s\"\n", publication);
    }
    while (CompCode == MQCC_OK);
    if (CompCode != MQCC_OK && Reason != MQRC_NO_MSG_AVAILABLE) break;
    MQCLOSE(Hconn, &Hsub, MQCO_NONE, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);
    if (CompCode != MQCC_OK) break;
    MQDISC(&Hconn, &CompCode, &Reason);
} while (0);
printf("Completion code %d and Return code %d\n", CompCode, Reason);
}
void inquireQname(MQHCONN Hconn, MQHOBJ Hobj, MQCHAR48 qName) {
#define _selectors 1
#define _intAttrs 1

    MQLONG select[_selectors] = {MQCA_Q_NAME}; /* Array of attribute selectors */
    MQLONG intAttrs[_intAttrs]; /* Array of integer attributes */
    MQLONG CompCode, Reason;
    MQINQ(Hconn, Hobj, _selectors, select, _intAttrs, intAttrs, MQ_Q_NAME_LENGTH, qName,
        &CompCode, &Reason);
    if (CompCode != MQCC_OK) {
        printf("MQINQ failed with Condition code %d and Reason %d\n", CompCode, Reason);
        strncpy(qName, "unknown queue", MQ_Q_NAME_LENGTH);
    }
    return;
}
```

Figure 49. Unmanaged MQ subscriber - part 3: code body.

Additional comments on the code in this example are as follows:

if (strlen(subscriptionQueue))

If there is no subscription queue name then the example uses MQHO_NONE as the value of Hobj.

MQOPEN(...);

The subscription queue is opened and the queue handle saved in Hobj.

MQSUB(Hconn, &sd, &Hobj, &Hsub, &CompCode, &Reason);

The subscription is opened using the Hobj passed from MQOPEN (or MQHO_NONE if there is no subscription queue name). An unmanaged queue can be resumed without explicitly opening it with an MQOPEN.

MQCLOSE(Hconn, &Hsub, MQCO_NONE, &CompCode, &Reason);

The subscription is closed using the subscription handle. Depending on whether the subscription is durable or not, the subscription is closed with an implicit MQCO_KEEP_SUB or MQCO_REMOVE_SUB. You can close a durable subscription with MQCO_REMOVE_SUB, but you *cannot* close a non-durable subscription with MQCO_KEEP_SUB. The action of MQCO_REMOVE_SUB is to remove the subscription which stops any further publications being sent to the subscription queue.

MQCLOSE(Hconn, &Hobj, MQCO_NONE, &CompCode, &Reason);

No special action is taken if the subscription is unmanaged. If the queue is managed and the subscription closed with either an explicit or implicit MQCO_REMOVE_SUB, then all publications are purged from the queue and queue deleted at this point.

gmo.MatchOptions = MQMO_MATCH_CORREL_ID;

memcpy(md.CorrelId, sd.SubCorrelId, MQ_CORREL_ID_LENGTH);

Ensure that the messages received are those for our subscription.

Results from the example illustrate aspects of publish/subscribe:

In [Figure 50 on page 287](#) the example starts by publishing 130 on the NYSE/IBM/PRICE topic.

```
W:\Subscribe3\Debug>...\Publish2\Debug\publishstock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0
```

Figure 50. Publish 130 to NYSE/IBM/PRICE

In [Figure 51 on page 287](#) execution of the example using default parameters receives the retained publication 130. The provided topic object and topic string are ignored, as shown in [Figure 55 on page 289](#). The topic object and topic string are always taken from the subscription object, when one is provided, and the topic string is immutable. The actual behavior of the example depends on the choice or combination of MQSO_CREATE, MQSO_RESUME, and MQSO_ALTER. In this example MQSO_RESUME is the option selected.

```
W:\Subscribe3\Debug>solution3
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(lter)|
C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8206
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0
```

Figure 51. Receive the retained publication

In [Figure 52 on page 288](#) no publications are received, because the durable subscription has already received the retained publication. In this example, the subscription is resumed by providing only the subscription name without the queue name. If the queue name was provided, the queue would be opened first and the handle passed to MQSUB.

Note: The 2038 error from MQINQ is due to the implicit MQOPEN of STOCKTICKER by MQSUB not including the MQOO_INQUIRE option. Avoid the 2038 return code from MQINQ by opening the queue explicitly.

```

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE IBMSTOCKPRICESUB / Resume
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|
C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "IBMSTOCKPRICESUB" "" sd.Options=8204
MQINQ failed with Condition code 2 and Reason 2038
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from unknown queue
Completion code 0 and Return code 0

```

Figure 52. Resume subscription

In Figure 53 on page 288 , the example creates a non-durable unmanaged subscription using STOCKTICKER as the destination. Because this is a new subscription, it receives the retained publication.

```

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE / STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|
C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "" "STOCKTICKER" sd.Options=8194
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0

```

Figure 53. Receive retained publication with new unmanaged non durable subscription

In Figure 54 on page 288 , to demonstrate overlapping subscriptions, another publication is sent, changing the retained publication. Next, a new non-durable, unmanaged subscription is created by not providing a subscription name. The retained publication is received twice, once for the new subscription, and once for the durable IBMSTOCKPRICESUB subscription that is still active on the STOCKTICKER queue. The example is an illustration that it is the queue that has subscriptions, and not the application. Despite not referring to the IBMSTOCKPRICESUB subscription in this invocation of the application, the application receives the publication twice: once from the durable subscription that was created administratively, and once from the non-durable subscription created by the application itself.

```

W:\Subscribe3\Debug>..\..\Publish2\Debug\publishstock
Provide parameters: TopicObject TopicString Publication
Publish "130" to topic "STOCKS" and topic string "IBM/PRICE"
Published "130" to topic string "NYSE/IBM/PRICE"
Completion code 0 and Return code 0

W:\Subscribe3\Debug>solution3 STOCKS IBM/PRICE / STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(1ter)|
C(reate)|R(esume)
Values "STOCKS" "IBM/PRICE" "" "STOCKTICKER" sd.Options=8194
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Received publication "130"
Completion code 0 and Return code 0

```

Figure 54. Overlapping subscriptions

In Figure 55 on page 289 the example demonstrates that providing a new topic string and an existing subscription does not result in a changed subscription.

1. In the first case, Resume resumes the existing subscription, as you might expect, and ignores the changed topic string.
2. In the second case, Alter causes an error, RC = 2510, Topic not alterable.
3. In the third example, Create causes an error RC = 2432, Sub already exists.


```

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Resume
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(lter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8204
Waiting 10 seconds for publications matching "NYSE/IBM/PRICE" from STOCKTICKER
Received publication "130"
Completion code 0 and Return code 0

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Alter
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(lter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8201
Completion code 2 and Return code 2510

W:\Subscribe3\Debug>solution3 "" NASDAQ/IBM/PRICE IBMSTOCKPRICESUB STOCKTICKER Create
Optional parameters: topicName, topicString, subscriptionName, subscriptionQueue, A(lter)|C(reate)|R(esume)
Values "" "NASDAQ/IBM/PRICE" "IBMSTOCKPRICESUB" "STOCKTICKER" sd.Options=8202
Completion code 2 and Return code 2432

```

Figure 55. Subscription topics cannot be changed

Related concepts

[“Example 1: MQ Publication consumer” on page 274](#)

The MQ Publication consumer is an IBM WebSphere MQ message consumer that does not subscribe to topics itself.

[“Example 2: Managed MQ subscriber” on page 276](#)

The managed MQ subscriber is the preferred pattern for most subscriber applications. The example requires *no* administrative definition of queues, topics or subscriptions.

[“Writing publisher applications” on page 267](#)

Get started with writing publisher applications by studying two examples. The first is modeled as closely as possible on a point to point application putting messages on a queue, and the second demonstrates creating topics dynamically - a more common pattern for publisher applications.

Publish/subscribe lifecycles

Consider the lifecycles of topics, subscriptions, subscribers, publications, publishers and queues in designing publish/subscribe applications.

The lifecycle of an object, such as a subscription, starts with its creation and ends with its deletion. It may also include other states and changes that it goes through, such as temporary suspension, having parent and children topics, expiration and deletion.

Traditionally WebSphere MQ objects such as queues are created administratively, or by administrative programs using Programmable Command Format (PCF). Publish/subscribe is different in providing the MQSUB and MQCLOSE API verbs to create and delete subscriptions, having the concept of managed subscriptions that not only create and delete queues, but also clean up unconsumed messages, and having associations between administratively created topic objects and programmatically or administratively created topic strings.

This functional richness caters for a wide range of publish/subscribe requirements, and also simplifies designing some common patterns of publish/subscribe application. Managed subscriptions, for example, simplify both the programming and administration of a subscription that is intended to last only as long as the program that created it. Unmanaged subscriptions simplify programming where there is a looser connection between subscribing and consuming publications. Centrally created subscriptions are useful where the pattern is one of routing publication traffic to consumers based on a centralized model of control, for example sending flight information to automated gates, whereas programmatically created subscriptions might be used if gate staff are responsible for subscribing to the passengers records for that flight, by entering a flight number at a gate.

In this last example a managed durable subscription might be appropriate: managed, because the subscriptions are being created very often, and have a clear endpoint when the gate closes and the subscription can be programmatically removed; durable, to avoid losing a passenger record due to the gate subscriber program going down for one reason or another². To initiate the publication of passenger records to the gate, a possible design would be for the gate application to both subscribe to the passenger records using the gate number, and publish the gate opening event using the gate number. The publisher responds to the gate opening event by publishing the passenger records - which might then also go to

² The publisher must send the passenger records as persistent messages to avoid other possible failures, of course.

other interested parties, such as billing, to record the flight is taking place, and to customer services, to text notifications to passengers' mobile phones of the gate number.

The centrally managed subscription might use a durable unmanaged model, routing passenger lists to the gate using a predefined queue for each gate.

The following three examples of publish/subscribe lifecycles illustrate how managed non-durable, managed durable, and unmanaged durable subscribers interact with subscriptions, topics, queues, publishers and the queue manager, and how the responsibilities might be divided between administration and the subscriber programs.

Managed non-durable subscriber

Figure 56 on page 291 shows an application creating a managed non-durable subscription, getting two messages that are published to the topic identified in the subscription, and terminating. The interactions labeled in an italic gray font with dotted arrows are implicit.

There are some points to note.

1. The application creates a subscription on a topic that has already been published to twice. When the subscriber receives its first publication, it receives the *second* publication which is the currently retained publication.
2. The queue manager creates a temporary subscription queue as well as creating a subscription for the topic.
3. The subscription has an expiry. When the subscription expires no more publications on the topic are sent to this subscription, but the subscriber continues to get messages published before the subscription expired. Publication expiry is not affected by subscription expiry.
4. The fourth publication is not placed on the subscription queue and consequently the last *MQGET* does not return a publication.
5. Although the subscriber closes its subscription, it does not close its connection to the queue or the queue manager.
6. The queue manager cleans up shortly after the application terminates. Because the subscription is managed and non-durable, the subscription queue is deleted.

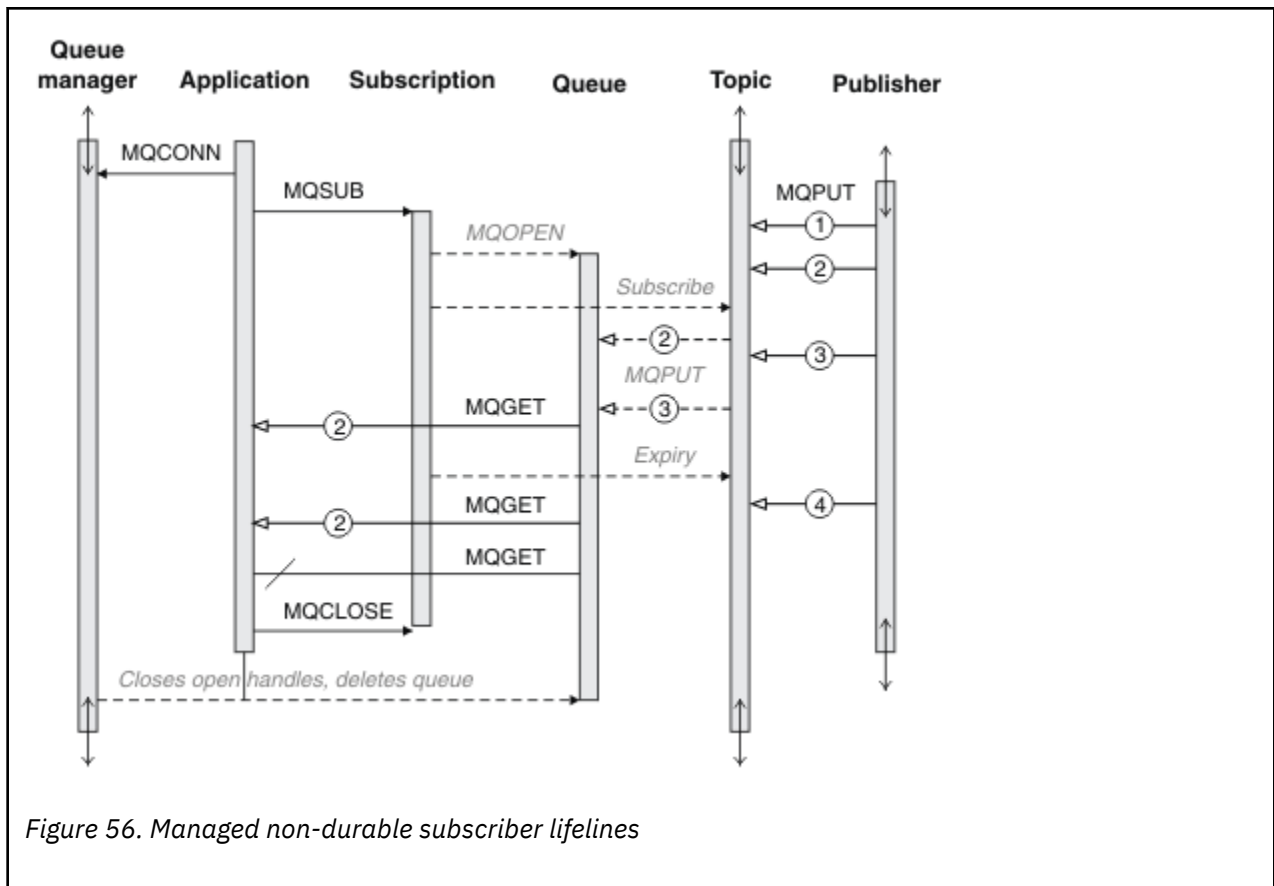


Figure 56. Managed non-durable subscriber lifelines

Managed durable subscriber

The managed durable subscriber takes the previous example a step further, and shows a managed subscription surviving the termination and restart of the subscribing application.

There are some new points to note.

1. In this example, unlike the last, the publication topic did not exist before it was defined in the subscription.
2. The first time the subscriber terminates, it closes the subscription with the option `MQCO_KEEP_SUB`. That is the default behavior for implicitly closing a managed durable subscription.
3. When the subscriber resumes the subscription, the subscription queue is reopened.
4. The new publication 2, placed on the queue before it is reopened, is available to `MQGET`, even after the subscription has been removed.

Even though the subscription is durable, the subscriber reliably receives all messages sent by the publisher only if *both* the subscription is durable and the messages persistent. Message persistence depends on the setting of the `Persistent` field in the `MQMD` of the message sent by the publisher. A subscriber has no control over this.

5. Closing the subscription with the flag `MQCO_REMOVE_SUB` removes the subscription, stopping any further publications being placed on the subscription queue. When the subscription queue is closed, then the queue manager removes the unread publication 3, and then deletes the queue. The action is equivalent to administratively deleting the subscription.

Note: Do not delete the queue manually, or issue `MQCLOSE` with the option `MQCO_DELETE`, or `MQCO_PURGE_DELETE`. The visible implementation details of a managed subscription is not part of the supported WebSphere MQ interface. The queue manager cannot manage a subscription reliably unless it has complete control.

tidy up later, either manually, or using a automated management program, when you can be sure that messages, subscriptions, and queues can be simply deleted, regardless of their state.

Publish/subscribe message properties

Several message properties relate to WebSphere MQ publish/subscribe messaging.

PubAccountingToken

This is the value that will be in the AccountingToken field of the Message Descriptor (MQMD) of all publication messages matching this subscription. AccountingToken is part of the identity context of the message. For more information about message context, see [“Message context” on page 37](#). For more information about the AccountingToken field in the MQMD, see [AccountingToken](#).

PubApplIdentityData

This is the value that will be in the ApplIdentityData field of the Message Descriptor (MQMD) of all publication messages matching this subscription. ApplIdentityData is part of the identity context of the message. For more information about message context, see [“Message context” on page 37](#). For more information about the ApplIdentityData field in the MQMD, see [ApplIdentityData](#).

If the option MQSO_SET_IDENTITY_CONTEXT is not specified, the ApplIdentityData which will be set in each message published for this subscription is blanks, as default context information.

If the option MQSO_SET_IDENTITY_CONTEXT is specified, the PubApplIdentityData is being generated by the user and this field is an input field which contains the ApplIdentityData to be set in each publication for this subscription.

PubPriority

This is the value that will be in the Priority field of the Message Descriptor (MQMD) of all publication messages matching this subscription. For more information about the Priority field in the MQMD, see [Priority](#).

The value must be greater than or equal to zero; zero is the lowest priority. The following special values can also be used:

- MQPRI_PRIORITY_AS_Q_DEF - When a subscription queue is provided in the Hobj field in the MQSUB call, and is not a managed handle, then the priority for the message is taken from the DefPriority attribute of this queue. If the queue so identified is a cluster queue, or there is more than one definition in the queue-name resolution path, the priority is determined when the publication message is put to the queue as described for [Priority](#) in the MQMD. If the MQSUB call uses a managed handle, the priority for the message is taken from the DefPriority attribute of the model queue associated with the topic subscribed to.
- MQPRI_PRIORITY_AS_PUBLISHED - The priority for the message is the priority of the original publication. This is the initial value of this field.

SubCorrelId



Attention: a correlation identifier can only be passed between queue managers in a publish/subscribe cluster, not a hierarchy.

All publications sent to match this subscription will contain this correlation identifier in the message descriptor. If multiple subscriptions use the same queue to get their publications from, using MQGET by correlation id allows only publications for a specific subscription to be obtained. This correlation identifier can either be generated by the queue manager or by the user.

If the option MQSO_SET_CORREL_ID is not specified, the correlation identifier is generated by the queue manager and this field is an output field which contains the correlation identifier which will be set in each message published for this subscription.

If the option MQSO_SET_CORREL_ID is specified, the correlation identifier is being generated by the user and this field is an input field which contains the correlation identifier to be set in each publication for this subscription. In this case, if the field contains MQCI_NONE, the correlation identifier which will be set in each message published for this subscription will be the correlation identifier created by the original put of the message.

If the option MQSO_GROUP_SUB is specified and the correlation identifier specified is the same as an existing grouped subscription using the same queue and an overlapping topic string, only the most significant subscription in the group is provided with a copy of the publication.

SubUserData

This is the subscription user data. The data provided on the subscription in this field will be included as the MQSubUserData message property of every publication sent to this subscription.

Publication properties

Table 44 on page 295 lists the publication properties that are provided with a publication message.

You can access these properties directly from the **MQRFH2** folder, or retrieve them using MQINQMP. MQINQMP accepts either the property name or **MQRFH2** name as the name of the property to inquire on.

Table 44. Publication properties			
Property name	MQRFH2 name	Type	Description
MQTopicString	mmps.Top	MQTYPE_STRING	Topic string
MQSubUserData	mmps.Sud	MQTYPE_STRING	Subscriber user data
MQIsRetained	mmps.Ret	MQTYPE_BOOLEAN	Retained publication
MQPubOptions	mmps.Pub	MQTYPE_INT32	Publication options
MQPubLevel	mmps.Pbl	MQTYPE_INT32	Publication level
MQPubTime	mmpse.Pts	MQTYPE_STRING	Publication time
MQPubSeqNum	mmpse.Seq	MQTYPE_INT32	Publication sequence number
MQPubStrIntData	mmpse.Sid	MQTYPE_STRING	String/Integer data added by the publisher
MQPubFormat	mmpse.Pfmt	MQTYPE_INT32	Message format: MQRFH1 MQRFH2 PCF

Message ordering

For a particular topic, messages are published by the queue manager in the same order as they are received from publishing applications (subject to reordering based on message priority).

Message ordering normally means that each subscriber receives messages from a particular queue manager, on a particular topic, from a particular publisher in the order that they are published by that publisher.

However, as with all WebSphere MQ messages, it is possible for messages, occasionally, to be delivered out of order. This can happen in the following situations:

- If a link in the network goes down and subsequent messages are rerouted along another link

- If a queue becomes temporarily full, or put-inhibited, so that a message is put to a dead-letter queue and therefore delayed, while subsequent messages pass straight through.
- If the administrator deletes a queue manager when publishers and subscribers are still operating, causing queued messages to be put to the dead-letter queue and subscriptions to be interrupted.

If these circumstances cannot occur, publications are always delivered in order.

Note: It is not possible to use grouped or segmented messages with Publish/Subscribe.

Intercepting publications

You can intercept a publication, modify it, and then republish it before it reaches any other subscriber.

You might want to intercept a publication before it reaches a subscriber in order to do one of the following actions:

- Attach additional information to the message
- Block the message
- Transform the message

You can perform the same operation on each message or vary the operation depending on the subscription, the message, or the message header.

Related reference

[MQ_PUBLISH_EXIT](#) - Publish exit

Subscription levels

Set the subscription level of a subscription to intercept a publication before it reaches its final subscribers. An intercepting subscriber subscribes at a higher subscription level, and republishes at a lower publication level. Build a chain of intercepting subscribers to perform message processing on a publication before it is delivered to final subscribers.




Figure 59. Sequence of intercepting subscribers

To intercept a publication, use the **MQSD** SubLevel attribute. After a message has been intercepted, it can be transformed and then republished at a lower publication level by changing the **MQPMO** PubLevel attribute. The message then goes to the final subscribers, or it is intercepted again by an intermediate subscriber at a lower subscription level.

The intercepting subscriber typically transforms a message before republishing it. A sequence of intercepting subscribers forms a message flow. Alternatively, you might not republish the intercepted publication: Subscribers at lower subscription levels would not receive the message.

Ensure that the interceptor receives publications before any other subscribers. Set the subscription level of the interceptor higher than other subscribers. By default, subscribers have a SubLevel of 1. The highest value is 9. A publication must start with a PubLevel at least as high as the highest SubLevel. Publish initially with the default PubLevel of 9.

- If you have one intercepting subscriber on a topic, set the SubLevel to 9.
- For multiple intercepting applications on a topic, set a lower SubLevel for each successive intercepting subscriber.
- You can implement a maximum of 8 intercepting applications, with subscription levels from 9 down to 2 inclusive. The final recipient of the message has a SubLevel of 1.

The interceptor with the highest subscription level that is equal to, or lower than, the `PubLevel` of the publication receives the publication first. Configure only one intercepting subscriber for a topic at a particular subscription level. Having multiple subscribers at a particular subscription level results in multiple copies of the publication being sent to the final set of subscribing applications.

A subscriber with a `SubLevel` of 0 is used as a catchall. It receives the publication if no final subscriber gets the message. A subscriber with `SubLevel` of 0 might be used to monitor the publications that no other subscribers received.

Programming an intercepting subscriber

Use the subscription options described in [Table 45 on page 297](#).

Table 45. Subscription options for intercepting subscribers	
Subscription option	Notes
<code>MQSO_SET_CORREL_ID</code> and <code>SubCorrelId</code> set to <code>MQCI_NONE</code>	Keep the <code>CorrelId</code> of the intercepted publication the same as the original publication. Note: You cannot pass the correlation identifier of a publication in a hierarchy. The field is used by the queue manager.
<code>PubPriority</code> set to <code>MQPRI_PRIORITY_AS_PUBLISHED</code>	Keep the priority of the intercepted publication the same as the original publication.

The options in [Table 45 on page 297](#) must be used by all the intercepting subscribers. The result is that the correlation identifier and message priority are not modified from the setting of the original publisher.

When the intercepting subscriber has processed the publication, it republishes the message to the same topic at a `PubLevel` one lower than the `SubLevel` of its own subscription. If the intercepting subscriber set a `SubLevel` of 9, it republishes the message with a `PubLevel` of 8.

To republish the message correctly, several pieces of information from the original publication are required. Reuse the same **MQMD** as in the original message and set `MQPMO_PASS_ALL_CONTEXT` to ensure all information in the **MQMD** is passed on to the next subscriber. Copy the values from the message properties shown in [Table 46 on page 297](#) into the corresponding fields of the republished message. The intercepting subscriber can change these values. Use the OR operator to add additional values to the **MQPMO.Options** field, to combine the put message options.

You must open the publication queue explicitly rather than use a managed publication queue. You cannot set `MQSO_SET_CORREL_ID` for a managed queue. You also cannot set `MQOO_SAVE_ALL_CONTEXT` on a managed queue. See the code fragments listed in [“Examples” on page 298](#).

Table 46. MQPUT values for republished messages	
Republish message using MQPUT	Information in publication message
<code>MQOD.ObjectString</code>	Message property <code>MQTopicString</code>
<code>MQPMO.Options</code>	Message property <code>MQPubOptions</code>

The final subscriber has the choice of setting its subscription options differently. For example, it might set the publication priority explicitly rather than to `MQPRI_PRIORITY_AS_PUBLISHED`. The settings of a final subscriber only affect publication from the final intercepting subscriber in the chain.

Retained publications

A retained publication must be preserved after it has been intercepted, by copying the original put-message options into the republished message.

The MQPMO_RETAIN option is set by the publisher. Each intercepting subscriber must transfer the MQPubOptions to the put-message options of the republished message as shown in [Table 46 on page 297](#). Copying the put-message options preserves the options set by the original publisher, including whether to retain the publication.

When a publication finishes its passage down the chain of intercepting subscribers, and is delivered to final subscribers, it is finally retained. New subscribers, at SubLevel 1, requesting the retained publication, receive it without any further interception. Subscribers at a SubLevel greater than 1 are not sent the retained publication. As a result, the retained publication is not modified by the chain of intercepting subscribers a second time round.

Examples

The examples are code fragments that can be combined to build an intercepting subscriber. The code is written to be brief, rather than of production quality.

The preprocessor directives in [Figure 60 on page 298](#) define the two properties to be extracted from the publication messages that are required by the MQINQMP MQI call.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <cmqc.h>
#define      MQPUBOPTIONS      (MQPTR)(char*) "MQPubOptions",\
                                0,\
                                12,\
                                MQVS_NULL_TERMINATED,\
                                MQCCSI_APPL
#define      MQTOPICSTRING     (MQPTR)(char*) "MQTopicString",\
                                0,\
                                13,\
                                MQVS_NULL_TERMINATED,\
                                MQCCSI_APPL
```

Figure 60. Preprocessor directives

Figure 61 on page 299 lists the declarations used in the code fragments. Except for the highlighted terms, the declarations are standard for a WebSphere MQ application.

The highlighted Put and Get options are initialized to pass all context. The highlighted MQTOPICSTRING and MQPUBOPTIONS are MQCHARV initializers for property names that are defined in the preprocessor directives. The names are passed to MQINQMP.

```

int main(int argc, char **argv) {
    MQLONG Reason = MQRC_NONE;
    MQLONG CompCode = MQCC_OK;
    MQHCONN Hcon = MQHC_UNUSABLE_HCONN;
    MQCHAR QMName[49] = "";
    MQCMHO CrtMsgHOpts = {MQCMHO_DEFAULT};
    MQHMSG Hmsg = MQHM_NONE;
    MQMD md = {MQMD_DEFAULT};
    MQHOBJ gHobj = MQHO_NONE;
    MQOD getOD = {MQOD_DEFAULT};
    MQGMO gmo = {MQGMO_DEFAULT};
    MQLONG GO_Options = MQOO_INPUT_AS_Q_DEF
        | MQOO_FAIL_IF QUIESCING
        | MQOO_SAVE_ALL_CONTEXT;
    MQLONG GC_Options = MQCO_DELETE_PURGE;
    MQHOBJ Hsub = MQHO_NONE;
    MQSD sd = {MQSD_DEFAULT};
    MQLONG SC_Options = MQCO_NONE;
    MQHOBJ pHobj = MQHO_NONE;
    MQOD putOD = {MQOD_DEFAULT};
    MQLONG PO_Options = MQOO_OUTPUT
        | MQOO_FAIL_IF QUIESCING
        | MQOO_PASS_ALL_CONTEXT;
    MQLONG PC_Options = MQCO_NONE;
    MQPMO pmo = {MQPMO_DEFAULT};
    MQIMPO InqPropOpts = {MQIMPO_DEFAULT};
    MQPD PropDesc = {MQPD_DEFAULT};
    MQLONG Type = MQTYPE_AS_SET;
    MQCHARV TopStrProp = {MQTOPICSTRING};
    MQCHARV PubOptProp = {MQPUBOPTIONS};
    MQLONG DataLength = 0;
    MQBYTE buffer[256] = "";
    MQLONG buflen = sizeof(buffer) - 1;
    MQLONG messlen = 0;
    char TopStrBuf[256] = "Initial value";
    int i = 0;
}

```

Figure 61. Declarations

Initializations that are not easily performed in declarations are shown in [Figure 62 on page 300](#). The highlighted values require explanation.

SYSTEM.NDURABLE.MODEL.QUEUE

In this example, instead of using MQSUB to open a managed non-durable subscription, the model queue, SYSTEM.NDURABLE.MODEL.QUEUE, is used to create a temporary dynamic queue. Its handle is passed to MQSUB. By opening the queue directly you are able to save all message context and set the subscription option, MQSO_SET_CORREL_ID.

MQGMO_CURRENT_VERSION

It is important to use the current version of most of the WebSphere MQ structures. Fields such as gmo.MsgHandle are only available in the latest version of the control structures.

MQGMO_PROPERTIES_IN_HANDLE

The topic string and put message options set in the original publication are to be retrieved by the intercepting subscriber using message properties. An alternative would be to read the **MQRFH2** structure in the message directly.

MQSO_SET_CORREL_ID

Use MQSO_SET_CORREL_ID in combination with,

```
memcpy(sd.SubCorrelId, MQCI_NONE, sizeof(sd.SubCorrelId));
```

The effect of these options is to pass on the correlation identifier. The correlation identifier set by the original publisher is placed in the correlation identifier field of the publication that is received by the intercepting subscriber. Each intercepting subscriber passes on the same correlation identifier. The final subscriber then has the option of receiving the same correlation identifier.

Note: If the publication is passed through a publish/subscribe hierarchy, the correlation identifier is never retained.

MQPRI_PRIORITY_AS_PUBLISHED


The publication is placed on the publication queue with the same message priority as it was published with.

```
strcpy(getOD.ObjectName, "SYSTEM.NDURABLE.MODEL.QUEUE",
      sizeof(getOD.ObjectName));
gmo.Version = MQGMO_VERSION_4;
gmo.Options = MQGMO_WAIT
             | MQGMO_PROPERTIES_IN_HANDLE
             | MQGMO_CONVERT;
gmo.WaitInterval = 30000;
sd.Options = MQSO_CREATE
            | MQSO_FAIL_IF_QUIESCING
            | MQSO_SET_CORREL_ID;
sd.PubPriority = MQPRI_PRIORITY_AS_PUBLISHED;
sd.Version = MQSD_VERSION_1;
memcpy(sd.SubCorrelId, MQCI_NONE, sizeof(sd.SubCorrelId));
putOD.ObjectType = MQOT_TOPIC;
putOD.ObjectString.VSPtr = &TopStrBuf;
putOD.ObjectString.VSBufSize = sizeof(TopStrBuf);
putOD.ObjectString.VSLength = MQVS_NULL_TERMINATED;
putOD.ObjectString.VSCCSID = MQCCSI_APPL;
putOD.Version = MQOD_VERSION_4;
pmo.Version = MQPMO_VERSION_3;
```

Figure 62. Initializations

Figure 63 on page 301 shows the code fragment to read command-line parameters, complete the initialization, and create the intercepting subscription.

Run the program with the command,



To make error handling as unobtrusive as possible, the reason code from each MQI call is stored in a different array element. After each call the completion code is tested, and if the value is MQCC_FAIL, control exits the `do { } while(0)` code block.

The two noteworthy lines of code are,

pmo.PubLevel = sd.SubLevel - 1;

Sets the publication level for the republished message to one less than the subscription level of the intercepting subscriber.

gmo.MsgHandle = Hmsg;

Provides a message handle for MQGET to return the message properties.

```

do {
    printf("Intercepting subscriber start\n");
    if (argc < 2) {
        printf("Required parameter missing - topic string\n");
        exit(99);
    } else {
        sd.ObjectString.VSPtr = argv[1];
        sd.ObjectString.VSLength = MQVS_NULL_TERMINATED;
        printf("TopicString = %s\n", sd.ObjectString.VSPtr);
    }
    if (argc > 2) {
        sd.SubLevel = atoi(argv[2]);
        pmo.PubLevel = sd.SubLevel - 1;
        printf("SubLevel is %d, PubLevel is %d\n", sd.SubLevel, pmo.PubLevel);
    }
    if (argc > 3)
        strncpy(QMName, argv[3], sizeof(QMName));
    MQCONN(QMName, &Hcon, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQOPEN(Hcon, &getOD, GO_Options, &gHobj, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQSUB(Hcon, &sd, &gHobj, &Hsub, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQCRTMH(Hcon, &CrtMsgHOpts, &Hmsg, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    gmo.MsgHandle = Hmsg;
}

```

Figure 63. Preparing to intercept publications

The main code fragment, [Figure 64 on page 302](#), gets messages from the publication queue. It queries the message properties and republishes the messages using the topic string, and the original **MQPMO**.option properties of the publication.

In this example, no transformation is performed on the publication. The topic string of the republished publication always matches the topic string the intercepting subscriber subscribed on. If the intercepting subscriber is responsible for intercepting multiple subscriptions sent to the same publication queue, it might be necessary to query the topic string to distinguish publications that match different subscriptions.

The calls to MQINQMP are highlighted. The topic string and publication put message options properties are written directly into the output control structures. The only reason for altering the MQCHARV length field of putOD.ObjectString from an explicit length to a null terminated string is to use `printf` to output the string.

```

while (CompCode != MQCC_FAILED) {
    memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
    memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
    md.Encoding = MQENC_NATIVE;
    md.CodedCharSetId = MQCCSI_Q_MGR;
    printf("MQGET : %d seconds wait time\n", gmo.WaitInterval/1000);
    MQGET(Hcon, gHobj, &md, &gmo, buflen, buffer, &messlen,
        &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    buffer[messlen] = '\0';
    MQINQMP(Hcon, Hmsg, &InqPropOpts, &TopStrProp, &PropDesc, &Type,
        putOD.ObjectString.VSBufSize, putOD.ObjectString.VSPtr,
        &(putOD.ObjectString.VSLength), &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    memset((void *)((MQLONG)(putOD.ObjectString.VSPtr)
        + putOD.ObjectString.VSLength), '\0', 1);
    putOD.ObjectString.VSLength = MQVS_NULL_TERMINATED;
    MQINQMP(Hcon, Hmsg, &InqPropOpts, &PubOptProp, &PropDesc, &Type,
        sizeof(pmo.Options), &(pmo.Options), &DataLength,
        &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQOPEN(Hcon, &putOD, PO_Options, &pHobj, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    printf("Republish message <%s> on topic <%s> with options %d\n",
        buffer, putOD.ObjectString.VSPtr, pmo.Options);
    MQPUT(Hcon, pHobj, &md, &pmo, messlen, buffer, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
    MQCLOSE(Hcon, &pHobj, PC_Options, &CompCode, &Reason);
    if (CompCode == MQCC_FAILED)
        break;
}

```

Figure 64. Intercept publication and republish

The final code fragment is shown in [Figure 65 on page 302](#).

```

} while (0);
if (CompCode == MQCC_FAILED && Reason != MQRC_NO_MSG_AVAILABLE)
    printf("MQI Call failed with reason code %d\n", Reason);
if (Hsub != MQHO_NONE)
    MQCLOSE(Hcon, &Hsub, SC_Options, &CompCode, &Reason);
if (Hcon != MQHC_UNUSABLE_HCONN)
    MQDISC(&Hcon, &CompCode, &Reason);
}

```

Figure 65. Completion

Intercepting publications and distributed publish/subscribe

Follow a simple pattern when you deploy intercepting subscribers or Publish exits to a distributed publish/subscribe topology. Deploy intercepting subscribers on the same queue managers as publishers, and Publish exits on the same queue managers as final subscribers.

[Figure 66 on page 303](#) shows two queue managers connected in a publish subscribe cluster. A publisher creates a publication to a cluster topic at publication level 9. The numbered arrows show the sequence of steps taken by the publication as it flows to subscribers to the cluster topic. The publication is intercepted by the subscriber with Sublevel 9 and republished with Publevel 8. It is intercepted again by a subscriber at Sublevel 8. The subscriber republishes at Publevel 7. The proxy subscriber provided by the queue manager forwards the publication to queue manager B, where a Publish exit has been deployed in addition to a final subscriber. The publication is processed by the Publish exit before it is finally received by the final subscriber at Sublevel 1. The intercepting subscribers and the publish exit are shown with broken outlines.

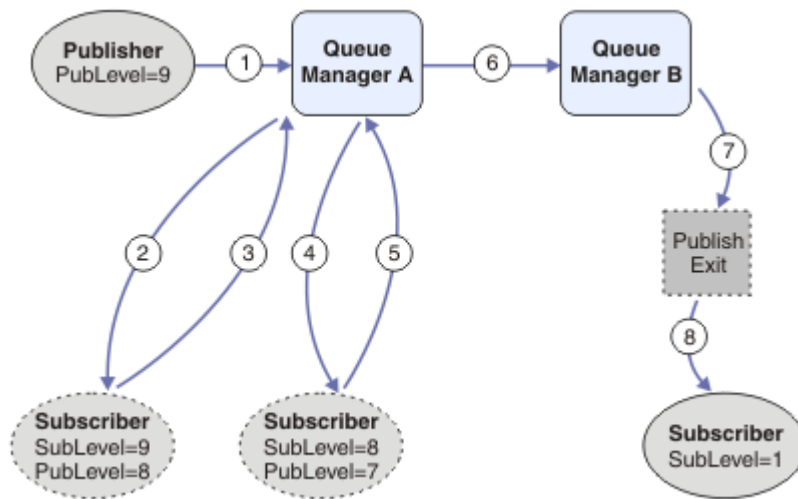


Figure 66. Interception and Publish exit in a cluster

The objective of the simple pattern is for every subscriber receiving a publication to receive the identical publication. The publication goes through the same sequence of transformations regardless of where the subscriber is connected. You probably want to avoid having the sequence of transformations varying, depending on where the publishers or final subscribers are connected. A reasonable exception would be to tailor the publication finally delivered to each individual subscriber. Use the Publish exit to do customize the publication based on the queue to which the publication is finally delivered.

You must consider carefully where to deploy intercepting subscribers and Publish exits in a distributed publish/subscribe topology. The straightforward pattern deploys intercepting subscribers to the same queue manager as the publishers, and Publish exits to the same queue managers as the final subscribers.

Anti-pattern

Figure 67 on page 304 shows how matters can go awry, if you do not follow a simple pattern. To complicate the deployment, a final subscriber is added to queue manager A and two additional intercepting subscribers are added to queue manager B.

The publication is forwarded to queue manager B at PubLevel 7, where it is intercepted by a subscriber at SubLevel 5 before being consumed by the final subscriber at SubLevel 1. The Publish exit intercepts the publication before it is passed to both the intercepting consumer and the final consumer at queue manager B. The publication reaches the final subscriber on queue manager A without being processed by the Publish exit.

In a publish/subscribe topology, proxy subscribers subscribe at SubLevel 1, and pass on the PubLevel set by the last intercepting subscriber. In Figure 67 on page 304, the result is that the publication is not intercepted by the subscriber using SubLevel 9 at queue manager B.

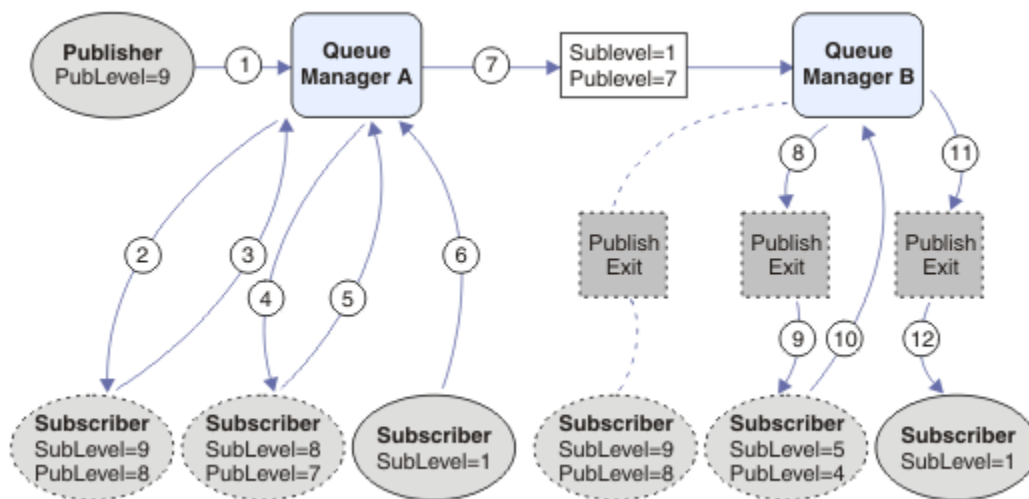


Figure 67. Complex deployment of intercepting subscribers

Publishing options

Several options are available that control the way messages are published.

Withholding reply-to information from subscribers

If you do not want subscribers to be able to reply to publications they receive, it is possible to withhold information in the ReplyToQ and ReplyToQmgr fields of the MQMD by using the MQPMO_SUPPRESS_REPLYTO put-message option. If this option is used, the queue manager removes that information from the MQMD when it receives the publication before forwarding it to any subscribers.

This option cannot be used in combination with a report option that needs a ReplyToQ, if this is attempted the call will fail with MQRC_MISSING_REPLY_TO_Q.

Publication level

Using publication levels is a way of controlling which subscribers receive the publication. The publication level denotes the level of subscription targeted by the publication. Only subscriptions with the highest subscription level less than or equal to the publication's publication level, will receive the publication. This value must be in the range zero to nine; zero is the lowest publication level. The initial value of this field is 9. One of the uses of publication and subscription levels is to [intercept publications](#).

Checking if a publication is not delivered to any subscribers

To check if a publication has not been delivered to any subscribers, use the MQPMO_WARN_IF_NO_SUBS_MATCHED put-message option with the MQPUT call. If a completion code of MQCC_WARNING and a reason code MQRC_NO_SUBS_MATCHED are returned by the put operation, the publication was not delivered to any subscriptions. If the MQPMO_RETAIN option is specified on the put operation, the message is retained and delivered to any subsequently defined matching subscription. In a distributed publish/subscribe system, the MQRC_NO_SUBS_MATCHED reason code is returned only if there are no proxy subscriptions registered for the topic on the queue manager.

Subscription options

Several options are available that control the way message subscriptions are handled.

Message persistence

Queue managers maintain the persistence of the publications they forward to subscribers as set by the publisher. The publisher sets the persistence to be one of the following options:

- 0 Nonpersistent
- 1 Persistent
- 2 Persistence as queue/topic definition

For publish/subscribe, the publisher resolves the topic object and **topicString** to a resolved topic object. If the publisher specifies Persistence as queue/topic definition, then the default persistence from the resolved topic object is set for the publication.

Retained publications

To control when retained publications are received, subscribers can use two subscription options:

Publish on request only, MQSO_PUBLICATIONS_ON_REQUEST

If you want a subscriber to have control of when it receives publications you can use the MQSO_PUBLICATIONS_ON_REQUEST subscription option. A subscriber can then control when it receives publications by using the MQSUBRQ call (specifying the Hsub handle that was returned from the original MQSUB call) to request that it is sent a topic's retained publication. Subscribers using the MQSO_PUBLICATIONS_ON_REQUEST subscription option, do not receive any non-retained publications.

If you specify MQSO_PUBLICATIONS_ON_REQUEST you must use MQSUBRQ to retrieve any publication. If you do not use MQSO_PUBLICATIONS_ON_REQUEST you get messages as they are published.

If a subscriber uses the MQSUBRQ call and uses wildcards in the subscription's topic, the subscription might match multiple topics or nodes on a topic tree, all of which with retained messages (if any exist) will be sent to the subscriber.

This option can be particularly helpful when used with durable subscriptions because a queue manager will continue to send publications to a subscriber if it subscribed durably even if that subscriber application is not running. This could lead to a buildup of messages on the subscriber queue. This build up can be avoided if the subscriber registers using the MQSO_PUBLICATIONS_ON_REQUEST option. Alternatively, you can use non-durable subscriptions if appropriate to your application to avoid a build up of unwanted messages.

If a subscription is durable and a publisher uses retained publications the subscriber application can use the MQSUBRQ call to refresh its state information after a restart. The subscriber must then refresh its state periodically using the MQSUBRQ call.

No publications will be sent as a result of the MQSUB call using this option. A durable subscription that has been resumed following disconnection will use the MQSO_PUBLICATIONS_ON_REQUEST option if the original subscription was configured to use this option.

New publications only, MQSO_NEW_PUBLICATIONS_ONLY

If a retained publication exists on a topic, any subscribers that make a subscription after the publication was made will receive a copy of that publication. If a subscriber does not want to receive any publications that were made earlier than the subscription being made, the subscriber can use the MQSO_NEW_PUBLICATIONS_ONLY subscription option.

Grouping subscriptions

Consider grouping subscriptions if you have set up a queue to receive publications and have a number of overlapping subscriptions feeding publications to the same queue. This situation is similar to the example in [Overlapping subscriptions](#).

You can avoid receiving duplicate publications by setting the option MQSO_GROUP_SUB when you subscribe to a topic. The result is that when more than one subscription in the group matches the topic

of a publication, only one subscription is responsible for placing the publication on the queue. The other subscriptions that matched the publication topic are ignored.

The subscription responsible for placing the publication on the queue is chosen on the basis that it has the longest matching topic string, before encountering any wildcards. It can be thought of as the closest matching subscription. Its properties are propagated to the publication, including whether it has the MQSO_NOT_OWN_PUBS property. If it does, no publication is delivered to the queue, even though other matching subscriptions might not have the MQSO_NOT_OWN_PUBS property.

You cannot place all your subscriptions in a single group to eliminate duplicate publications. Grouped subscriptions must fulfill these conditions:

1. None of the subscriptions are managed.
2. A group of subscriptions deliver publications to the same queue.
3. Each subscription must be at the same subscription level.
4. The publication message for each subscription in the group has the same correlation identifier.

To ensure each subscription results in a publication message with the same correlation identifier, set MQSO_SET_CORREL_ID to create your own correlation identifier in the publication, and set the same value in the **SubCorrelId** field in each subscription. Do not set **SubCorrelId** to the value MQCI_NONE.

See [../com.ibm.mq.ref.dev.doc/q100080_.dita#q100080_/mqso_group_sub](http://com.ibm.mq.ref.dev.doc/q100080_.dita#q100080_/mqso_group_sub) for more information.


Inquiring about and setting object attributes

Attributes are the properties that define the characteristics of a WebSphere MQ object.

They affect the way that a queue manager processes an object. The attributes of each type of WebSphere MQ object are described in detail in [Attributes of objects](#).

Some attributes are set when the object is defined, and can be changed only by using the WebSphere MQ commands; an example of such an attribute is the default priority for messages put on a queue. Other attributes are affected by the operation of the queue manager and can change over time; an example is the current depth of a queue.

You can inquire about the current values of most attributes using the MQINQ call. The MQI also provides an MQSET call with which you can change some queue attributes. You cannot use the MQI calls to change the attributes of any other type of object; instead you must use:

 **For WebSphere MQ for Windows, UNIX and Linux platforms**
The MQSC facility, described in [MQSC reference](#).

Note: The names of the attributes of objects are shown in this documentation in the form that you use them with the MQINQ and MQSET calls. When you use WebSphere MQ commands to define, alter, or display the attributes, you must identify the attributes using the keywords shown in the descriptions of the commands in the topic links.

Both the MQINQ and the MQSET calls use arrays of selectors to identify those attributes that you want to inquire about or set. There is a selector for each attribute that you can work with. The selector name has a prefix, determined by the nature of the attribute:

MQCA_	These selectors refer to attributes that contain character data (for example, the name of a queue).
MQIA_	These selectors refer to attributes that contain either numeric values (such as <i>CurrentQueueDepth</i> , the number of messages on a queue) or a constant value (such as <i>SyncPoint</i> , whether the queue manager supports syncpoints).

Before you use the MQINQ or MQSET calls your application must be connected to the queue manager, and you must use the MQOPEN call to open the object for setting or inquiring about attributes. These

operations are described in [“Connecting to and disconnecting from a queue manager” on page 198](#) and [“Opening and closing objects” on page 206](#).

Use the following links to find out more about getting inquiring about and setting object attributes:

- [“Inquiring about the attributes of an object” on page 307](#)
- [“Some cases where the MQINQ call fails” on page 308](#)
- [“Setting queue attributes” on page 308](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Inquiring about the attributes of an object

Use the MQINQ call to inquire about the attributes of any type of IBM WebSphere MQ.

As input to this call, you must supply:

- A connection handle.
- An object handle.
- The number of selectors.
- An array of attribute selectors, each selector having the form MQCA_* or MQIA_*. Each selector represents an attribute with a value that you want to inquire about, and each selector must be valid for the type of object that the object handle represents. You can specify selectors in any order.
- The number of integer attributes that you are inquiring about. Specify zero if you are not inquiring about integer attributes.
- The length of the character attributes buffer in *CharAttrLength*. This must be at least the sum of the lengths required to hold each character attribute string. Specify zero if you are not inquiring about character attributes.

The output from MQINQ is:

- A set of integer attribute values copied into the array. The number of values is determined by *IntAttrCount*. If either *IntAttrCount* or *SelectorCount* is zero, this parameter is not used.
- The buffer in which character attributes are returned. The length of the buffer is given by the *CharAttrLength* parameter. If either *CharAttrLength* or *SelectorCount* is zero, this parameter is not used.

- A completion code. If the completion code gives a warning, this means that the call completed only partially. In this case, examine the reason code.
- A reason code. There are three partial-completion situations:
 - The selector does not apply to the queue type
 - There is not enough space allowed for integer attributes
 - There is not enough space allowed for character attributes

If more than one of these situations arise, the first one that applies is returned.

If you open a queue for output or inquire and it resolves to a non-local cluster queue you can only inquire the queue name, queue type, and common attributes. The values of the common attributes are those of the chosen queue if `MQOO_BIND_ON_OPEN` was used. The values are those of an arbitrary one of the possible cluster queues if either `MQOO_BIND_NOT_FIXED` or `MQOO_BIND_ON_GROUP` was used or `MQOO_BIND_AS_Q_DEF` was used and the *DefBind* queue attribute was `MQBND_BIND_NOT_FIXED`. See “MQOPEN and clusters” on page 332 and [MQOPEN](#) for more information.

Note: The values returned by the call are a snapshot of the selected attributes. The attributes can change before your program acts on the returned values.

There is a description of the `MQINQ` call in [MQINQ](#).

Some cases where the MQINQ call fails

If you open an alias to inquire about its attributes, you are returned the attributes of the alias queue (the WebSphere MQ object used to access another queue), not those of the base queue.

However, the definition of the base queue to which the alias resolves is also opened by the queue manager, and if another program changes the usage of the base queue in the interval between your `MQOPEN` and `MQINQ` calls, your `MQINQ` call fails and returns the `MQRC_OBJECT_CHANGED` reason code. The call also fails if the attributes of the alias queue object are changed.

Similarly, when you open a remote queue to inquire about its attributes, you are returned the attributes of the local definition of the remote queue only.

If you specify one or more selectors that are not valid for the type of queue attributes upon which you are inquiring, the `MQINQ` call completes with a warning and sets the output as follows:

- For integer attributes, the corresponding elements of *IntAttrs* are set to `MQIAV_NOT_APPLICABLE`.
- For character attributes, the corresponding portions of the *CharAttrs* string are set to asterisks.

If you specify one or more selectors that are not valid for the type of object attributes upon which you are inquiring, the `MQINQ` call fails and returns the `MQRC_SELECTOR_ERROR` reason code.

You cannot call `MQINQ` to look at a model queue; use either the MQSC facility or the commands available on your platform.

Setting queue attributes

Use this information to learn how to set queue attributes using the `MQSET` call.

You can set only the following queue attributes using the `MQSET` call:

- *InhibitGet* (but not for remote queues)
- *DistList* (not on z/OS)
- *InhibitPut*
- *TriggerControl*
- *TriggerType*
- *TriggerDepth*
- *TriggerMsgPriority*
- *TriggerData*

The MQSET call has the same parameters as the MQINQ call. However, for MQSET, all parameters except the completion code and reason code are input parameters. There are no partial-completion situations.

Note: You cannot use the MQI to set the attributes of WebSphere MQ objects other than locally defined queues.

For more details about the MQSET call, see [MQSET](#).

Committing and backing out units of work

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

The following terms are used in this topic:

- Commit
- Back out
- Syncpoint coordination
- Syncpoint
- Unit of work
- Single-phase commit
- Two-phase commit

If you are familiar with these transaction processing terms, you can skip to [“Syncpoint considerations in IBM WebSphere MQ applications”](#) on page 310.

Commit and back out

When a program puts a message on a queue within a unit of work, that message is made visible to other programs only when the program commits the unit of work. To commit a unit of work, all updates must be successful to preserve data integrity. If the program detects an error and decides that the put operation is not permanent, it can back out the unit of work. When a program performs a backout, IBM WebSphere MQ restores the queue by removing the messages that were put on the queue by that unit of work. The way in which the program performs the commit and back out operations depends on the environment in which the program is running.

Similarly, when a program gets a message from a queue within a unit of work, that message remains on the queue until the program commits the unit of work, but the message is not available to be retrieved by other programs. The message is permanently deleted from the queue when the program commits the unit of work. If the program backs out the unit of work, IBM WebSphere MQ restores the queue by making the messages available to be retrieved by other programs.

Syncpoint coordination, syncpoint, unit of work

Syncpoint coordination is the process by which units of work are either committed or backed out with data integrity.

The decision to commit or back out the changes is taken, in the simplest case, at the end of a transaction. However, it can be more useful for an application to synchronize data changes at other logical points within a transaction. These logical points are called *syncpoints* (or *synchronization points*) and the period of processing a set of updates between two syncpoints is called a *unit of work*. Several MQGET calls and MQPUT calls can be part of a single unit of work. The maximum number of messages within a unit of work can be controlled by the MAXUMSGS attribute of the ALTER QMGR command on other platforms, except z/OS. See the [MQSC reference](#) *WebSphere MQ Script (MQSC) Command Reference* for details of these commands.

Single-phase commit

A *single-phase commit* process is one in which a program can commit updates to a queue without coordinating its changes with other resource managers.

Two-phase commit

A *two-phase commit* process is one in which updates that a program has made to IBM WebSphere MQ queues can be coordinated with updates to other resources (for example, databases under the control of DB2). Under such a process, updates to all resources are committed or backed out together.

To help handle units of work, IBM WebSphere MQ provides the *BackoutCount* attribute. This is incremented each time that a message within a unit of work is backed out. If the message repeatedly causes the unit of work to abnormally end, the value of the *BackoutCount* finally exceeds that of the *BackoutThreshold*. This value is set when the queue is defined. In this situation, the application can remove the message from the unit of work and put it onto another queue, as defined in *BackoutRequeueQName*. When the message is moved, the unit of work can commit.

Use the following links to find out more about committing and backing out units of work:

- [“Syncpoint considerations in IBM WebSphere MQ applications” on page 310](#)
- [“Syncpoints in IBM WebSphere MQ in UNIX, Linux, and Windows systems” on page 311](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Syncpoint considerations in IBM WebSphere MQ applications

Use this information to learn about using syncpoints in IBM WebSphere MQ applications.

Two-phase commit is supported under:

- WebSphere MQ for AIX
- WebSphere MQ for HP-UX
- WebSphere MQ for Linux
- WebSphere MQ for Solaris
- WebSphere MQ for Windows
- CICS for MVS/ESA 4.1
- CICS Transaction Server for z/OS
- TXSeries
- IMS/ESA®
- Other external coordinators using the X/Open XA interface

Single-phase commit is supported under:

- WebSphere MQ on UNIX systems
- WebSphere MQ for Windows

Note: For further details on external interfaces see [“Interfaces to external syncpoint managers”](#) on page 313, and the XA documentation *CAE Specification Distributed Transaction Processing: The XA Specification*, published by The Open Group. Transaction managers (such as CICS, IMS, Encina, and Tuxedo) can participate in two-phase commit, coordinated with other recoverable resources. This means that the queuing functions provided by WebSphere MQ can be brought within the scope of a unit of work, managed by the transaction manager.

Samples shipped with WebSphere MQ show WebSphere MQ coordinating XA-compliant databases. For further information about these samples, see [“Sample programs for distributed platforms”](#) on page 92.

In your WebSphere MQ application, you can specify on every put and get call whether you want the call to be under syncpoint control. To make a put operation operate under syncpoint control, use the MQPMO_SYNCPOINT value in the *Options* field of the MQPMO structure when you call MQPUT. For a get operation, use the MQGMO_SYNCPOINT value in the *Options* field of the MQGMO structure. If you do not explicitly choose an option, the default action depends on the platform. The syncpoint control default is no.

When an MQPUT1 call is issued with MQPMO_SYNCPOINT, the default behavior changes, so that the put operation is completed asynchronously. This might cause a change in the behavior of some applications that rely on certain fields in the MQOD and MQMD structures being returned, but which now contain undefined values. An application can specify MQPMO_SYNC_RESPONSE to ensure that the put operation is performed synchronously and that all of the appropriate field values are completed.

When your application receives an MQRC_BACKED_OUT reason code in response to an MQPUT or MQGET under syncpoint, the application should normally back out the current transaction using MQBACK and then, if appropriate, try the entire transaction again. If the application receives MQRC_BACKED_OUT in response to an MQCMIT or MQDISC call, it does not need to call MQBACK.

Every time an MQGET call is backed out, the *BackoutCount* field of the MQMD structure of the affected message is incremented. A high *BackoutCount* indicates a message that has been repeatedly backed out. This might indicate a problem with this message, which you should investigate. See [BackoutCount](#) for details of *BackoutCount*.

If a program issues the MQDISC call while there are uncommitted requests, an implicit syncpoint occurs. If the program ends abnormally, an implicit backout occurs.

Changes to queue attributes (either by the MQSET call or by commands) are not affected by the committing or backing out of units of work.

Syncpoints in IBM WebSphere MQ in UNIX, Linux, and Windows systems

Syncpoint support operates on two types of units of work: local and global.

A *local* unit of work is one in which the only resources updated are those of the WebSphere MQ queue manager. Here syncpoint coordination is provided by the queue manager itself using a single-phase commit procedure.

A *global* unit of work is one in which resources belonging to other resource managers, such as databases, are also updated. WebSphere MQ can coordinate such units of work itself. They can also be coordinated by an external commitment controller such as another transaction manager or the IBM i commitment controller.

For full integrity, use a two-phase commit procedure. Two-phase commit can be provided by XA-compliant transaction managers and databases such as IBM's TXSeries and UDB. WebSphere MQ products (except WebSphere MQ for IBM i and WebSphere MQ for z/OS) can coordinate global units of work using a two-phase commit process.

Local units of work

Units of work that involve only the queue manager are called *local* units of work. Syncpoint coordination is provided by the queue manager itself (internal coordination) using a single-phase commit process.

To start a local unit of work, the application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate syncpoint option. The unit of work is committed using MQCMIT or rolled back using MQBACK. However, the unit of work also ends when the connection between the application and the queue manager is broken, intentionally or unintentionally.

If an application disconnects (MQDISC) from a queue manager while a global unit of work coordinated by WebSphere MQ is still active, an attempt is made to commit the unit of work. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

Global units of work

Use global units of work when you also need to include updates to resources belonging to other resource managers.

Here the coordination can be internal or external to the queue manager:

Internal syncpoint coordination

Queue manager coordination of global units of work is not supported by WebSphere MQ for IBM i or WebSphere MQ for z/OS. It is not supported in a WebSphere MQ MQI client environment.

Here, WebSphere MQ does the coordination. To start a global unit of work, the application issues the MQBEGIN call.

As input to the MQBEGIN call, you must supply the connection handle (*Hconn*) that is returned by the MQCONN or MQCONNX call. This handle represents the connection to the WebSphere MQ queue manager.

The application issues MQGET, MQPUT, or MQPUT1 requests specifying the appropriate syncpoint option. This means that you can use MQBEGIN to initiate a global unit of work that updates local resources, resources belonging to other resource managers, or both. Updates made to resources belonging to other resource managers are made using the API of that resource manager. However, you cannot use the MQI to update queues that belong to other queue managers. Issue MQCMIT or MQBACK before starting further units of work (local or global).

The global unit of work is committed using MQCMIT; this initiates a two-phase commit of all the resource managers involved in the unit of work. A two-phase commit process is used whereby resource managers (for example, XA-compliant database managers such as DB2, Oracle, and Sybase) are first all asked to prepare to commit. Only if all are prepared are they asked to commit. If any resource manager signals that it cannot commit, each is asked to back out instead. Alternatively, you can use MQBACK to roll back the updates of all the resource managers.

If an application disconnects (MQDISC) while a global unit of work is still active, the unit of work is committed. If, however, the application terminates without disconnecting, the unit of work is rolled back as the application is deemed to have terminated abnormally.

The output from MQBEGIN is a completion code and a reason code.

When you use MQBEGIN to start a global unit of work, all the external resource managers that have been configured with the queue manager are included. However, the call starts a unit of work but completes with a warning if:

- There are no participating resource managers (that is, no resource managers have been configured with the queue manager)

or

- One or more resource managers are not available.

In these cases, the unit of work must include updates to only those resource managers that were available when the unit of work was started.

If one of the resource managers cannot commit its updates, all the resource managers are instructed to roll back their updates, and MQCMIT completes with a warning. In unusual circumstances (typically, operator intervention), an MQCMIT call might fail if some resource managers commit their updates but others roll them back; the work is deemed to have completed with a *mixed* outcome. Such occurrences are diagnosed in the error log of the queue manager so that remedial action may be taken.

An MQCMIT of a global unit of work succeeds if all the resource managers involved commit their updates.

For a description of the MQBEGIN call, see [MQBEGIN](#).

External syncpoint coordination

This occurs when a syncpoint coordinator other than WebSphere MQ has been selected; for example, CICS, Encina, or Tuxedo.

In this situation, WebSphere MQ on UNIX and Linux systems and WebSphere MQ for Windows register their interest in the outcome of the unit of work with the syncpoint coordinator so that they can commit or roll back any uncommitted get or put operations as required. The external syncpoint coordinator determines whether one- or two-phase commitment protocols are provided.

When you use an external coordinator, MQCMIT, MQBACK, and MQBEGIN cannot be issued. Calls to these functions fail with the reason code MQRC_ENVIRONMENT_ERROR.

The way in which an externally coordinated unit of work is started depends on the programming interface provided by the syncpoint coordinator. An explicit call might be required. If an explicit call is required, and you issue an MQPUT call specifying the MQPMO_SYNCPOINT option when a unit of work is not started, the completion code MQRC_SYNCPOINT_NOT_AVAILABLE is returned.

The scope of the unit of work is determined by the syncpoint coordinator. The state of the connection between the application and the queue manager affects the success or failure of MQI calls that an application issues, not the state of the unit of work. An application can, for example, disconnect and reconnect to a queue manager during an active unit of work and perform further MQGET and MQPUT operations inside the same unit of work. This is known as a pending disconnect.

You can use WebSphere MQ API calls in CICS programs, whether you choose to use the XA abilities of CICS. If you do not use XA, then the puts and gets of messages to and from queues will not be managed within CICS atomic units of work. One reason for choosing this method is that the overall consistency of the unit of work is not important to you.

If the integrity of your units of work is important to you, then you must use XA. When you use XA, CICS uses a two-phase commit protocol to ensure all resources within the unit of work are updated together.

For more information about setting up transactional support, see [“Transactional support scenarios” on page 39](#), and also TXSeries CICS documentation, for example, *TXSeries for Multiplatforms CICS Administration Guide for Open Systems*.

Interfaces to external syncpoint managers

WebSphere MQ on UNIX and Linux systems, and WebSphere MQ for Windows support coordination of transactions by external syncpoint managers that use the X/Open XA interface.

Some XA transaction managers (TXSeries) require that each XA resource manager supplies its name. This is the string called name in the XA switch structure. The resource manager for WebSphere MQ on UNIX, Linux and Windows systems is named MQSeries_XA_RMI. For further details on XA interfaces refer to the XA documentation *CAE Specification Distributed Transaction Processing: The XA Specification*, published by The Open Group.

In an XA configuration, WebSphere MQ on UNIX, Linux, and Windows systems fulfill the role of an XA Resource Manager. An XA syncpoint coordinator can manage a set of XA Resource Managers, and

synchronize the commit or backout of transactions in both Resource Managers. This is how it works for a statically-registered resource manager:

1. An application notifies the syncpoint coordinator that it wants to start a transaction.
2. The syncpoint coordinator issues a call to any resource managers that it knows of, to notify them of the current transaction.
3. The application issues calls to update the resources managed by the resource managers associated with the current transaction.
4. The application requests that the syncpoint coordinator either commit or roll back the transaction.
5. The syncpoint coordinator issues calls to each resource manager using two-phase commit protocols to complete the transaction as requested.

The XA specification requires each Resource Manager to provide a structure called an *XA Switch*. This structure declares the capabilities of the Resource Manager, and the functions that are to be called by the syncpoint coordinator.

There are two versions of this structure:

MQRMIXASwitch	Static XA resource management
MQRMIXASwitchDynamic	Dynamic XA resource management

For a list of the libraries containing this structure see [“The IBM WebSphere MQ XA switch structure” on page 67](#).

The method that must be used to link them to an XA syncpoint coordinator is defined by the coordinator; consult the documentation provided by that coordinator to determine how to enable WebSphere MQ to cooperate with your XA syncpoint coordinator.

The *xa_info* structure that is passed on any *xa_open* call by the syncpoint coordinator can be the name of the queue manager that is to be administered. This takes the same form as the queue manager name passed to MQCONN or MQCONNX, and can be blank if the default queue manager is to be used. However, you can use the two extra parameters TPM and AXLIB

TPM allows you to specify to WebSphere MQ the transaction manager name, for example, CICS. AXLIB allows you to specify the actual library name in the transaction manager where the XA AX entry points are located.

If you use either of these parameters or a non default queue manager you must specify the queue manager name using the QMNAME parameter. For further information see [The CHANNEL, TRPTYPE, CONNAME, and QMNAME parameters of the xa_open string](#).

Restrictions

1. Global units of work are not allowed with a shared Hconn (as described in [“Shared \(thread independent\) connections with MQCONNX” on page 204](#)).
2. On Windows systems, all functions declared in the XA switch are declared as `_cdecl` functions.
3. An external syncpoint coordinator can administer only one queue manager at a time. This is because the coordinator has an effective connection to each queue manager, and is therefore subject to the rule that only one connection is allowed at a time.

Note: Note: A JMS client application (CLIENT JEE application) running in a JEE server does not have this restriction, so a single JEE server-managed transaction can coordinate multiple queue managers in the same transaction. However, a JMS server application, running in bindings mode, is still subject to the rule that only one connection is allowed at a time.

4. All applications that are run using the syncpoint coordinator can connect only to the queue manager that is administered by the coordinator because they are already effectively connected to that queue manager. They must issue MQCONN or MQCONNX to obtain a connection handle and must issue MQDISC before they exit. Alternatively, they can use the exit UE014015 for TXSeries CICS.

Starting IBM WebSphere MQ applications using triggers

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

Some WebSphere MQ applications that serve queues run continuously, so they are always available to retrieve messages that arrive on the queues. However, you might not want this when the number of messages arriving on the queues is unpredictable. In this case, applications could be consuming system resources even when there are no messages to retrieve.

WebSphere MQ provides a facility that enables an application to be started automatically when there are messages available to retrieve. This facility is known as *triggering*.

For information about triggering channels see [Triggering channels](#).

What is triggering?

The queue manager defines certain conditions as constituting *trigger events*.

If triggering is enabled for a queue and a trigger event occurs, the queue manager sends a *trigger message* to a queue called an *initiation queue*. The presence of the trigger message on the initiation queue indicates that a trigger event has occurred.

Trigger messages generated by the queue manager are not persistent. This reduces logging (resulting in improving performance), and minimizing duplicates during restart, so improving restart time.

The program that processes the initiation queue is called a *trigger-monitor application*, and its function is to read the trigger message and take appropriate action, based on the information contained in the trigger message. Typically this action is to start some other application to process the queue that generated the trigger message. From the point of view of the queue manager, there is nothing special about the trigger-monitor application; it is simply another application that reads messages from a queue (the initiation queue).

If triggering is enabled for a queue, you can create a *process-definition object* associated with it. This object contains information about the application that processes the message that caused the trigger event. If the process definition object is created, the queue manager extracts this information and places it in the trigger message, for use by the trigger-monitor application. The name of the process definition associated with a queue is given by the *ProcessName* local-queue attribute. Each queue can specify a different process definition, or several queues can share the same process definition.

If you want to trigger the start of a channel, you do not need to define a process definition object. The transmission queue definition is used instead.

Triggering is supported by WebSphere MQ clients running in the following environments:

- UNIX and Linux systems
- Windows systems

An application running in a client environment is the same as one running in a full WebSphere MQ environment, except that you link it with the client libraries. However the trigger monitor and the application to be started must both be in the same environment.

Triggering involves:

Application queue

An *application queue* is a local queue that, when it has triggering set on and when the conditions are met, requires that trigger messages are written.

Process definition

An application queue can have a *process definition object* associated with it that holds details of the application that will get messages from the application queue. (See [Attributes for process definitions](#) for a list of attributes.)

Remember that if you want a trigger to start a channel, you do not need to define a process definition object.

Transmission queue

You need a transmission queue if you want a trigger to start a channel.

For a transmission queue on AIX, HP-UX, IBM i, Solaris, z/OS, or Windows systems, the *TriggerData* attribute of the transmission queue can specify the name of the channel to be started. This can replace the process definition for triggering channels, but is used only when a process definition is not created.

Trigger event

A *trigger event* is an event that causes a trigger message to be generated by the queue manager. This is typically a message arriving on an application queue, but it can also occur at other times (see [“Conditions for a trigger event”](#) on page 320). WebSphere MQ has a range of options to allow you to control the conditions that cause a trigger event (see [“Controlling trigger events”](#) on page 324).

Trigger message

The queue manager creates a *trigger message* when it recognizes a trigger event (see [“Conditions for a trigger event”](#) on page 320). It copies into the trigger message information about the application to be started. This information comes from the application queue and the process definition object associated with the application queue. Trigger messages have a fixed format (see [“Format of trigger messages”](#) on page 330).

Initiation queue

An *initiation queue* is a local queue on which the queue manager puts trigger messages. Note, that an initiation queue cannot be an alias queue or a model queue. A queue manager can own more than one initiation queue, and each one is associated with one or more application queues. A shared queue, a local queue accessible by queue managers in a queue-sharing group, can be an initiation queue on WebSphere MQ for z/OS.

Trigger monitor

A *trigger monitor* is a continuously running program that serves one or more initiation queues. When a trigger message arrives on an initiation queue, the trigger monitor retrieves the message. The trigger monitor uses the information in the trigger message. It issues a command to start the application that is to retrieve the messages arriving on the application queue, passing it information contained in the trigger message header, which includes the name of the application queue.

On all platforms, a special trigger monitor known as the channel initiator is responsible for starting channels. On z/OS, the channel initiator is typically started manually, or it can be done automatically when a queue manager starts by changing CSQINP2 in the queue manager startup JCL. On other platforms, it is automatically started when the queue manager starts or it can be started manually with the runmqchi command.

(For more information, see [“Initiation queue processing by trigger monitors”](#) on page 327.)

To understand how triggering works, consider [Figure 68 on page 317](#), which is an example of trigger type FIRST (MQTT_FIRST).




Figure 68. Flow of application and trigger messages

In Figure 68 on page 317, the sequence of events is:

1. Application A, which can be either local or remote to the queue manager, puts a message on the application queue. No application has this queue open for input. However, this fact is relevant only to trigger type FIRST and DEPTH.
2. The queue manager checks to see if the conditions are met under which it has to generate a trigger event. They are, and a trigger event is generated. Information held within the associated process definition object is used when creating the trigger message.
3. The queue manager creates a trigger message and puts it on the initiation queue associated with this application queue, but only if an application (trigger monitor) has the initiation queue open for input.
4. The trigger monitor retrieves the trigger message from the initiation queue.
5. The trigger monitor issues a command to start application B (the server application).
6. Application B opens the application queue and retrieves the message.

Note:

1. If the application queue is open for input, by any program, and has triggering set for FIRST or DEPTH, no trigger event occurs because the queue is already being served.
2. If the initiation queue is not open for input, the queue manager does not generate any trigger messages; it waits until an application opens the initiation queue for input.
3. When using triggering for channels, use trigger type FIRST or DEPTH.
4. Triggered applications run under the user ID and group of the user who started the trigger monitor, the CICS user, or the user who started the queue manager.

So far, the relationship between the queues within triggering has been only on a one to one basis. Consider Figure 69 on page 318.




Figure 69. Relationship of queues within triggering

An application queue has a process definition object associated with it that holds details of the application that will process the message. The queue manager places the information in the trigger message, so only one initiation queue is necessary. The trigger monitor extracts this information from the trigger message and starts the relevant application to deal with the message on each application queue.

Remember that, if you want to trigger the start of a channel, you do not need to define a process definition object. The transmission queue definition can determine the channel to be triggered.

Use the following links to find out more about starting WebSphere MQ applications using triggers:

- [“Prerequisites for triggering” on page 319](#)
- [“Conditions for a trigger event” on page 320](#)
- [“Controlling trigger events” on page 324](#)
- [“Designing an application that uses triggered queues” on page 326](#)

- [“Initiation queue processing by trigger monitors” on page 327](#)
- [“Properties of trigger messages” on page 329](#)
- [“When triggering does not work” on page 331](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Working with the MQI and clusters” on page 331](#)

There are special options on calls and return codes that relate to clustering.

Prerequisites for triggering

Use this information to learn about the steps to take before using triggering.

Before your application can take advantage of triggering, complete the following steps:

1. Either:

a. Create an initiation queue for your application queue. For example:

```
DEFINE QLOCAL (initiation.queue) REPLACE +
      LIKE (SYSTEM.DEFAULT.INITIATION.QUEUE) +
      DESCR ('initiation queue description')
```

or

b. Determine the name of a local queue that exists and can be used by your application (typically, this name is SYSTEM.DEFAULT.INITIATION.QUEUE or, if you are starting channels with triggers, SYSTEM.CHANNEL.INITQ), and specify its name in the *InitiationQName* field of the application queue.

2. Associate the initiation queue with the application queue. A queue manager can own more than one initiation queue. You might want some of your application queues to be served by different programs, in which case, you can use one initiation queue for each serving program, although you do not have to. Here is an example of how to create an application queue:

```
DEFINE QLOCAL (application.queue) REPLACE +
      LIKE (SYSTEM.DEFAULT.LOCAL.QUEUE) +
      DESCR ('appl queue description') +
      INITQ ('initiation.queue') +
      PROCESS ('process.name') +
      TRIGGER +
      TRIGTYPE (FIRST)
```


3. If you are triggering an application, create a process definition object to contain information relating to the application that is to serve your application queue. For example, to trigger-start a CICS payroll transaction called PAYR:

```
DEFINE PROCESS (process.name) +  
  REPLACE +  
  DESCR ('process description') +  
  APPLICID ('PAYR') +  
  APPLTYPE (CICS) +  
  USERDATA ('Payroll data')
```

When the queue manager creates a trigger message, it copies information from the attributes of the process definition object into the trigger message.

Platform	To create a process definition object
UNIX, Linux, and Windows systems	Use DEFINE PROCESS or use SYSTEM.DEFAULT.PROCESS and modify using ALTER PROCESS

4. Optional: Create a transmission queue definition and use blanks for the *ProcessName* attribute.

The *TrigData* attribute can contain the name of the channel to be triggered or it can be left blank. Except on IBM WebSphere MQ for z/OS, if it is left blank, the channel initiator searches the channel definition files until it finds a channel that is associated with the named transmission queue. When the queue manager creates a trigger message, it copies information from the *TrigData* attribute of the transmission queue definition into the trigger message.

5. If you have created a process definition object to specify properties of the application that is to serve your application queue, associate the process object with your application queue by naming it in the *ProcessName* attribute of the queue.

Platform	Use commands
UNIX, Linux, and Windows systems	ALTER QLOCAL

6. Start instances of the trigger monitors that are to serve the initiation queues you have defined. See [“Initiation queue processing by trigger monitors” on page 327](#) for more information.

If you want to be aware of any undelivered trigger messages, make sure that your queue manager has a dead-letter (undelivered-message) queue defined. Specify the name of the queue in the *DeadLetterQName* queue manager field.

You can then set the trigger conditions that you require, using the attributes of the queue object that defines your application queue. For more information, see [“Controlling trigger events” on page 324](#).

Conditions for a trigger event

References to shared queues in this topic mean shared queues in a queue-sharing group, only available on WebSphere MQ for z/OS.

The queue manager creates a trigger message when the following conditions are satisfied:

1. A message is *put* on a queue.
2. The message has a priority greater than or equal to the threshold trigger priority of the queue. This priority is set in the *TriggerMsgPriority* local queue attribute; if it is set to zero, any message qualifies.
3. The number of messages on the queue with priority greater than or equal to *TriggerMsgPriority* was previously, depending on *TriggerType*:
 - Zero (for trigger type MQTT_FIRST)

- Any number (for trigger type MQTT_EVERY)
- *TriggerDepth* minus 1 (for trigger type MQTT_DEPTH)

Note:

- a. For non-shared local queues, the queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist. Consequently an application might be started when there are no messages for it to retrieve because the messages on the queue have not been committed. In this situation, consider using the wait option with a suitable *WaitInterval*, so that the application waits for its messages to arrive.
 - b. For local shared queues, the queue manager counts committed messages only.
4. For triggering of type FIRST or DEPTH, no program has the application queue open for removing messages (that is, the *OpenInputCount* local queue attribute is zero).

Note:

- a. For shared queues, special conditions apply when multiple queue managers have trigger monitors running against a queue. In this situation, if one or more queue managers have the queue open for input shared, the trigger criteria on the other queue managers are treated as *TriggerType* MQTT_FIRST and *TriggerMsgPriority* zero. When all the queue managers close the queue for input, the trigger conditions revert to those conditions specified in the queue definition.

An example scenario affected by this condition is multiple queue managers QM1, QM2, and QM3 with a trigger monitor running for an application queue A. A message arrives on A satisfying the conditions for triggering, and a trigger message is generated on the initiation queue. The trigger monitor on QM1 gets the trigger message and triggers an application. The triggered application opens the application queue for shared input. From this point on the trigger conditions for application queue A are evaluated as *TriggerType* MQTT_FIRST, and *TriggerMsgPriority* zero on queue managers QM2 and QM3, until QM1 closes the application queue.

- b. For shared queues, this condition is applied for each queue manager. That is, a queue manager's *OpenInputCount* for a queue must be zero for a trigger message to be generated for the queue by that queue manager. However, if any queue manager in the queue-sharing group has the queue open using the MQOO_INPUT_EXCLUSIVE option, no trigger message is generated for that queue by any of the queue managers in the queue-sharing group.

The change in how the trigger conditions are evaluated occurs when the triggered application opens the queue for input. In scenarios where there is only one trigger monitor running, other applications can have the same effect because they similarly open the application queue for input. It does not matter whether the application queue was opened by an application that is started by a trigger monitor, or by some other application; it is the fact that the queue is open for input on another queue manager that causes the change in trigger criteria.

5. On WebSphere MQ for z/OS, if the application queue is one with a *Usage* attribute of MQUS_NORMAL, get requests for it are not inhibited (that is, the *InhibitGet* queue attribute is MQQA_GET_ALLOWED). Also, if the triggered application queue is one with a *Usage* attribute of MQUS_XMITQ, get requests for it are not inhibited.
6. Either:
 - The *ProcessName* local queue attribute for the queue is not blank, and the process definition object identified by that attribute has been created, or
 - The *ProcessName* local queue attribute for the queue is all blank, but the queue is a transmission queue. As the process definition is optional, the *TriggerData* attribute might also contain the name of the channel to be started. In this case, the trigger message contains attributes with the following values:
 - *QName*: queue name
 - *ProcessName*: blanks
 - *TriggerData*: trigger data
 - *AppType*: MQAT_UNKNOWN

- *ApplId*: blanks
 - *EnvData*: blanks
 - *UserData*: blanks
7. An initiation queue has been created, and has been specified in the *InitiationQName* local queue attribute. Also:
 - Get requests are not inhibited for the initiation queue (that is, the *InhibitGet* queue attribute is MQQA_GET_ALLOWED).
 - Put requests must not be inhibited for the initiation queue (that is, the *InhibitPut* queue attribute must be MQQA_PUT_ALLOWED).
 - The *Usage* attribute of the initiation queue must be MQUS_NORMAL.
 - In environments where dynamic queues are supported, the initiation queue must not be a dynamic queue that has been marked as logically deleted.
 8. A trigger monitor currently has the initiation queue open for removing messages (that is, the *OpenInputCount* local queue attribute is greater than zero).
 9. The trigger control (*TriggerControl* local queue attribute) for the application queue is set to MQTC_ON. To do this, set the *trigger* attribute when you define your queue, or use the ALTER QLOCAL command.
 10. The trigger type (*TriggerType* local queue attribute) is not MQTT_NONE.

If all the required conditions are met, and the message that caused the trigger condition is put as part of a unit of work, the trigger message does not become available for retrieval by the trigger monitor application until the unit of work completes, whether the unit of work is committed or, for trigger type MQTT_FIRST or MQTT_DEPTH, backed out.
 11. A suitable message is placed on the queue, for a *TriggerType* of MQTT_FIRST or MQTT_DEPTH, and the queue:
 - Was not previously empty (MQTT_FIRST), or
 - Had *TriggerDepth* or more messages (MQTT_DEPTH)

and conditions “2” on page 320 through “10” on page 322 (excluding “3” on page 320) are satisfied, if in the case of MQTT_FIRST a sufficient interval (*TriggerInterval* queue-manager attribute) has elapsed since the last trigger message was written for this queue.

This is to allow for a queue server that ends before processing all the messages on the queue. The purpose of the trigger interval is to reduce the number of duplicate trigger messages that are generated.

Note: If you stop and restart the queue manager, the *TriggerInterval* timer is reset. There is a small window during which it is possible to produce two trigger messages. The window exists when the trigger attribute of the queue is set to enabled at the same time as a message arrives and the queue was not previously empty (MQTT_FIRST) or had *TriggerDepth* or more messages (MQTT_DEPTH).
 12. The only application serving a queue issues an MQCLOSE call, for a *TriggerType* of MQTT_FIRST or MQTT_DEPTH, and there is at least:
 - One (MQTT_FIRST), or
 - *TriggerDepth* (MQTT_DEPTH)

messages on the queue of sufficient priority (condition “2” on page 320), and conditions “6” on page 321 through “10” on page 322 are also satisfied.

This is to allow for a queue server that issues an MQGET call, finds the queue empty, and so ends; however, in the interval between the MQGET and the MQCLOSE calls, one or more messages arrive.

Note:

- a. If the program serving the application queue does not retrieve all the messages, this can cause a closed loop. Each time that the program closes the queue, the queue manager creates another trigger message that causes the trigger monitor to start the server program again.
 - b. If the program serving the application queue backs out its get request (or if the program abends) before it closes the queue, the same happens. However, if the program closes the queue before backing out the get request, and the queue is otherwise empty, no trigger message is created.
 - c. To prevent such a loop occurring, use the *BackoutCount* field of MQMD to detect messages that are repeatedly backed out. For more information, see [“Messages that are backed out” on page 36](#).
13. The following conditions are satisfied using MQSET or a command:
- a. • *TriggerControl* is changed to MQTC_ON, or
 - *TriggerControl* is already MQTC_ON and the value of either *TriggerType*, *TriggerMsgPriority*, or *TriggerDepth* (if relevant) is changed,
 and there is at least:
 - One (MQTT_FIRST or MQTT_EVERY), or
 - *TriggerDepth* (MQTT_DEPTH)
 messages on the queue of sufficient priority (condition “2” on page 320), and conditions “4” on page 321 through “10” on page 322 (excluding “8” on page 322) are also satisfied.
 This is to allow for an application or operator changing the triggering criteria, when the conditions for a trigger to occur are already satisfied.
 - b. The *InhibitPut* queue attribute of an initiation queue changes from MQQA_PUT_INHIBITED to MQQA_PUT_ALLOWED, and there is at least:
 - One (MQTT_FIRST or MQTT_EVERY), or
 - *TriggerDepth* (MQTT_DEPTH)
 messages of sufficient priority (condition “2” on page 320) on any of the queues for which this is the initiation queue, and conditions “4” on page 321 through “10” on page 322 are also satisfied. (One trigger message is generated for each such queue satisfying the conditions.)
 This is to allow for trigger messages not being generated because of the MQQA_PUT_INHIBITED condition on the initiation queue, but this condition now having been changed.
 - c. The *InhibitGet* queue attribute of an application queue changes from MQQA_GET_INHIBITED to MQQA_GET_ALLOWED, and there is at least:
 - One (MQTT_FIRST or MQTT_EVERY), or
 - *TriggerDepth* (MQTT_DEPTH)
 messages of sufficient priority (condition “2” on page 320) on the queue, and conditions “4” on page 321 through “10” on page 322, excluding “5” on page 321, are also satisfied.
 This allows applications to be triggered only when they can retrieve messages from the application queue.
 - d. A trigger-monitor application issues an MQOPEN call for input from an initiation queue, and there is at least:
 - One (MQTT_FIRST or MQTT_EVERY), or
 - *TriggerDepth* (MQTT_DEPTH)
 messages of sufficient priority (condition “2” on page 320) on any of the application queues for which this is the initiation queue, and conditions “4” on page 321 through “10” on page 322 (excluding “8” on page 322) are also satisfied, and no other application has the initiation queue open for input (one trigger message is generated for each such queue satisfying the conditions).
 This is to allow for messages arriving on queues while the trigger monitor is not running, and for the queue manager restarting and trigger messages (which are nonpersistent) being lost.

14. MSGDLVSQ is set correctly. If you set MSGDLVSQ=FIFO, messages are delivered to the queue in a First In First Out basis. The priority of the message is ignored and the default priority of the queue is assigned to the message. If *TriggerMsgPriority* is set to a higher value than the default priority of the queue, no messages are triggered. If *TriggerMsgPriority* is set equal to or lower than the default priority of the queue, triggering occurs for type FIRST, EVERY, and DEPTH. For information about these types, see the description of the *TriggerType* field under [“Controlling trigger events”](#) on page 324.

If you set MSGDLVSQ=PRIORITY and the message priority is equal to or greater than the *TriggerMsgPriority* field, messages only *count* towards a trigger event. In this case, triggering occurs for type FIRST, EVERY, and DEPTH. As an example, if you put 100 messages of lower priority than the *TriggerMsgPriority*, the effective queue depth for triggering purposes is still zero. If you then put another message on the queue, but this time the priority is greater than or equal to the *TriggerMsgPriority*, the effective queue depth increases from zero to one and the condition for *TriggerType* FIRST is satisfied.

Note:

1. From step [“12”](#) on page 322 (where trigger messages are generated as a result of some event other than a message arriving on the application queue), the trigger message is not put as part of a unit of work. Also, if the *TriggerType* is MQTT_EVERY, and if there are one or more messages on the application queue, only one trigger message is generated.
2. If WebSphere MQ segments a message during MQPUT, a trigger event will not be processed until all the segments have been successfully placed on the queue. However, once message segments are on the queue, WebSphere MQ treats them as individual messages for triggering purposes. For example, a single logical message split into three pieces causes only one trigger event to be processed when it is first MQPUT and segmented. However, each of the three segments causes their own trigger events to be processed as they are moved through the WebSphere MQ network.

Controlling trigger events

You control trigger events using some of the attributes that define your application queue. This information also gives examples of using the trigger types: EVERY, FIRST, and DEPTH.

You can enable and disable triggering, and you can select the number or priority of the messages that count toward a trigger event. There is a full description of these attributes in [Attributes of objects](#).

The relevant attributes are:

TriggerControl

Use this attribute to enable and disable triggering for an application queue.

TriggerMsgPriority

The minimum priority that a message must have for it to count toward a trigger event. If a message of priority less than *TriggerMsgPriority* arrives on the application queue, the queue manager ignores the message when it determines whether to create a trigger message. If *TriggerMsgPriority* is set to zero, all messages count toward a trigger event.

TriggerType

In addition to the trigger type NONE (which disables triggering just like setting the *TriggerControl* to OFF), you can use the following trigger types to set the sensitivity of a queue to trigger events:

EVERY	A trigger event occurs every time that a message arrives on the application queue. Use this type of trigger if you want multiple instances of an application started.
FIRST	A trigger event occurs only when the number of messages on the application queue changes from zero to one. Use this type of trigger if you want a serving program to start when the first message arrives on a queue, continue until there are no more messages to process, then end. You must always process the queue until it is empty. Also see “Special case of trigger type FIRST” on page 325 .

DEPTH

A trigger event occurs only when the number of messages on the application queue reaches the value of the *TriggerDepth* attribute. A typical use of this type of triggering is to start a program when all the replies to a set of requests are received.

Triggering by depth: With triggering by depth, the queue manager disables triggering (using the `<xph><pv>TriggerControl</pv></xph>` attribute) after it creates a trigger message. Your application must re-enable triggering itself (by using the MQSET call) after this has happened.

The action of disabling triggering is not under syncpoint control, so triggering cannot be re-enabled by backing out a unit of work. If a program backs out a put request that caused a trigger event, or if the program abends, you must re-enable triggering by using the MQSET call or the ALTER QLOCAL command.

TriggerDepth

The number of messages on a queue that causes a trigger event when using triggering by depth.

The conditions that must be satisfied for a queue manager to create a trigger message are described in [“Conditions for a trigger event” on page 320](#).

Example of the use of trigger type EVERY

Consider an application that generates requests for motor insurance. The application might send request messages to a number of insurance companies, specifying the same reply-to queue each time. It might set a trigger of type EVERY on this reply-to queue so that each time a reply arrives, the reply might trigger an instance of the server to process the reply.

Example of the use of trigger type FIRST

Consider an organization with a number of branch offices that each transmit details of the days business to the head office. They all do this at the same time, at the end of the working day, and at the head office there is an application that processes the details from all the branch offices. The first message to arrive at the head office could cause a trigger event that starts this application. This application would continue processing until there are no more messages on its queue.

Example of the use of trigger type DEPTH

Consider a travel agency application that creates a single request to confirm a flight reservation, to confirm a reservation for a hotel room, to rent a car, and to order some travelers checks. The application might separate these items into four request messages, sending each to a separate destination. It might set a trigger of type DEPTH on its reply-to queue (with the depth set to the value 4), so that it is restarted only when all four replies have arrived.

If another message (possibly from a different request) arrives on the reply-to queue before the last of the four replies, the requesting application is triggered early. To avoid this, when using DEPTH triggering to collect multiple replies to a request, always use a new reply-to queue for each request.

Special case of trigger type FIRST

With trigger type FIRST, if there is already a message on the application queue when another message arrives, the queue manager does not typically create another trigger message.

However, the application serving the queue might not actually open the queue (for example, the application might end, possibly because of a system problem). If an incorrect application name has been put into the process definition object, the application serving the queue will not pick up any of the messages. In these situations, if another message arrives on the application queue, there is no server running to process this message (and any other messages on the queue).

To deal with this, the queue manager creates further trigger messages under the following circumstances:

- If another message arrives on the application queue, but only if a predefined time interval has elapsed since the queue manager created the last trigger message for that queue. This time interval is defined in the queue manager attribute *TriggerInterval*. Its default value is 999 999 999 milliseconds.
- On WebSphere MQ for z/OS, application queues that name an open initiation queue are scanned periodically. If *TRIGINT* milliseconds have passed since the last trigger message was sent and the queue satisfies the conditions for a trigger event and *CURDEPTH* is greater than zero, a trigger message is generated. This process is called backstop triggering.

Consider the following points when deciding on a value for the trigger interval to use in your application:

- If you set *TriggerInterval* to a low value, and there is no application serving the application queue, trigger type *FIRST* might behave like trigger type *EVERY*. This depends on the rate that messages are being put onto the application queue, which in turn might depend on other system activity. This is because, if the trigger interval is very small, another trigger message is generated each time that a message is put onto the application queue, even though the trigger type is *FIRST*, not *EVERY*. (Trigger type *FIRST* with a trigger interval of zero is equivalent to trigger type *EVERY*.)
- On WebSphere MQ for z/OS if you set *TRIGINT* to a low value, and there is no application serving the trigger type *FIRST* application queue, backstop triggering will generate a trigger message each time the periodic scan of application queues that name open initiation queues takes place.
- If a unit of work is backed out (see [Trigger messages and units of work](#)) and the trigger interval has been set to a high value (or the default value), one trigger message is generated when the unit of work is backed out. However, if you have set the trigger interval to a low value or to zero (causing trigger type *FIRST* to behave like trigger type *EVERY*) many trigger messages can be generated. If the unit of work is backed out, all the trigger messages are still made available. The number of trigger messages that are generated depends on the trigger interval. If the trigger interval is set to zero, the maximum number of messages are generated.

Designing an application that uses triggered queues

You have seen how to set up, and control, triggering for your applications. Here are some tips to consider when you design your application.

Trigger messages and units of work

Trigger messages created because of trigger events that are not part of a unit of work are put on the initiation queue, outside any unit of work, with no dependence on any other messages, and are available for retrieval by the trigger monitor immediately.

Trigger messages created because of trigger events that are part of a unit of work are made available on the initiation queue when the UOW is resolved, whether the unit of work is committed or backed out

If the queue manager fails to put a trigger message on an initiation queue, it will be put on the dead-letter (undelivered-message) queue.

Note:

1. The queue manager counts both committed and uncommitted messages when it assesses whether the conditions for a trigger event exist.

With triggering of type *FIRST* or *DEPTH*, trigger messages are made available even if the unit of work is backed out so that a trigger message is always available when the required conditions are met. For example, consider a put request within a unit of work for a queue that is triggered with trigger type *FIRST*. This causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has now changed from one to two, which does not satisfy the conditions for a trigger event. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created.

However, this means that trigger messages are sometimes created when the conditions for a trigger event are *not* satisfied. Applications that use triggering must always be prepared to handle this situation. It is recommended that you use the wait option with the MQGET call, setting the *WaitInterval* to a suitable value.

Created trigger messages are always made available, whether the unit of work is backed out or committed.

2. For local shared queues (that is, shared queues in a queue-sharing group) the queue manager counts committed messages only.

Getting messages from a triggered queue

When you design applications that use triggering, be aware that there might be a delay between a trigger monitor starting a program and other messages becoming available on the application queue. This can happen when the message that causes the trigger event is committed before the others.

To allow time for messages to arrive, always use the wait option when you use the MQGET call to remove messages from a queue for which trigger conditions are set. The *WaitInterval* must be sufficient to allow for the longest reasonable time between a message being put and that put call being committed. If the message is arriving from a remote queue manager, this time is affected by:

- The number of messages that are put before being committed
- The speed and availability of the communication link
- The sizes of the messages

For an example of a situation where you should use the MQGET call with the wait option, consider the same example that we used when describing units of work. This was a put request within a unit of work for a queue that is triggered with trigger type FIRST. This event causes the queue manager to create a trigger message. If another put request occurs, from another unit of work, this does not cause another trigger event because the number of messages on the application queue has not changed from zero to one. Now if the first unit of work is backed out, but the second is committed, a trigger message is still created. So the trigger message is created at the time that the first unit of work is backed out. If there is a significant delay before the second message is committed, the triggered application might need to wait for it.

With triggering of type DEPTH, a delay can occur even if all relevant messages are eventually committed. Suppose that the *TriggerDepth* queue attribute has the value 2. When two messages arrive on the queue, the second causes a trigger message to be created. However, if the second message is the first to be committed, it is at that time that the trigger message becomes available. The trigger monitor starts the server program, but the program can retrieve only the second message until the first one is committed. So the program might need to wait for the first message to be made available.

Design your application so that it terminates if no messages are available for retrieval when your wait interval expires. If one or more messages arrive later, rely on your application being retriggered to process them. This method prevents applications being idle, and unnecessarily using resources.

Initiation queue processing by trigger monitors

To a queue manager, a trigger monitor is like any other application that serves a queue. However, a trigger monitor serves initiation queues.

A trigger monitor is usually a continuously-running program. When a trigger message arrives on an initiation queue, the trigger monitor retrieves that message. It uses information in the message to issue a command to start the application that is to process the messages on the application queue.

The trigger monitor must pass sufficient information to the program that it is starting so that the program can perform the correct actions on the correct application queue.

A channel initiator is an example of a special type of trigger monitor for message channel agents. In this situation however, you must use either trigger type FIRST or DEPTH.

Trigger monitors on UNIX and Windows systems

This topic contains information about trigger monitors provided on UNIX and Windows systems.

The following trigger monitors are provided for the server environment:

amqstrg0

This is a sample trigger monitor that provides a subset of the function provided by **runmqtrm**. See [“Sample programs for distributed platforms”](#) on page 92 for more information about amqstrg0.

runmqtrm

The syntax of this command is **runmqtrm** [-m *QMgrName*] [-q *InitQ*], where *QMgrName* is the queue manager and *InitQ* is the initiation queue. The default queue is SYSTEM.DEFAULT.INITIATION.QUEUE on the default queue manager. It calls programs for the appropriate trigger messages. This trigger monitor supports the default application type.

The command string passed by the trigger monitor to the operating system is built as follows:

1. The *AppLId* from the relevant PROCESS definition (if created)
2. The MQTMC2 structure, enclosed in double quotation marks
3. The *EnvData* from the relevant PROCESS definition (if created)

where *AppLId* is the name of the program to run as it would be entered on the command line.

The parameter passed is the MQTMC2 character structure. A command string is invoked that has this string, exactly as provided, in double quotation marks, in order that the system command will accept it as one parameter.

The trigger monitor does not look to see if there is another message on the initiation queue until the completion of the application that it has just started. If the application has much processing to do, the trigger monitor might not be able to keep up with the number of trigger messages arriving. You have two options:

- Have more trigger monitors running
- Run the started applications in the background

If you have more trigger monitors running, you can control the maximum number of applications that can run at any one time. If you run applications in the background, there is no restriction imposed by WebSphere MQ on the number of applications that can run.

To run the started application in the background on Windows systems, within the *AppLId* field, prefix the name of your application with a START command. For example:

```
START ?B AMQSECHA
```

To run the started application in the background on UNIX systems, put an & at the end of the *EnvData* of the PROCESS definition.

Note: Where a Windows path has spaces as a part of the path name, these should be enclosed in quotation marks (") to ensure that it is handled as a single argument. For example, " C:\Program Files\Application Directory\Application.exe".

The following is an example of an APPLICID string where the file name includes spaces as a part of the path:

```
START "" /B "C:\Program Files\Application Directory\Application.exe"
```

The syntax of the Windows START command in the example includes an empty string enclosed in double quotation marks. START specifies that the first argument in quotation marks will be treated as the title of the new command. To ensure that Windows does not mistake the application path for a 'title' argument, add a title string enclosed in double quotation marks to the command before the application name.

The following trigger monitors are provided for the WebSphere MQ client:

runmqtmc

This is the same as runmqtrm except that it links with the WebSphere MQ MQI client libraries.

For CICS

The amqltmc0 trigger monitor is provided for CICS. It works in the same way as the standard trigger monitor, runmqtrm, but you run it in a different way and it triggers CICS transactions.

This topic applies only to Windows, UNIX, and Linux systems.

It is supplied as a CICS program; define it with a 4-character transaction name. Enter the 4-character name to start the trigger monitor. It uses the default queue manager (as named in the qm.ini file or, on WebSphere MQ for Windows, the registry), and the SYSTEM.CICS.INITIATION.QUEUE.

If you want to use a different queue manager or queue, build the trigger monitor MQTMC2 structure: this requires you to write a program using the EXEC CICS START call, because the structure is too long to add as a parameter. Then, pass the MQTMC2 structure as data to the START request for the trigger monitor.

When you use the MQTMC2 structure, you need to supply only the *StrucId*, *Version*, *QName*, and *QMgrName* parameters to the trigger monitor as it does not reference any other fields.

Messages are read from the initiation queue and used to start CICS transactions, using EXEC CICS START, assuming the APPL_TYPE in the trigger message is MQAT_CICS. The reading of messages from the initiation queue is performed under CICS syncpoint control.

Messages are generated when the monitor starts and stops, and when an error occurs. These messages are sent to the CSMT transient data queue.

Here are the available versions of the trigger monitor:

Version	Use
amqltmc0	TXSeries for AIX, HP-UX, and Sun Solaris Version 5.1
amqltmc4	TXSeries for Windows, Version 5.1
amqltmcc	Client bound version of the CICS trigger monitor

If you need a trigger monitor for other environments, write a program that can process the trigger messages that the queue manager puts on the initiation queues. Such a program should perform the following actions:

1. Use the MQGET call to wait for a message to arrive on the initiation queue.
2. Examine the fields of the MQTM structure of the trigger message to find the name of the application to start, and the environment in which it runs.
3. Issue an environment-specific start command.
4. Convert the MQTM structure to the MQTMC2 structure if required.
5. Pass either the MQTMC2 or MQTM structure to the started application. This can contain user data.
6. Associate with your application queue the application that is to serve that queue. You do this by naming the process definition object (if created) in the *ProcessName* attribute of the queue.

For more information about the trigger monitor interface, see [MQTMC2](#).

Properties of trigger messages

The following topics describe some other properties of trigger messages.

- [“Persistence and priority of trigger messages” on page 330](#)
- [“Queue manager restart and trigger messages” on page 330](#)
- [“Trigger messages and changes to object attributes” on page 330](#)
- [“Format of trigger messages” on page 330](#)

Persistence and priority of trigger messages

Trigger messages are not persistent because there is no requirement for them to be so.

However, the conditions for generating triggering events do persist, so trigger messages are generated whenever these conditions are met. If a trigger message is lost, the continued existence of the application message on the application queue guarantees that the queue manager generates a trigger message as soon as all the conditions are met.

If a unit of work is rolled back, any trigger messages it generated are always delivered.

Trigger messages take the default priority of the initiation queue.

Queue manager restart and trigger messages

Following the restart of a queue manager, when an initiation queue is next opened for input, a trigger message can be put to this initiation queue if an application queue associated with it has messages on it, and is defined for triggering.

Trigger messages and changes to object attributes

Trigger messages are created according to the values of the trigger attributes in force at the time of the trigger event.

If the trigger message is not made available to a trigger monitor until later (because the message that caused it to be generated was put within a unit of work), any changes to the trigger attributes in the meantime have no effect on the trigger message. In particular, disabling triggering does not prevent a trigger message being made available once it has been created. Also, the application queue might no longer exist at the time that the trigger message is made available.

Format of trigger messages

The format of a trigger message is defined by the MQTM structure.

This has the following fields, which the queue manager fills when it creates the trigger message, using information in the object definitions of the application queue and of the process associated with that queue:

StrucId

The structure identifier.

Version

The version of the structure.

QName

The name of the application queue on which the trigger event occurred. When the queue manager creates a trigger message, it fills this field using the *QName* attribute of the application queue.

ProcessName

The name of the process definition object that is associated with the application queue. When the queue manager creates a trigger message, it fills this field using the *ProcessName* attribute of the application queue.

TriggerData

A free-format field for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *TriggerData* attribute of the application queue. On any WebSphere MQ product except WebSphere MQ for z/OS, this field can be used to specify the name of the channel to be triggered.

ApplType

The type of the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the *ApplType* attribute of the process definition object identified in *ProcessName*.

AppId

A character string that identifies the application that the trigger monitor is to start. When the queue manager creates a trigger message, it fills this field using the *AppId* attribute of the process definition object identified in *ProcessName*. When you use trigger monitor CKTI or CSQQTRMN supplied by WebSphere MQ for z/OS, the *AppId* attribute of the process definition object is a CICS or IMS transaction identifier.

EnvData

A character field containing environment-related data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *EnvData* attribute of the process definition object identified in *ProcessName*. The WebSphere MQ for z/OS-supplied trigger monitors (CKTI or CSQQTRMN) do not use this field, but other trigger monitors might choose to use it.

UserData

A character field containing user data for use by the trigger monitor. When the queue manager creates a trigger message, it fills this field using the *UserData* attribute of the process definition object identified in *ProcessName*. This field can be used to specify the name of the channel to be triggered.

There is a full description of the trigger message structure in [MQTM](#).

When triggering does not work

A program is not triggered if the trigger monitor cannot start the program or the queue manager cannot deliver the trigger message. For example, the applid in the process object must specify that the program is to be started in the background; otherwise, the trigger monitor cannot start the program.

If a trigger message is created but cannot be put on the initiation queue (for example, because the queue is full or the length of the trigger message is greater than the maximum message length specified for the initiation queue), the trigger message is put instead on the dead-letter (undelivered message) queue.

If the put operation to the dead-letter queue cannot complete successfully, the trigger message is discarded and a warning message is sent to the z/OS console or to the system operator or is put on the error log.

Putting the trigger message on the dead-letter queue might generate a trigger message for that queue. This second trigger message is discarded if it adds a message to the dead-letter queue.

If the program is triggered successfully but abends before it receives the message from the queue, use a trace utility (for example, CICS AUXTRACE if the program is running under CICS) to find the cause of the failure.

Working with the MQI and clusters

There are special options on calls and return codes that relate to clustering.

Use the following links to find out more about the options available on the calls and return codes for use with clusters:

- [“MQOPEN and clusters” on page 332](#)
- [“MQPUT, MQPUT1 and clusters” on page 333](#)
- [“MQINQ and clusters” on page 333](#)
- [“MQSET and clusters” on page 334](#)
- [“Return codes” on page 334](#)

Related concepts

[“The Message Queue Interface overview” on page 187](#)

Learn about the Message Queue Interface (MQI) components.

[“Connecting to and disconnecting from a queue manager” on page 198](#)

To use WebSphere MQ programming services, a program must have a connection to a queue manager. Use this information to learn how to connect to and disconnect from a queue manager.

[“Opening and closing objects” on page 206](#)

This information provides an insight into opening and closing WebSphere MQ objects.

[“Putting messages on a queue” on page 216](#)

Use this information to learn how to put messages on a queue.

[“Getting messages from a queue” on page 230](#)

Use this information to learn about getting messages from a queue.

[“Inquiring about and setting object attributes” on page 306](#)

Attributes are the properties that define the characteristics of a WebSphere MQ object.

[“Committing and backing out units of work” on page 309](#)

This information describes how to commit and back out any recoverable get and put operations that have occurred in a unit of work.

[“Starting IBM WebSphere MQ applications using triggers” on page 315](#)

Learn about triggers and how to start IBM WebSphere MQ applications using triggers.

MQOPEN and clusters

The queue to which a message is put to, or read from, when a cluster queue is opened, depends on the MQOPEN call.

Selecting the target queue

If you do not provide a queue manager name in the object descriptor, MQOD, the queue manager selects the queue manager to send the message to. If you do provide a queue manager name in the object descriptor, then messages are always sent to the queue manager you have selected.

If the queue manager is selecting the target queue manager, the selection depends on the binding options, MQ00_BIND_*, and if a local queue exists. If there is a local instance of the queue, it is always opened in preference to a remote instance, unless the CLWLUSEQ attribute is set to ANY. Otherwise, the selection depends on the binding options. Either MQ00_BIND_ON_OPEN or MQ00_BIND_ON_GROUP must be specified when using [message groups](#) with clusters to ensure that all messages in the group are processed at the same destination.

If the queue manager is selecting the target queue manager, it does so in a round-robin fashion, using the workload management algorithm; see [Workload balancing](#).

MQ00_BIND_ON_OPEN

The MQ00_BIND_ON_OPEN option on the MQOPEN call specifies that the target queue manager is to be fixed. Use the MQ00_BIND_ON_OPEN option if there are multiple instances of the same queue within a cluster. All messages put to the queue specifying the object handle returned from the MQOPEN call are directed to the same queue manager.

- Use the MQ00_BIND_ON_OPEN option if messages have affinities. For example, if a batch of messages is all to be processed by the same queue manager, specify MQ00_BIND_ON_OPEN when you open the queue. IBM WebSphere MQ fixes the queue manager and the route to be taken by all messages put to that queue.
- If MQ00_BIND_ON_OPEN option is specified, the queue must be reopened for a new instance of the queue to be selected.

MQ00_BIND_NOT_FIXED

The MQ00_BIND_NOT_FIXED option on the MQOPEN call specifies that the target queue manager is not fixed. Messages written to the queue specifying the object handle returned from the MQOPEN call are routed to a queue manager at MQPUT time on a message-by-message basis. Use the MQ00_BIND_NOT_FIXED option if you do not want to force all your messages to be written to the same destination.

- Do not specify MQ00_BIND_NOT_FIXED and MQMF_SEGMENTATION_ALLOWED at the same time. If you do, the segments of your message might be delivered to different queue managers, scattered throughout the cluster.

MQOO_BIND_ON_GROUP

Allows an application to request that a group of messages is allocated to the same destination instance. This option is valid only for queues, and affects only cluster queues. If specified for a queue that is not a cluster queue, the option is ignored.

- Groups are only routed to a single destination when MQPMO_LOGICAL_ORDER is specified on the MQPUT. When MQOO_BIND_ON_GROUP is specified, but a message is not part of a group, BIND_NOT_FIXED behaviour is used instead.

MQOO_BIND_AS_Q_DEF

If you do not specify either MQOO_BIND_ON_OPEN, MQOO_BIND_NOT_FIXED or MQOO_BIND_ON_GROUP, the default option is MQOO_BIND_AS_Q_DEF. Using MQOO_BIND_AS_Q_DEF causes the binding that is used for the queue handle to be taken from the DefBind queue attribute.

Relevance of MQOPEN options

The MQOPEN options MQOO_BROWSE, MQOO_INPUT_*, or MQOO_SET require a local instance of the cluster queue for MQOPEN to succeed.

The MQOPEN options MQOO_OUTPUT, MQOO_BIND_*, or MQOO_INQUIRE do not require a local instance of the cluster queue to succeed.

Resolved queue manager name

When a queue manager name is resolved at MQOPEN time, the resolved name is returned to the application. If the application tries to use this name on a subsequent MQOPEN call, it might find that it is not authorized to access the name.

MQPUT, MQPUT1 and clusters

If MQOO_BIND_NOT_FIXED is specified on an MQOPEN the workload management routines chooses which destination MQPUT or MQPUT1 select.

If MQOO_BIND_NOT_FIXED is specified on an MQOPEN call, each subsequent MQPUT call invokes the workload management routine to determine which queue manager to send the message to. The destination and route to be taken are selected on a message-by-message basis. The destination and route might change after the message has been put if conditions in the network change. The MQPUT1 call always operates as though MQOO_BIND_NOT_FIXED were in effect, that is, it always invokes the workload management routine.

When the workload management routine has selected a queue manager, the local queue manager completes the put operation. The message can be placed on different queues:

1. If the destination is the local instance of the queue, the message is placed on the local queue.
2. If the destination is a queue manager in a cluster, the message is placed on a cluster transmission queue.
3. If the destination is a queue manager outside a cluster, the message is placed on a transmission queue with the same name as the target queue manager.

If MQOO_BIND_ON_OPEN is specified on the MQOPEN call, MQPUT calls do not invoke the workload management routine because the destination and route have already been selected.

MQINQ and clusters

Which cluster queue is inquired upon depends upon the options you combine with MQOO_INQUIRE.

Before you can inquire on a queue, open it using the MQOPEN call and specify MQOO_INQUIRE.

To inquire on a cluster queue, use the MQOPEN call and combine other options with MQOO_INQUIRE. The attributes that can be inquired depend on whether there is a local instance of the cluster queue, and on how the queue is opened:

- Combining MQ00_BROWSE, MQ00_INPUT_*, or MQ00_SET with MQ00_INQUIRE requires a local instance of the cluster queue for the open to succeed. In this case you can inquire on all the attributes that are valid for local queues.
- Combining MQ00_OUTPUT with MQ00_INQUIRE, and specifying none of the preceding options, the instance opened is either:
 - The instance on the local queue manager, if there is one. In this case you can inquire on all the attributes that are valid for local queues.
 - An instance elsewhere in the cluster, if there is no local queue-manager instance. In this case only the following attributes can be inquired on. The QType attribute has the value MQQT_CLUSTER in this case.
 - DefBind
 - DefPersistence
 - DefPriority
 - InhibitPut
 - QDesc
 - QName
 - QType

To inquire on the DefBind attribute of a cluster queue, use the MQINQ call with the selector MQIA_DEF_BIND. The value returned is either MQBND_BIND_ON_OPEN or MQBND_BIND_NOT_FIXED, or MQBND_BIND_ON_GROUP. Either MQBND_BIND_ON_OPEN or MQBND_BIND_ON_GROUP must be specified when using groups with clusters.

To inquire on the CLUSTER and CLUSNL attributes of the local instance of a queue, use the MQINQ call with the selector MQCA_CLUSTER_NAME or the selector MQCA_CLUSTER_NAMELIST.

Note: If you open a cluster queue without fixing the queue that MQOPEN has bound to, successive MQINQ calls might inquire on different instances of the cluster queue.

Related concepts

[“MQOPEN option for cluster queue” on page 211](#)

The binding used for the queue handle is taken from the *DefBind* queue attribute, which can take the value MQBND_BIND_ON_OPEN, MQBND_BIND_NOT_FIXED, or MQBND_BIND_ON_GROUP.

MQSET and clusters

The MQOPEN option MQ00_SET option requires there to be a local instance of a cluster queue for MQSET to succeed.

You cannot use the MQSET call to set the attributes of a queue elsewhere in the cluster.

You can open a local alias or remote queue defined with the cluster attribute and use the MQSET call. You can set the attributes of the local alias or remote queue. It does not matter if the target queue is a cluster queue defined on a different queue manager.

Return codes

Return codes specific to clusters

MQRC_CLUSTER_EXIT_ERROR (2266 X'8DA')

An MQOPEN, MQPUT, or MQPUT1 call is issued to open a cluster queue or put a message on it. The cluster workload exit, defined by the ClusterWorkloadExit attribute of a queue-manager, fails unexpectedly or does not respond in time.

A message is written to the system log on WebSphere MQ for z/OS giving more information about this error.

Subsequent MQOPEN, MQPUT, and MQPUT1 calls for this queue handle are processed as though the ClusterWorkloadExit attribute were blank.

MQRC_CLUSTER_EXIT_LOAD_ERROR (2267 X'8DB')

On z/OS, the cluster workload exit cannot be loaded.

A message is written to the system log and processing continues as though the ClusterWorkloadExit attribute is blank.

On platforms other than z/OS, an MQCONN or MQCONNEX call is issued to connect to a queue manager. The call fails because the cluster workload exit, defined by the queue-manager ClusterWorkloadExit attribute of the queue manager, cannot be loaded.

MQRC_CLUSTER_PUT_INHIBITED (2268 X'8DC')

An MQOPEN call with the MQOO_OUTPUT and MQOO_BIND_ON_OPEN options in effect is issued for a cluster queue. All the instances of the queue in the cluster are currently put-inhibited by having the InhibitPut attribute set to MQQA_PUT_INHIBITED. Because there are no queue instances available to receive messages, the MQOPEN call fails.

This reason code occurs only when both of the following are true:

- There is no local instance of the queue. If there is a local instance, the MQOPEN call succeeds, even if the local instance is put-inhibited.
- There is no cluster workload exit for the queue, or there is a cluster workload exit but it does not choose a queue instance. (If the cluster workload exit chooses a queue instance, the MQOPEN call succeeds, even if that instance is put-inhibited.)

If the MQOO_BIND_NOT_FIXED option is specified on the MQOPEN call, the call can succeed even if all the queues in the cluster are put-inhibited. However, a subsequent MQPUT call might fail if all the queues are still put-inhibited at the time of that call.

MQRC_CLUSTER_RESOLUTION_ERROR (2189 X'88D')

1. An MQOPEN, MQPUT, or MQPUT1 call is issued to open a cluster queue or put a message on it. The queue definition cannot be resolved correctly because a response is required from the full repository queue manager but none is available.
2. An MQOPEN, MQPUT, MQPUT1 or MQSUB call is issued for a topic object specifying PUBSCOPE(ALL) or SUBSCOPE(ALL). The cluster topic definition cannot be resolved correctly because a response is required from the full repository queue manager but none is available.

MQRC_CLUSTER_RESOURCE_ERROR (2269 X'8DD')

An MQOPEN, MQPUT, or MQPUT1 call is issued for a cluster queue. An error occurs while trying to use a resource required for clustering.

MQRC_NO_DESTINATIONS_AVAILABLE (2270 X'8DE')

An MQPUT or MQPUT1 call is issued to put a message on a cluster queue. At the time of the call, there are no longer any instances of the queue in the cluster. The MQPUT fails and the message is not sent.

The error can occur if MQOO_BIND_NOT_FIXED is specified on the MQOPEN call that opens the queue, or MQPUT1 is used to put the message.

MQRC_STOPPED_BY_CLUSTER_EXIT (2188 X'88C')

An MQOPEN, MQPUT, or MQPUT1 call is issued to open or put a message on a cluster queue. The cluster workload exit rejects the call.

Writing client applications

What you need to know to write client applications on WebSphere MQ.

Applications can be built and run in the WebSphere MQ client environment. The application must be built and linked to the WebSphere MQ MQI client used. The way in which applications are built and linked

varies according to the platform and programming language used. For information on how to build client applications, see [“Building applications for WebSphere MQ MQI clients” on page 341.](#)

You can run a WebSphere MQ application both in a full WebSphere MQ environment and in a WebSphere MQ MQI client environment without changing your code, provided that certain conditions are met. For more information on running your applications in the WebSphere MQ client environment, see [“Running applications in the IBM WebSphere MQ MQI client environment” on page 343.](#)

If you use the message queue interface (MQI) to write applications to run in a WebSphere MQ MQI client environment there are some additional controls to impose during an MQI call to ensure that the WebSphere MQ application processing is not disrupted. For more information about these controls, see [“Using the message queue interface \(MQI\) in a client application” on page 336.](#)

See the following topics for information preparing and running other application types as client applications:

- [“Preparing and running CICS and Tuxedo applications” on page 354](#)
- [“Preparing and running Microsoft Transaction Server applications” on page 39](#)
- [“Preparing and running WebSphere MQ JMS applications” on page 357](#)

Related concepts

[“Application development concepts” on page 8](#)

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Sample WebSphere MQ programs” on page 91](#)

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

[“Writing a queuing application” on page 186](#)

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Using Web services in WebSphere MQ” on page 904](#)

You can develop IBM WebSphere MQ applications for Web services using the IBM WebSphere MQ transport for SOAP or the IBM WebSphere MQ bridge for HTTP.

[“Writing publish/subscribe applications” on page 266](#)

Start writing publish/subscribe WebSphere MQ applications.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

[“Handling program errors” on page 523](#)

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Using the message queue interface (MQI) in a client application

This collection of topics considers the differences between writing your WebSphere MQ application to run in a WebSphere MQ MQI client environment and to run in the full WebSphere MQ queue manager environment.

When you design an application, consider what controls you need to impose during an MQI call to ensure that the WebSphere MQ application processing is not disrupted.

Limiting the size of a message in a client application

A queue manager has a maximum message length, but the maximum size of message you can transmit from a client application is limited by the channel definition.

The maximum message length (MaxMsgLength) attribute of a queue manager is the maximum length of a message that can be handled by that queue manager.

On platforms other than z/OS, you can increase the maximum message length attribute of a queue manager. Details are given in [ALTER QMGR](#).

You can find out the value of MaxMsgLength for a queue manager by using the MQINQ call.

If the MaxMsgLength attribute is changed, no check is made that there are not already queues, and even messages, with a length greater than the new value. After you change this attribute, restart applications and channels in order to ensure that the change has taken effect. It is then not possible for any new messages to be generated that exceed the MaxMsgLength of either the queue manager or the queue (unless queue manager segmentation is allowed).

The maximum message length in a channel definition limits the size of a message that you can transmit along a client connection. If a WebSphere MQ application tries to use the MQPUT call or the MQGET call with a message larger than this, an error code is returned to the application. The maximum message size parameter of the channel definition does not affect the maximum message size which can be consumed using MQCB over a client connection.

Choosing client or server coded character set identifier (CCSID)

Use the local CCSID for the client. The queue manager performs necessary conversion. Use the MQCCSID environment variable to override the CCSID. If your application performs multiple PUTs, the CCSID and encoding fields of the MQMD can be overwritten after completion of the first PUT.

The data passed across the MQI from the application to the client stub must be in the local CCSID, encoded for the WebSphere MQ MQI client. If the connected queue manager requires the data to be converted, then conversion is done by the client support code on the queue manager.

The Java client in V7, however, can do the conversion if the queue manager is unable to do so. See [“WebSphere MQ classes for Java client connections”](#) on page 639

The client code assumes that the character data crossing the MQI in the client is in the CCSID configured for that workstation. If this CCSID is an unsupported CCSID or is not the required CCSID, it can be overridden with the MQCCSID environment variable by using one of these commands:

- On Windows:

```
SET MQCCSID=850
```

- On UNIXsystems:

```
export MQCCSID=850
```

If this parameter is set in the profile, all MQI data is assumed to be in code page 850.

Note: The assumption about code page 850 does not apply to application data in the message.

If your application is performing multiple PUTs that include WebSphere MQ headers after the message descriptor (MQMD), be aware that the CCSID and encoding fields of the MQMD are overwritten after completion of the first PUT.

After the first PUT, these fields contain the value used by the connected queue manager to convert the WebSphere MQ headers. Ensure that your application resets the values to the values it requires.

Using MQINQ in a client application

Some values queried using MQINQ are modified by the client code.

CCSID

is set to the client CCSID, not that of the queue manager.

MaxMsgLength

is reduced if it is restricted by the channel definition. This will be the lower of:

- The value defined in the queue definition, or
- The value defined in the channel definition

For more information, see the [MQINQ](#).

Using sync point coordination in a client application

An application running on the base client can issue MQCMIT and MQBACK, but the scope of the sync point control is limited to the MQI resources. You can use an external transaction manager with an extended transactional client.

Within WebSphere MQ, one of the roles of the queue manager is sync point control within an application. If an application runs on a WebSphere MQ base client, it can issue MQCMIT and MQBACK, but the scope of the sync point control is limited to the MQI resources. The WebSphere MQ verb MQBEGIN is not valid in a base client environment.

Applications running in the full queue manager environment on the server can coordinate multiple resources (for example databases) via a transaction monitor. On the server you can use the Transaction Monitor supplied with WebSphere MQ products, or another transaction monitor such as CICS. You cannot use a transaction monitor with a base client application.

You can use an external transaction manager with a WebSphere MQ extended transactional client. See [What is an extended transactional client?](#) for details.

Using read ahead in a client application

You can use read ahead on a client to allow non persistent messages to be sent to a client without the client application having to request the messages.

When a client requires a message from a server, it sends a request to the server. It sends a separate request for each of the messages it consumes. To improve the performance of a client consuming non persistent messages by avoiding having to send these request messages, a client can be configured to use read ahead. Read ahead allows messages to be sent to a client without an application having to request them.

Using read ahead can improve performance when consuming non persistent messages from a client application. This performance improvement is available to both MQI and JMS applications. Client applications using MQGET or asynchronous consumption benefit from the performance improvements when consuming non persistent messages.

When you call MQOPEN with MQOO_READ_AHEAD, the WebSphere MQ client only enables readahead if certain conditions are met. These conditions include:

- Both the client and remote queue manager must be at WebSphere MQ Version 7 or later.
- The client application must be compiled and linked against the threaded WebSphere MQ MQI client libraries.
- The client channel must be using TCP/IP protocol
- The channel must have a non-zero SharingConversations (SHARECNV) setting in both the client and server channel definitions.

When read ahead is enabled, messages are sent to a memory buffer on the client called the read ahead buffer. The client has a read ahead buffer for each queue it has open with read ahead enabled. The

messages in the read ahead buffer are not persisted. The client periodically updates the server with information about the amount of data it has consumed.

Not all client application designs are suited to using read ahead because not all options are supported for use. Some options are required to be consistent between MQGET calls when read ahead is enabled. If a client alters its selection criteria between MQGET calls, messages being stored in the read ahead buffer remain stranded in the client read ahead buffer. For more information, see [“Improving performance of non-persistent messages” on page 247](#)

Read ahead configuration is controlled by three attributes, MaximumSize, PurgeTime, and UpdatePercentage, which are specified in the MessageBuffer stanza of the WebSphere MQ client configuration file.

Using asynchronous put in a client application

Using asynchronous put, an application can put a message to a queue without waiting for a response from the queue manager. You can use this to improve messaging performance in some situations.

Normally, when an application puts a message or messages on a queue, using MQPUT or MQPUT1, the application has to wait for the queue manager to confirm that it has processed the MQI request. You can improve messaging performance, particularly for applications that use client bindings, and applications that put large numbers of small messages to a queue, by choosing instead to put messages asynchronously. When an application puts a message asynchronously, the queue manager does not return the success or failure of each call, but you can instead check for errors periodically.

To put a message on a queue asynchronously, use the MQPMO_ASYNC_RESPONSE option in the *Options* field of the MQPMO structure.

If a message is not eligible for asynchronous put, it is put to a queue synchronously.

When requesting asynchronous put response for MQPUT or MQPUT1, a CompCode and Reason of MQCC_OK and MQRC_NONE does not necessarily mean that the message was successfully put to a queue. Although the success or failure of each individual MQPUT or MQPUT1 call might not be returned immediately, the first error that occurred under an asynchronous call can be determined later through a call to MQSTAT.

For more details on MQPMO_ASYNC_RESPONSE, see [MQPMO options](#).

The Asynchronous Put sample program demonstrates some of the features available. For details of the features and design of the program, and how to run it, see [“The Asynchronous Put sample program” on page 109](#).

Using sharing conversations in a client application

In an environment where sharing conversations is permitted, conversations can share an MQI channel instance.

Sharing conversations is controlled by two fields, both called SharingConversations, one of which is part of the channel definition (MQCD) structure and one of which is part of the channel exit parameter (MQCXP) structure. The SharingConversations field in the MQCD is an integer value, determining the maximum number of conversations that can share a channel instance associated with the channel. The SharingConversations field in the MQCXP is a boolean value, indicating whether the channel instance is currently shared.

In an environment where sharing conversations is not permitted, new client connections specifying identical MQCDs will not share a channel instance.

A new client application connection will share the channel instance when the following conditions are true:

- Both the client-connection and server-connection ends of the channel instance are configured for sharing conversations, and these values are not overridden by channel exits.
- The client connection MQCD value (supplied on the client MQCONN call or from the client channel definition table (CCDT)) exactly matches the client connection MQCD value supplied on the client

MQCONN call or from the CCDT when the existing channel instance was first established. Note that the original MQCD might have been subsequently altered by exits or by channel negotiation, but that the match is made against the value which was supplied to the client system before these changes were made.

- The sharing conversations limit on the server side is not exceeded.

If a new client application connection matches the criteria to run sharing a channel instance with other conversations, this decision is made before any exits are called on that conversation. Exits on such a conversation cannot alter the fact that it is sharing the channel instance with other conversations. If there are no existing channel instances matching the new channel definition, a new channel instance is connected.

Channel negotiation only occurs for the first conversation on a channel instance; the negotiated values for the channel instance are fixed at that stage and cannot be altered when subsequent conversations start. TLS/SSL authentication also only occurs for the first conversation.

If the MQCD SharingConversations value is altered during the initialization of any security, send or receive exits for the first conversation on the socket at either the client-connection or the server-connection end of the channel instance, the new value it has after all these exits are initialized is used to determine the sharing conversations value for the channel instance (the lowest value takes precedence).

If the negotiated value for sharing conversations is zero, the channel instance is never shared. Further exit programs that set this field to zero similarly run on their own channel instance.

If the negotiated value for sharing conversations is greater than zero then MQCXP SharingConversations is set to TRUE for subsequent calls to exits, indicating that other exit programs on this channel instance can be entered simultaneously with this one.

When you write a channel exit program, consider whether it will run on a channel instance that might involve sharing conversations. If the channel instance might involve sharing conversations, consider the effect on other instances of the channel exit of changing MQCD fields; all MQCD fields have common values across all the sharing conversations. After the channel instance is established, if exit programs try to alter MQCD fields they might encounter problems because other instances of exit programs running on the channel instance could be attempting to alter the same fields at the same time. If this situation could arise with your exit programs, you must serialize access to the MQCD in your exit code.

If you are working with a channel which is defined to share conversations, but you do not want sharing to occur on a particular channel instance, set the MQCD value of SharingConversations to 1 or 0 when you initialize a channel exit on the first conversation on the channel instance. See [SharingConversations](#) for an explanation of the values of SharingConversations.

Example

Sharing conversations is enabled.

You are using a client-connection channel definition which specifies an exit program.

The first time that this channel starts, the exit program alters some of the MQCD parameters when it is initialized. These are acted on by the channel, so the definition that the channel is running with is now different from the one that was originally supplied. The MQCXP SharingConversations parameter is set to TRUE.

The next time that the application connects using this channel, the conversation runs on the channel instance which was started previously, because it has the same original channel definition. The channel instance the application connects to the second time is the same instance as the first time it connected. Consequently, it uses the definitions that have been altered by the exit program. When the exit program is initialized for the second conversation, although it can alter MQCD fields, they are *not* acted on by the channel. These same characteristics apply to any subsequent conversations which share the channel instance.

Using MQCONNX

You can use the MQCONNX call to specify a channel definition (MQCD) structure in the MQCNO structure.

This allows the calling client application to specify the definition of the client-connection channel at run time. For more information, see [Using the MQCNO structure on an MQCONNX call](#). When you use MQCONNX, the call issued at the server depends on the server level and listener configuration.

When you use MQCONNX from a client, the following options are ignored:

- MQCNO_STANDARD_BINDING
- MQCNO_FASTPATH_BINDING

The MQCD structure you can use depends on the MQCD version number you are using. For information on MQCD versions (MQCD_VERSION), see [MQCD Version](#). You can use the MQCD structure, for example, to pass channel-exit programs to the server. If you are using MQCD Version 3 or later, you can use the structure to pass an array of exits to the server. You can use this function to perform more than one operation on the same message, such as encryption and compression, by adding an exit for each operation, rather than modifying an existing exit. If you do not specify an array in the MQCD structure, the single exit fields will be checked. For more information on channel-exit programs, see [“Channel-exit programs for messaging channels”](#) on page 378.

Shared connection handles on MQCONNX

You can share handles between different threads within the same process, using shared connection handles.

When you specify a shared connection handle, the connection handle returned from the MQCONNX call can be passed in subsequent MQI calls on any thread in the process.

Note: You can use a shared connection handle on a WebSphere MQ MQI client to connect to a server queue manager that does not support shared connection handles.

For more information, see [“Using MQCONNX”](#) on page 341.

Building applications for WebSphere MQ MQI clients

Applications can be built and run in WebSphere MQ MQI client environment. The application must be built and linked to the WebSphere MQ MQI client used. The way in which applications are built and linked varies according to the platform and programming language used.

If an application is to run in a client environment, you can write it in the languages shown in the following table:

<i>Table 47. Programming languages supported in client environments</i>						
Client platform	C	C++	COBOL	pTAL	RPG	Visual Basic
AIX	Yes	Yes	Yes			
HP Integrity NonStop Server	Yes		Yes	Yes		
HP-UX	Yes	Yes	Yes			
Linux	Yes	Yes	Yes			
Solaris	Yes	Yes	Yes			
Windows	Yes	Yes	Yes			Yes

See the related topics for instructions for linking or building client applications in these languages.

Linking C applications with the WebSphere MQ MQI client code

Having written your WebSphere MQ application that you want to run on the WebSphere MQ MQI client, you must link it to a queue manager.

You can link your application to a queue manager in two ways:

1. Directly, in which case the queue manager must be on the same workstation as your application
2. To a client library file, which gives you access to queue managers on the same or on a different workstation

WebSphere MQ provides a client library file for each environment:

AIX

libmqic.a library for non-threaded applications, or libmqic_r.a library for threaded applications.

HP-UX

libmqic.sl library for non-threaded applications, or libmqic_r.sl library for threaded applications.

Linux

libmqic.so library for non-threaded applications, or libmqic_r.so library for threaded applications.

Solaris

libmqic.so.

If you want to use the programs on a workstation that has only the WebSphere MQ MQI client for Solaris installed, you must recompile the programs to link them with the client library:

```
$ /opt/SUNWsprio/bin/cc -o <prog> <prog> c -mt -lmqic \  
-lsocket -lc -lnsl -ldl
```

The parameters must be entered in the correct order, as shown.

Windows

MQIC32.LIB.

Linking C++ applications with the WebSphere MQ MQI client code

You can write applications to run on the client in C++. Build methods vary according to the environment.

For information about how to link your C++ applications, see [Building WebSphere MQ C++ programs](#).

For full details of all aspects of using C++, see [Using C++](#)

Linking COBOL applications with the IBM WebSphere MQ MQI client code

Having written a COBOL application that you want to run on the IBM WebSphere MQ MQI client, you must link it with an appropriate library.

IBM WebSphere MQ provides a client library file for each environment:

AIX

Link your non-threaded COBOL application with the library libmqicb.a or threaded COBOL application with libmqicb_r.a.

HP-UX

Link your non-threaded COBOL application with the library libmqicb.sl or threaded COBOL application with libmqicb_r.sl.

Linux

Link your non-threaded COBOL application with the library libmqicb.so or threaded COBOL application with libmqicb_r.so.

Solaris

Link your non-threaded COBOL application with the library libmqicb.so or threaded COBOL application with libmqicb_r.so.

Windows

Link your application code with the MQICCB library for 32-bit COBOL. The IBM WebSphere MQ MQI client for Windows does not support 16-bit COBOL.

Linking Visual Basic applications with the WebSphere MQ MQI client code

You can link Visual Basic applications with the WebSphere MQ MQI client code on Windows.

Link your Visual Basic application with the following include files:

CMQB.bas

MQI

CMQBB.bas

MQAI

CMQCFB.bas

PCF commands

CMQXB.bas

Channels

Set `mqtype=2` for the client in the Visual Basic compiler, to ensure the correct automatic selection of the client dll:

MQIC32.dll

Windows 2000, Windows XP and Windows 2003

Running applications in the IBM WebSphere MQ MQI client environment

You can run an IBM WebSphere MQ application both in a full IBM WebSphere MQ environment and in an IBM WebSphere MQ MQI client environment without changing your code, provided that certain conditions are met.

These conditions are that:

- The application does not need to connect to more than one queue manager concurrently.
- The queue manager name is not prefixed with an asterisk (*) on an MQCONN or MQCONNX call.
- The application does not need to use any of the exceptions listed in [What applications run on a IBM WebSphere MQ MQI client?](#)

Note: The libraries that you use at link-edit time determine the environment in which your application must run.

When working in the IBM WebSphere MQ MQI client environment, remember that:

- Each application running in the IBM WebSphere MQ MQI client environment has its own connections to servers. An application establishes one connection to a server each time it issues an MQCONN or MQCONNX call.
- An application sends and gets messages synchronously. This implies a wait between the time the call is issued at the client and the return of a completion code and reason code across the network.
- All data conversion is done by the server, but see also [MQCCSID](#) for information about overriding the machine's configured CCSID.

Connecting IBM WebSphere MQ MQI client applications to queue managers

An application running in an IBM WebSphere MQ MQI client environment can connect to a queue manager in various ways. You can use environmental variables, the MQCNO structure, or a client definition table.

When an application running in an IBM WebSphere MQ client environment issues an MQCONN or MQCONNX call, the client identifies how it is to make the connection. When an MQCONNX call is issued by an application on an IBM WebSphere MQ client, the MQI client library searches for the client channel information in the following order:

1. Using the contents of the *ClientConnOffset* or *ClientConnPtr* fields of the MQCNO structure (if supplied). These fields identify the channel definition structure (MQCD) to be used as the definition of the client connection channel. Connection details can be overridden by using a pre-connect exit. For more information, see [“Referencing connection definitions using a pre-connect exit from a repository” on page 404.](#)
2. If the MQSERVER environment variable is set, the channel it defines is used.
3. If an mqclient.ini file is defined and contains a ServerConnectionParms, the channel that it defines is used. For more information, see [Configuring a client using a configuration file](#) and [CHANNELS stanza of the client configuration file.](#)
4. If the MQCHLLIB and MQCHLTAB environment variables are set, the client channel definition table they point to is used.
5. If an mqclient.ini file is defined and contains ChannelDefinitionDirectory and ChannelDefinitionFile attributes, these attributes are used to locate the client channel definition table. For more information, see [Configuring a client using a configuration file](#) and [CHANNELS stanza of the client configuration file.](#)
6. Finally, if the environment variables are not set, the client searches for a client channel definition table with a path and name that are established from the DefaultPrefix in the mqs.ini file. If the search for a client definition table fails, the client uses the following paths:
 - UNIX and Linux systems: /var/mqm/AMQCLCHL.TAB
 - Windows: C:\Program Files\IBM\WebSphere MQ\amqclchl.tab

The first of the options described in the previous list (using the *ClientConnOffset* or *ClientConnPtr* fields of MQCNO) is supported only by the MQCONN call. If the application is using MQCONN rather than MQCONNEX, the channel information is searched for in the remaining five ways in the order shown in the list. If the client fails to find the channel information, the MQCONN or MQCONNEX call fails.

The channel name (for the client connection) must match the server-connection channel name defined on the server for the MQCONN or MQCONNEX call to succeed.

If you receive an MQRC_Q_MGR_NOT_AVAILABLE return code from your application with an error message in the error log file of AMQ9517 - File damaged, see [Migration and client channel definition tables \(CCDT\).](#)

Related concepts

[Client channel definition table](#)

Related tasks

[Configuring connections between the server and client](#)

Related reference

[MQSERVER](#)

[MQCHLLIB](#)

[MQCHLTAB](#)

Connecting client applications to queue managers using environment variables

Client channel information can be supplied to an application running in a client environment by the MQSERVER, MQCHLLIB, and MQCHLTAB environment variables.

See [MQSERVER](#), [MQCHLLIB](#) and [MQCHLTAB](#) for details of these variables.

Connecting client applications to queue managers using the MQCNO structure

You can specify the definition of the channel in a channel definition structure (MQCD), which is supplied using the MQCNO structure of the MQCONN call.

For more information see [Using the MQCNO structure on an MQCONN call.](#)

Connecting client applications to queue managers using a client channel definition table

If you use the MQSC DEFINE CHANNEL command, the details you provide are placed in the client channel definition table (ccdt). The contents of the *QMgrName* parameter of the MQCONN or MQCONNX call determines which queue manager the client connects to.

This file is accessed by the client to determine the channel an application will use. Where there is more than one suitable channel definition, the choice of channel is influenced by the client channel weight (CLNTWGHT) and connection affinity (AFFINITY) channel attributes.

Role of the client channel definition table

The client channel definition table (CCDT) contains definitions of client connection channels. It is particularly useful if your client applications might need to connect to a number of alternative queue managers.

The client channel definition table is created when you define a queue manager.

Note: The same file can be used by more than one IBM WebSphere MQ client. You access different versions of this file using the MQCHLLIB and MQCHLTAB IBM WebSphere MQ environment variables. See [Using WebSphere MQ environment variables](#) for information about environment variables.

Queue manager groups in the CCDT

You can define a set of connections in the client channel definition table (CCDT) as a *queue manager group*. You can connect an application to a queue manager that is part of a queue manager group. This can be done by prefixing the queue manager name on an MQCONN or MQCONNX call with an asterisk.

You might choose to define connections to more than one server machine because:

- You want to connect a client to any one of a set of queue managers that is running, to improve availability.
- You want to reconnect a client to the same queue manager it connected to successfully last time, but connect to a different queue manager if the connection fails.
- You want to be able to retry a client connection to a different queue manager if the connection fails, by issuing the MQCONN in the client program again.
- You want to automatically reconnect a client connection to another queue manager if the connection fails, without writing any client code.
- You want to automatically reconnect a client connection to a different instance of a multi-instance queue manager if a standby instance takes over, without writing any client code.
- You want to balance your client connections across a number of queue managers, with more clients connecting to some queue managers than others.
- You want to spread the reconnection of many client connections over multiple queue managers and over time, in case the high volume of connections causes a failure.
- You want to be able to move your queue managers without changing any client application code.
- You want to write client application programs that do not need to know queue manager names.

It is not always appropriate to connect to different queue managers. An extended transactional client or a Java client in WebSphere Application Server, for example, might need to connect to a predictable queue manager instance. Automatic client reconnect is not supported by WebSphere MQ classes for Java.

A queue manager group is a set of connections defined in the client channel definition table (CCDT). The set is defined by its members having the same value of the **QMNAME** attribute in their channel definitions.

Figure 70 on page 346 is a graphical representation of a client connection table, showing three queue manager groups, two named queue manager groups written in the CCDT as **QMNAME**(QM1) and **QMNAME**(QMGrp1), and one blank or default group written as **QMNAME**(' ').

1. Queue manager group QM1 has three client connection channels, connecting it to queue managers QM1 and QM2. QM1 might be a multi-instance queue manager located on two different servers.

2. The default queue manager group has six client connection channels connecting it to all the queue managers.
3. QMGrp1 has client connection channels to two queue managers, QM4 and QM5.

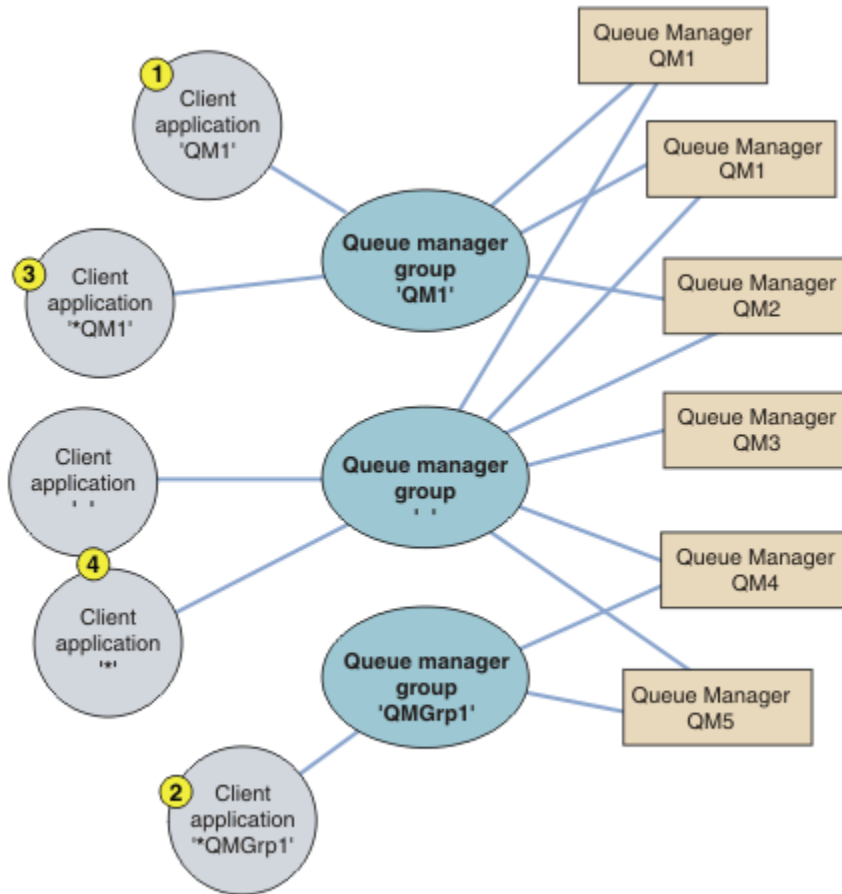


Figure 70. Queue manager groups

Four examples of using this client connection table are described with the help of the numbered client applications in Figure 70 on page 346.

1. In the first example, the client application passes a queue manager name, QM1, as the **QmgrName** parameter to its MQCONN or MQCONNEX MQI call. The WebSphere MQ client code selects the matching queue manager group, QM1. The group contains three connection channels, and the WebSphere MQ MQI client tries to connect to QM1 using each of these channels in turn until it finds an WebSphere MQ listener for the connection attached to a running queue manager called QM1.

The order of connection attempts depends on the value of the client connection AFFINITY attribute and the client channel weightings. Within these constraints, the order of connection attempts is randomized, both over the three possible connections, and over time, in order to spread out the load of making connections.

The MQCONN or MQCONNEX call issued by the client application succeeds when a connection is established to a running instance of QM1.

2. In the second example, the client application passes a queue manager name prefixed with an asterisk, *QMGrp1 as the **QmgrName** parameter to its MQCONN or MQCONNEX MQI call. The WebSphere MQ client selects the matching queue manager group, QMGrp1. This group contains two client connection channels, and the WebSphere MQ MQI client tries to connect to *any* queue manager using each channel in turn. In this example, the WebSphere MQ MQI client needs to make a successful connection; the name of the queue manager that it connects to does not matter.

The rule for the order of making connection attempts is the same as before. The only difference is that by prefixing the queue manager name with an asterisk, the client indicates that the name of the queue manager is not relevant.

The MQCONN or MQCONNX call issued by the client application succeeds when a connection is established to a running instance of any queue manager connected to by the channels in the QMgrp1 queue manager group.

3. The third example is essentially the same as the second because the **QmgrName** parameter is prefixed by an asterisk, *QM1. The example illustrates that you cannot determine which queue manager a client channel connection is going to connect to by inspecting the QMNAME attribute in one channel definition by itself. The fact that the **QMNAME** attribute of the channel definition is QM1, is not sufficient to require a connection is made to a queue manager called QM1. If your client application prefixes its **QmgrName** parameter with an asterisk then any queue manager is a possible connection target.

In this case the MQCONN or MQCONNX calls issued by the client application succeed when a connection is established to a running instance of either QM1 or QM2.

4. The fourth example illustrates use of the default group. In this case the client application passes an asterisk, '*', or blank ' ', as the **QmgrName** parameter to its MQCONN or MQCONNX MQI call. By convention in the client channel definition, a blank **QMNAME** attribute signifies the default queue manager group and either a blank or asterisk **QmgrName** parameter matches a blank **QMNAME** attribute.

In this example the default queue manager group has client channel connections to all the queue managers. By selecting the default queue manager group the application might be connected to any queue manager in the group.

The MQCONN or MQCONNX call issued by the client application succeeds when a connection is established to a running instance of any queue manager.

Note: The default group is different from a default queue manager, although an application uses a blank **QmgrName** parameter to connect to either the default queue manager group or to the default queue manager. The concept of a default queue manager group is only relevant to a client application, and a default queue manager to a server application.

Define your client connection channels on one queue manager only, including those channels that connect to a second or third queue manager. Do *not* define them on two queue managers and then try to merge the two client channel definition tables. Only one client channel definition table can be accessed by the client.

Examples

Look again at the [list](#) of reasons for using queue manager groups at the beginning of the topic. How does using a queue manager group provide those capabilities?

Connect to any one of a set of queue managers.

Define a queue manager group with connections to all the queue managers in the set, and connect to the group using the **QmgrName** parameter prefixed by an asterisk.

Reconnect to the same queue manager, but connect to a different one, if the queue manager connected to last time is unavailable.

Define a queue manager group as before but set the attribute, **AFFINITY** (PREFERRED) on each client channel definition.

Retry a connection to another queue manager if a connection fails.

Connect to a queue manager group, and reissue the MQCONN or MQCONNX MQI call if the connection is broken or the queue manager fails.

Automatically reconnect to another queue manager if a connection fails.

Connect to a queue manager group using the MQCONNX **MQCNO** option MQCNO_RECONNECT.

Automatically reconnect to a different instance of a multi-instance queue manager.

Do the same as the preceding example. In this case, if you want to restrict the queue manager group to connect to the instances of a particular multi-instance queue manager, define the group with connections to only the multi-instance queue manager instances.

You can also ask the client application to issue its MQCONN or MQCONNX MQI call with no asterisk prefixed to the **QmgrName** parameter. That way the client application can only connect to the named queue manager. Finally, you can set the **MQCNO** option to MQCNO_RECONNECT_Q_MGR. This option accepts reconnections to the same queue manager that was previously connected. You can also use this value to restrict reconnections to the same instance of a normal queue manager.

Balance client connections across queue managers, with more clients connected to some queue managers than others.

Define a queue manager group, and set the **CLNTWGHT** attribute on each client channel definition to distribute the connections unevenly.

Spread the client reconnection load unevenly, and spread it over time, after a connection or queue manager failure.

Do the same as the preceding example. The WebSphere MQ MQI client randomizes reconnections across queue managers and spreads the reconnections over time.

Move your queue managers without changing any client code.

The CCDT isolates your client application from the location of the queue manager.

You have a choice of either distributing the client connection table to each client or placing the CCDT on a shared file system for each client to refer to. Alternatively use the programmatic version of the CCDT supported in the MQCONNX MQI call and call a service to pass the CCDT to the client application.

Write a client application that does not know queue manager names.

Use queue manager group names and establish a naming convention for queue manager group names that is relevant to your client applications in your organization, and reflects the architecture of your solutions rather than the naming of queue managers.

Connecting to queue-sharing groups

You can connect your application to a queue manager that is part of a queue-sharing group. This can be done by using the queue-sharing group name instead of the queue manager name on the MQCONN or MQCONNX call.

Queue-sharing groups have a name of up to four characters. The name must be unique in your network, and must be different from any queue manager names.

The client channel definition should use the queue sharing group generic interface to connect to an available queue manager in the group. For more information, see [Connecting a client to a queue-sharing group](#). A check is made to ensure that the queue manager the listener connects to is a member of the queue sharing group.

Examples of channel weighting and affinity

These examples illustrate how client-connection channels are selected when non-zero ClientChannelWeights are used.

The ClientChannelWeight and ConnectionAffinity channel attributes control how client-connection channels are selected when more than one suitable channel is available for a connection. These channels are configured to connect to different queue managers in order to provide higher availability, workload balancing, or both. MQCONN calls that could result in a connection to one of several queue managers must prefix the queue manager name with an asterisk as described in: [Examples of MQCONN calls: Example 1. Queue manager name includes an asterisk \(*\)](#).

Applicable candidate channels for a connection are those where the QMNAME attribute matches the queue manager name specified in the MQCONN call. If all applicable channels for a connection have a ClientChannelWeight of zero (the default) then they are selected in alphabetical order as in the example: [Examples of MQCONN calls: Example 1. Queue manager name includes an asterisk \(*\)](#).

The following examples illustrate what happens when non-zero ClientChannelWeights are used. Note that, since this feature involves pseudo-random channel selection, the examples show a sequence of actions that might happen rather than what definitely will.

Example 1. Selecting channels when ConnectionAffinity is set to PREFERRED

This example illustrates how a WebSphere MQ MQI client selects a channel from a CCDT, where the ConnectionAffinity is set to PREFERRED.

In this example, a number of client machines use a Client Channel Definition Table (CCDT) provided by a queue manager. The CCDT includes client connection channels with the following attributes (shown using the syntax of the DEFINE CHANNEL command):

```
CHANNEL(A) QMNAME(DEV) CONNAME(devqm.it.company.example)
CHANNEL(B) QMNAME(CORE) CONNAME(core1.ops.company.example) CLNTWGHT(5) +
AFFINITY(PREFERRED)
CHANNEL(C) QMNAME(CORE) CONNAME(core2.ops.company.example) CLNTWGHT(3) +
AFFINITY(PREFERRED)
CHANNEL(D) QMNAME(CORE) CONNAME(core3.ops.company.example) CLNTWGHT(2) +
AFFINITY(PREFERRED)
```

The application issues MQCONN(*CORE)

Channel A is not a candidate for this connection, because the QMNAME attribute does not match. Channels B, C and D are identified as candidates, and are placed in an order of preference based on their weighting. In this example the order might be C, B, D. The client attempts to connect to the queue manager at core2.ops.company.example. The name of the queue manager at that address is not checked, because the MQCONN call included an asterisk in the queue manager name.

It is important to note that, with AFFINITY(PREFERRED), each time this particular client machine connects it will place the channels in the same initial order of preference. This applies even when the connections are from different processes or at different times.

In this example, the queue manager at core2.ops.company.example cannot be reached. The client attempts to connect to core1.ops.company.example because channel B is next in the order of preference. In addition, channel C is demoted to become the least preferred.

A second MQCONN(*CORE) call is issued by the same application. Channel C was demoted by the previous connection, so the most preferred channel is now B. This connection is made to core1.ops.company.example.

A second machine sharing the same Client Channel Definition Table might place the channels in a different initial order of preference. For example, D, B, C. Under normal circumstances, with all channels working, applications on this machine are connected to core3.ops.company.example while those on the first machine are connected to core2.ops.company.example. This allows workload balancing of large numbers of clients across multiple queue managers while allowing each individual client to connect to the same queue manager if it is available.

Example 2. Selecting channels when ConnectionAffinity is set to NONE

This example illustrates how a WebSphere MQ MQI client selects a channel from a CCDT, where the ConnectionAffinity is set to NONE.

In this example, a number of clients use a Client Channel Definition Table (CCDT) provided by a queue manager. The CCDT includes client connection channels with the following attributes (shown using the syntax of the DEFINE CHANNEL command):

```
CHANNEL(A) QMNAME(DEV) CONNAME(devqm.it.company.example)
CHANNEL(B) QMNAME(CORE) CONNAME(core1.ops.company.example) CLNTWGHT(5) +
AFFINITY(NONE)
CHANNEL(C) QMNAME(CORE) CONNAME(core2.ops.company.example) CLNTWGHT(3) +
AFFINITY(NONE)
CHANNEL(D) QMNAME(CORE) CONNAME(core3.ops.company.example) CLNTWGHT(2) +
AFFINITY(NONE)
```

The application issues MQCONN(*CORE). As in the previous example, channel A is not considered because the QMNAME does not match. Channel B, C, or D are selected based on their weighting, with probabilities of 50%, 30%, or 20%. In this example, channel B might be selected. There is no persistent order of preference created.

A second MQCONN(*CORE) call is made. Again, one of the three applicable channels is selected, with the same probabilities. In this example, channel C is chosen. However, core2.ops.company.example does not

respond, so another choice is made between the remaining candidate channels. Channel B is selected and the application is connected to core1.ops.company.example.

With AFFINITY(NONE), each MQCONN call is independent of any other. Therefore when this example application makes a third MQCONN(*CORE), it might once more attempt to connect through the broken channel C, before choosing one of B or D.

Examples of MQCONN calls

Examples of using MQCONN to connect to a specific queue manager, or to one of a group of queue managers.

In each of the following examples, the network is the same; there is a connection defined to two servers from the same WebSphere MQ MQI client. (In these examples, the MQCONN call could be used instead of the MQCONN call.)

There are two queue managers running on the server machines, one named SALE and the other named SALE_BACKUP.

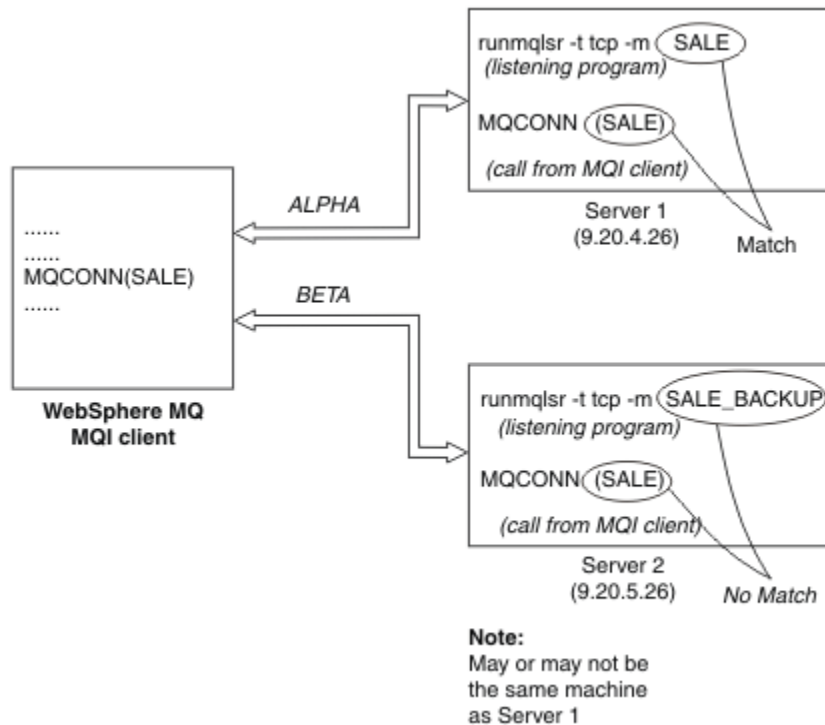


Figure 71. MQCONN example

The definitions for the channels in these examples are:

SALE definitions:

```
DEFINE CHANNEL(ALPHA) CHLTYPE(SVRCONN) TRPTYPE(TCP) +
DESCR('Server connection to WebSphere MQ MQI client')

DEFINE CHANNEL(ALPHA) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(9.20.4.26) DESCR('WebSphere MQ MQI client connection to server 1') +
QMNAME(SALE)

DEFINE CHANNEL(BETA) CHLTYPE(CLNTCONN) TRPTYPE(TCP) +
CONNNAME(9.20.5.26) DESCR('WebSphere MQ MQI client connection to server 2') +
QMNAME(SALE)
```

SALE_BACKUP definition:

```
DEFINE CHANNEL(BETA) CHLTYPE(SVRCONN) TRPTYPE(TCP) +  
DESCR('Server connection to WebSphere MQ MQI client')
```

The client channel definitions can be summarized as follows:

Name	CHLTYPE	TRPTYPE	CONNAME	QMNAME
ALPHA	CLNTCONN	TCP	9.20.4.26	SALE
BETA	CLNTCONN	TCP	9.20.5.26	SALE

What the MQCONN examples demonstrate

The examples demonstrate the use of multiple queue managers as a backup system.

Suppose the communication link to Server 1 is temporarily broken. The use of multiple queue managers as a backup system is demonstrated.

Each example covers a different MQCONN call and gives an explanation of what happens in the specific example presented, by applying the following rules:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name order for a queue manager name (QMNAME field) corresponding to the one given in the MQCONN call.
2. If a match is found, the channel definition is used.
3. An attempt is made to start the channel to the machine identified by the connection name (CONNAME). If this is successful, the application continues. It requires:
 - A listener to be running on the server.
 - The listener to be connected to the same queue manager as the one the client wants to connect to (if specified).
4. If the attempt to start the channel fails and there is more than one entry in the client channel definition table (in this example there are two entries), the file is searched for a further match. If a match is found, processing continues at step 1.
5. If no match is found, or there are no more entries in the client channel definition table and the channel has failed to start, the application is unable to connect. An appropriate reason code and completion code are returned in the MQCONN call. The application can take action based on the reason and completion codes returned.

Example 1. Queue manager name includes an asterisk ()*

In this example the application is not concerned about which queue manager it connects to. The application issues an MQCONN call for a queue manager name including an asterisk. A suitable channel is chosen.

The application issues:

```
MQCONN (*SALE)
```

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned for the queue manager name SALE, matching with the application MQCONN call.
2. Channel definitions for ALPHA and BETA are found.
3. If one channel has a CLNTWGHT value of 0, this channel is selected. If both have a CLNTWGHT value of 0, channel ALPHA is selected because it is first in alphabetical sequence. If both channels have a non-zero CLNTWGHT value, one channel is randomly selected, based on its weighting.
4. An attempt to start the channel is made.
5. If channel BETA was selected, the attempt to start it is successful.

6. If channel ALPHA was selected, the attempt to start it is NOT successful because the communication link is broken. The following steps then apply:
 - a. The only other channel for the queue manager name SALE is BETA.
 - b. An attempt to start this channel is made - this is successful.
7. A check to see that a listener is running shows that there is one running. It is not connected to the SALE queue manager, but because the MQI call parameter has an asterisk (*) included in it, no check is made. The application is connected to the SALE_BACKUP queue manager and continues processing.

Example 2. Queue manager name specified

In this example the application must connect to a particular queue manager. The application issues an MQCONN call for that queue manager name. A suitable channel is chosen.

The application requires a connection to a specific queue manager, named SALE, as seen in the MQI call:

```
MQCONN (SALE)
```

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name sequence, for the queue manager name SALE, matching with the application MQCONN call.
2. The first channel definition found to match is ALPHA.
3. An attempt to start the channel is made - this is *not* successful because the communication link is broken.
4. The client channel definition table is again scanned for the queue manager name SALE and the channel name BETA is found.
5. An attempt to start the channel is made - this is successful.
6. A check to see that a listener is running shows that there is one running, but it is not connected to the SALE queue manager.
7. There are no further entries in the client channel definition table. The application cannot continue and receives return code MQRC_Q_MGR_NOT_AVAILABLE.

Example 3. Queue manager name is blank or an asterisk ()*

In this example the application is not concerned about which queue manager it connects to. The application issues an MQCONN specifying a blank queue manager name or an asterisk. A suitable channel is chosen.

This is treated in the same way as [“Example 1. Queue manager name includes an asterisk \(*\)” on page 351.](#)

Note: If this application were running in an environment other than a WebSphere MQ MQI client, and the name was blank, it would be attempting to connect to the default queue manager. This is *not* the case when it is run from a client environment; the queue manager accessed is the one associated with the listener to which the channel connects.

The application issues:

```
MQCONN (" ")
```

or

```
MQCONN (*)
```

Following the rules, this is what happens in this instance:

1. The client channel definition table (CCDT) is scanned in alphabetical channel name sequence, for a queue manager name that is blank, matching with the application MQCONN call.

2. The entry for the channel name ALPHA has a queue manager name in the definition of SALE. This does *not* match the MQCONN call parameter, which requires the queue manager name to be blank.
3. The next entry is for the channel name BETA.
4. The queue manager name in the definition is SALE. Once again, this does *not* match the MQCONN call parameter, which requires the queue manager name to be blank.
5. There are no further entries in the client channel definition table. The application cannot continue and receives return code MQRC_Q_MGR_NOT_AVAILABLE.

Triggering in the client environment

Messages sent by WebSphere MQ applications running on WebSphere MQ MQI clients contribute to triggering in exactly the same way as any other messages, and they can be used to trigger programs on both the server and the client.

Triggering is explained in detail in the [Triggering channels](#).

The trigger monitor and the application to be started must be on the same system.

The default characteristics of the triggered queue are the same as those in the server environment. In particular, if no MQPMO sync point control options are specified in a client application putting messages to a triggered queue that is local to a z/OS queue manager, the messages are put within a unit of work. If the triggering condition is then met, the trigger message is put on the initiation queue within the same unit of work and cannot be retrieved by the trigger monitor until the unit of work ends. The process that is to be triggered is not started until the unit of work ends.

Process definition

You must define the process definition on the server, because this is associated with the queue that has triggering set on.

The process object defines what is to be triggered. If the client and server are not running on the same platform, any processes started by the trigger monitor must define *AppType*, otherwise the server takes its default definitions (that is, the type of application that is normally associated with the server machine) and causes a failure.

For example, if the trigger monitor is running on a Windows client and wants to send a request to a server on another operating system, MQAT_WINDOWS_NT must be defined otherwise the other operating system uses its default definitions and the process fails.

Trigger monitor

The trigger monitor provided by non-z/OS WebSphere MQ products runs in the client environments for UNIX, Linux and Windows systems.

To run the trigger monitor, issue one of these commands:

-  On Windows, UNIX and Linux platforms:

```
runmqtrmc [-m QMgrName] [-q InitQ]
```

The default initiation queue is SYSTEM.DEFAULT.INITIATION.QUEUE on the default queue manager. The initiation queue is where the trigger monitor looks for trigger messages. It then calls programs for the appropriate trigger messages. This trigger monitor supports the default application type and is the same as runmqtrmc except that it links the client libraries.

The command string, built by the trigger monitor, is as follows:

1. The *ApplicId* from the relevant process definition. *ApplicId* is the name of the program to run, as it would be entered on the command line.
2. The MQTMC2 structure, enclosed in quotation marks, obtained from the initiation queue. A command string is started that has this string, exactly as provided, in quotation marks in order that the system command accepts it as one parameter.
3. The *EnvrData* from the relevant process definition.

The trigger monitor does not look to see if there is another message on the initiation queue until the completion of the application it has started. If the application has much processing to do, the trigger monitor might not keep up with the number of trigger messages arriving. There are two ways to deal with this situation:

1. Have more trigger monitors running

If you choose to have more trigger monitors running, you can control the maximum number of applications that can run at any one time.

2. Run the started applications in the background

If you choose to run applications in the background, WebSphere MQ imposes no restriction on the number of applications that can run.

To run the started application in the background on UNIX and Linux systems, you must put an & (ampersand) at the end of the *EnvrData* of the process definition.

CICS applications (non-z/OS)

A non-z/OS CICS application program that issues an MQCONN or MQCONNX call must be defined to CEDA as RESIDENT. If you relink a CICS server application as a client, you risk losing sync point support.

A non-z/OS CICS application program that issues an MQCONN or MQCONNX call must be defined to CEDA as RESIDENT. To make the resident code as small as possible, you can link to a separate program to issue the MQCONN or MQCONNX call.

If the MQSERVER environment variable is used to define the client connection, it must be specified in the CICSENV.CMD file.

WebSphere MQ applications can be run in a WebSphere MQ server environment or on a WebSphere MQ client without changing code. However, in a WebSphere MQ server environment, CICS can act as sync point coordinator, and you use EXEC CICS SYNCPOINT and EXEC CICS SYNCPOINT ROLLBACK rather than MQCMIT and MQBACK. If a CICS application is simply relinked as a client, sync point support is lost. MQCMIT and MQBACK must be used for the application running on a WebSphere MQ MQI client.

Preparing and running CICS and Tuxedo applications

To run CICS and Tuxedo applications as client applications, you use different libraries from those you use with server applications. The user ID under which the application runs is also different.

To prepare CICS and Tuxedo applications to run as WebSphere MQ MQI client applications, follow the instructions in the [Configuring an extended transactional client](#).

Note, however, that the information that deals specifically with preparing CICS and Tuxedo applications, including the sample programs supplied with WebSphere MQ, assumes that you are preparing applications to run on a WebSphere MQ server system. As a result, the information refers only to WebSphere MQ libraries that are intended for use on a server system. When you are preparing your client applications, you must do the following things:

- Use the appropriate client system library for the language bindings that your application uses. For example, for applications written in C on AIX, HP-UX, or Solaris, use the library libmqic instead of libmqm. On Windows systems, use the library mqic.lib instead of mqm.lib.
- Instead of the server system libraries shown in [Table 48 on page 354](#), for AIX, HP-UX, and Solaris, and [Table 49 on page 355](#), for Windows systems, use the equivalent client system libraries. If a server system library is not listed in these tables, use the same library on a client system.

<i>Table 48. Client system libraries on AIX, HP-UX, and Solaris</i>	
Library for a WebSphere MQ server system	Equivalent library to use on a WebSphere MQ client system
libmqmxa	libmqcxa

Table 49. Client system libraries on Windows systems

Library for a WebSphere MQ server system	Equivalent library to use on a WebSphere MQ client system
mqmxa.lib	mqcxa.lib
mqmtux.lib	mqcxa.lib
mqmenc.lib	mqcxa.lib
mqmcics4.lib	mqccics4.lib

The user ID used by a client application

When you run a WebSphere MQ server application under CICS, it normally switches from the CICS user to the user ID of the transaction. However, when you run a WebSphere MQ MQI client application under CICS, it retains the CICS privileged authority.

CICS and Tuxedo sample programs

CICS and Tuxedo sample programs for use on AIX, HP-UX, Solaris, and Windows systems.

Table 50 on page 355 lists the CICS and Tuxedo sample programs that are supplied for use on AIX, HP-UX, and Solaris client systems. Table 51 on page 355 lists the equivalent information for Windows client systems. The tables also list the files that are used for preparing and running the programs. For a description of the sample programs, see [“The CICS transaction sample”](#) on page 112 and [“TUXEDO samples”](#) on page 149.

Table 50. Sample programs for AIX, HP-UX, and Solaris client systems

Description	Source	Executable module
CICS program	amqscic0.ccs	amqscicc
Header file for the CICS program	amqscih0.h	-
Tuxedo client program to put messages	amqstxpx.c	-
Tuxedo client program to get messages	amqstxgx.c	-
Tuxedo server program for the two client programs	amqstxsx.c	-
UBBCONFIG file for the Tuxedo programs	ubbstxcx.cfg	-
Field table file for the Tuxedo programs	amqstxvx.flds	-
View description file for the Tuxedo programs	amqstxvx.v	-

Table 51. Sample programs for Windows client systems

Description	Source	Executable module
CICS transaction	amqscic0.ccs	amqscicc
Header file for the CICS transaction	amqscih0.h	-
Tuxedo client program to put messages	amqstxpx.c	-
Tuxedo client program to get messages	amqstxgx.c	-
Tuxedo server program for the two client programs	amqstxsx.c	-
UBBCONFIG file for the Tuxedo programs	ubbstxcx.cfg	-

Table 51. Sample programs for Windows client systems (continued)

Description	Source	Executable module
Field table file for the Tuxedo programs	amqstxvx.fld	-
View description file for the Tuxedo programs	amqstxvx.v	-
Makefile for the Tuxedo programs	amqstxmc.mak	-
ENVFILE file for the Tuxedo programs	amqstxen.env	-

Error message AMQ5203, as modified for CICS and Tuxedo applications

When you run CICS or Tuxedo applications that use an extended transactional client, you might see standard diagnostic messages. One of these has been modified for use with an extended transactional client

The messages that you might see in the WebSphere MQ error log files are documented in [Diagnostic messages: AMQ4000-9999](#). Message AMQ5203 has been modified for use with an extended transactional client. Here is the text of the modified message:

AMQ5203: An error occurred calling the XA interface.

Explanation

The error number is &2 where a value of 1 indicates the supplied flags value of &1 was invalid, 2 indicates that there was an attempt to use threaded and non-threaded libraries in the same process, 3 indicates that there was an error with the supplied queue manager name '&3', 4 indicates that the resource manager id of &1 was invalid, 5 indicates that an attempt was made to use a second queue manager called '&3' when another queue manager was already connected, 6 indicates that the Transaction Manager has been called when the application isn't connected to a queue manager, 7 indicates that the XA call was made while another call was in progress, 8 indicates that the xa_info string '&4' in the xa_open call contained an invalid parameter value for parameter name '&5', and 9 indicates that the xa_info string '&4' in the xa_open call is missing a required parameter, parameter name '&5'.

User response

Correct the error and try the operation again.

Preparing and running Microsoft Transaction Server applications

To prepare an MTS application to run as a WebSphere MQ MQI client application, follow these instructions as appropriate for your environment.

For general information about how to develop Microsoft Transaction Server (MTS) applications that access WebSphere MQ resources, see the section on MTS in the WebSphere MQ Help Center.

To prepare an MTS application to run as a WebSphere MQ MQI client application, do one of the following for each component of the application:

- If the component uses the C language bindings for the MQI, follow the instructions in [“Preparing C programs in Windows”](#) on page 439 but link the component with the library mqicxa.lib instead of mqic.lib.
- If the component uses the WebSphere MQ C++ classes, follow the instructions in [“Building C++ programs on Windows”](#) on page 624 but link the component with the library imqx23vn.lib instead of imqc23vn.lib.
- If the component uses the Visual Basic language bindings for the MQI, follow the instructions in the [“Preparing Visual Basic programs in Windows”](#) on page 443 but when you define the Visual Basic project, type MqType=3 in the **Conditional Compilation Arguments** field.

- If the component uses the WebSphere MQ Automation Classes for ActiveX (MQAX), define an environment variable, GMQ_MQ_LIB, with the value mqic32xa.dll .

You can define the environment variable from within your application, or you can define it so that its scope is system wide. However, defining it as system wide can cause any existing MQAX application, that does not define the environment variable from within the application, to behave incorrectly.

Preparing and running WebSphere MQ JMS applications

You can run WebSphere MQ JMS applications in client mode, with WebSphere Application Server as your transaction manager. You might see certain warning messages.

To prepare and run WebSphere MQ JMS applications in client mode, with WebSphere Application Server as your transaction manager, follow the instructions in [“Using WebSphere MQ classes for JMS” on page 684](#).

When you run a WebSphere MQ JMS client application, you might see the following warning messages:

MQJE080

Insufficient license units - run setmqcap

MQJE081

File containing the license unit information is in the wrong format - run setmqcap

MQJE082

File containing the license unit information could not be found - run setmqcap

User exits, API exits, and WebSphere MQ installable services

You can extend queue manager facilities by using user exits, API exits, or installable services. This topic contains links to information about using and developing these programs.

For an introduction to how you can use user exits, API exits, and installable services to extend queue manager facilities, see [Extending queue manager facilities](#).

For information on writing and compiling exits and installable services, see [“Writing and compiling exits and installable services” on page 357](#).

Related concepts

[Channel-exit programs for MQI channels](#)

Related reference


[API exit reference](#)

[Installable services interface reference information](#)

Writing and compiling exits and installable services

You can write and compile exits without linking to any IBM WebSphere MQ libraries on UNIX, Linux and Windows.

About this task

 This topic applies to Windows, UNIX and Linux systems only. For details on writing exits and installable services for other platforms, see the relevant platform specific topics.

If IBM WebSphere MQ is installed in a non-default location you must write and compile your exits without linking to any IBM WebSphere MQ libraries.

You can write and compile exits on Windows, UNIX and Linux systems without linking any of these IBM WebSphere MQ libraries:

- mqmzf
- mqm

- mqmvx
- mqmvxd
- mqic
- mqutl

Existing exits that are linked to these libraries continue to work, providing that on UNIX and Linux systems IBM WebSphere MQ is installed in the default location.

Procedure

1. Include the cmqec.h header file.

Including this header file automatically includes the cmqc.h, cmqxc.h and cmqzc.h header files.

2. Write the exit so that MQI and DCI calls are made through the MQIEP structure. For more information about the MQIEP structure, see [MQIEP structure](#).

- Installable services
 - Use the **Hconfig** parameter to point to the MQZEP call.
 - You must check that the first 4 bytes of **Hconfig** match the **StrucId** of the MQIEP structure before using the **Hconfig** parameter.
 - For more information about writing installable service components, see [MQIEP](#).
- API exits
 - Use the **Hconfig** parameter to point to the MQXEP call.
 - You must check that the first 4 bytes of **Hconfig** match the **StrucId** of the MQIEP structure before using the **Hconfig** parameter.
 - For more information about writing API exits, see [“Writing API exits” on page 371](#).
- Channel exits
 - Use the **pEntryPoints** parameter of the MQCXP structure to point to MQI and DCI calls.
 - You must check that the MQCXP version number is at version 8 or higher before using **pEntryPoints**.
 - For more information about writing channel exits, see [“Writing channel-exit programs” on page 381](#).
- Data conversion exits
 - Use the **pEntryPoints** parameter of the MQDXP structure to point to MQI and DCI calls.
 - You must check that the MQDXP version number is at version 2 or higher before using **pEntryPoints**.
 - You can use the **crtmqcvx** command and the amqsvfc0.c source file to create data conversion code that uses the **pEntryPoints** parameter. See [“Writing a data-conversion exit for WebSphere MQ for Windows” on page 402](#) and [“Writing a data-conversion exit for WebSphere MQ on UNIX and Linux systems” on page 398](#).
 - If you have existing data conversion exits that were generated using the **crtmqcvx** command, you must regenerate the exit using the updated command.
 - For more information about writing data conversion exits, see [“Writing data-conversion exits” on page 397](#).
- Pre-connect exits
 - Use the **pEntryPoints** parameter of the MQNXP structure to point to MQI and DCI calls.
 - You must check that the MQNXP version number is at version 2 or higher before using **pEntryPoints**.
 - For more information about writing pre-connect exits, see [“Referencing connection definitions using a pre-connect exit from a repository” on page 404](#).

- Publish exits
 - Use the **pEntryPoints** parameter of the MQPSXP structure to point to MQI and DCI calls.
 - You must check that the MQPSXP version number is at version 2 or higher before using **pEntryPoints**.
 - For more information about writing publish exits, see [“Writing and compiling publish exits” on page 406](#).
- Cluster workload exits
 - Use the **pEntryPoints** parameter of the MQWXP structure to point to MQXCLWLN calls.
 - You must check that the MQWXP version number is at version 4 or higher before using **pEntryPoints**.
 - For more information about writing cluster workload exits, see [“Writing and compiling cluster workload exits” on page 408](#).

For example, in a channel exit calling MQPUT:

```
pChannelExitParms -> pEntryPoints -> MQPUT_Call(pChannelExitParms -> Hconn,
                                                Hobj,
                                                &md,
                                                &pmo,
                                                messlen,
                                                buffer,
                                                &CompCode,
                                                &Reason);
```

Further examples can be seen in the [“Sample WebSphere MQ programs” on page 91](#).

3. Compile the exit:

- Do not link to the IBM WebSphere MQ libraries.
- Do not include an embedded RPath to any IBM WebSphere MQ libraries in your exit.
- For more information about compiling your exit, see one of the following topics:
 - API exits: [“Compiling API exits” on page 373](#).
 - Channel exits, publish exits, Cluster workload exits: [“Compiling channel exit programs on Windows, UNIX and Linux systems” on page 396](#).
 - Data conversion exits: [“Writing data-conversion exits” on page 397](#).

4. Put the exit in one of the following places:

- A path of your choosing that you fully qualify when configuring the exit
- The default exit path, in a specific installation directory. For example, *MQ_DATA_PATH/exits/installation2*.
- The default exit path

The default exit path is *MQ_DATA_PATH/exits* for 32 bit exits, and *MQ_DATA_PATH/exits64* for 64 bit exits. You can change these paths in the *qm.ini* or *mqclient.ini* file. For more information, see [Exit path](#). On Windows and Linux, you can use the WebSphere MQ Explorer to change the path:

- Right click the queue manager name
- Click **Properties...**
- Click **Exits**
- In the exits default path field, specify the path name of the directory that holds the exit program.

If an exit is placed in both a specific installation directory and in the default path directory, the specific installation directory exit is used by the installation of WebSphere MQ named in the path. For example, the exit is placed in */exits/installation2* and in */exits*, but not in */exits/installation1*. The WebSphere MQ installation *installation2* uses the exit from */exits/installation2*. The WebSphere MQ installation *installation1* uses the exit from the */exits* directory.

5. If necessary, configure the exit:

- Installable services: [“Configuring services and components” on page 367.](#)
- API exits: [“Configuring API exits” on page 376.](#)
- Channel exits: [“Configuring channel exits” on page 396.](#)
- Publish exits: [“Configuring publish exits” on page 407.](#)
- Pre-connect exits: [“PreConnect stanza of the client configuration file” on page 406.](#)

Installable services and components for UNIX, Linux and Windows

This section introduces the installable services and the functions and components associated with them. The interface to these functions is documented so that you, or software vendors, can supply components.

The main reasons for providing WebSphere MQ installable services are:

- To provide you with the flexibility of choosing whether to use components provided by WebSphere MQ products, or replace or augment them with others.
- To allow vendors to participate, by providing components that might use new technologies, without making internal changes to WebSphere MQ products.
- To allow WebSphere MQ to exploit new technologies faster and cheaper, and so provide products earlier and at lower prices.

Installable services and *service components* are part of the WebSphere MQ product structure. At the center of this structure is the part of the queue manager that implements the function and rules associated with the Message Queue Interface (MQI). This central part requires a number of service functions, called *installable services*, in order to perform its work. The installable services are:

- Authorization service
- Name service

Each installable service is a related set of functions implemented using one or more *service components*. Each component is invoked using a properly-architected, publicly-available interface. This enables independent software vendors and other third parties to provide installable components to augment or replace those provided by the WebSphere MQ products. [Table 52 on page 360](#) summarizes the services and components that can be used.

Table 52. Installable service components summary			
Installable service	Supplied component	Function	Requirements
Authorization service	object authority manager (OAM)	Provides authorization checking on commands and MQI calls. Users can write their own component to augment or replace the OAM. For example, to check that a user ID has authority to open a queue.	(Appropriate platform authorization facilities are assumed)
Name service	None	Provides support to the queue manager for looking up the name of the queue manager that owns a specified queue. • User defined Note: Shared queues must have their <i>Scope</i> attribute set to CELL.	• A third-party or user-written name manager

The installable services interface is described in [Installable services interface reference information](#).

Writing a service component

This section describes the relationship between services, components, entry points, and return codes.

Functions and components

Each service consists of a set of related functions. For example, the name service contains function for:

- Looking up a queue name and returning the name of the queue manager where the queue is defined
- Inserting a queue name into the service's directory
- Deleting a queue name from the service's directory

It also contains initialization and termination functions.

An installable service is provided by one or more service components. Each component can perform some or all of the functions that are defined for that service. For example, in WebSphere MQ for AIX, the supplied authorization service component, the OAM, performs all the available functions. See “Authorization service interface” on page 364 for more information. The component is also responsible for managing any underlying resources or software (for example, an LDAP directory) that it needs to implement the service. Configuration files provide a standard way of loading the component and determining the addresses of the functional routines that it provides.

Figure 72 on page 361 shows how services and components are related:

- A service is defined to a queue manager by stanzas in a configuration file.
- Each service is supported by supplied code in the queue manager. Users cannot change this code and therefore cannot create their own services.
- Each service is implemented by one or more components; these can be supplied with the product or user-written. Multiple components for a service can be invoked, each supporting different facilities within the service.
- Entry points connect the service components to the supporting code in the queue manager.

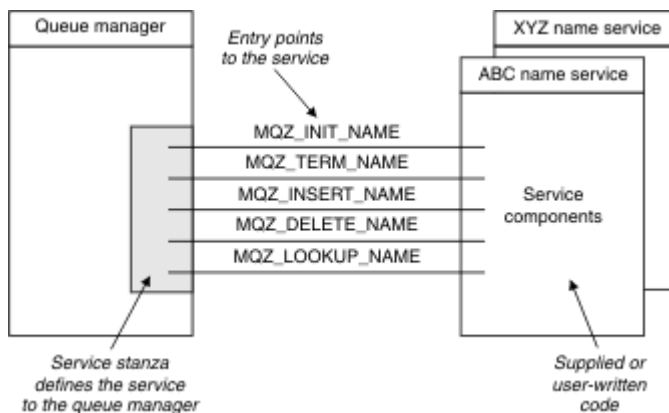


Figure 72. Understanding services, components, and entry points

Entry-points

Each service component is represented by a list of the entry-point addresses of the routines that support a particular installable service. The installable service defines the function to be performed by each routine.

The ordering of the service components when they are configured defines the order in which entry-points are called in an attempt to satisfy a request for the service.

In the supplied header file `cmqzc.h`, the supplied entry points to each service have an `MQZID_` prefix.

If the services are present, the services are loaded in a predefined order. The following list shows the services, and the order in which they are initialized.

1. NameService
2. AuthorizationService
3. UserIDentifierService

The AuthorizationService is the only service that is configured by default. Configure the NameService and UserIDentifierService manually if you want to use them.

Services, and service components have a one-to-one or one-to-many mapping. Multiple service components can be defined for each service. On UNIX and Linux systems, the ServiceComponent stanza's Service value must match the Service stanza's Name value in the qm.ini file. On Windows, the ServiceComponent's Service registry key value must match the Name registry key value, and is defined as:

HKEY_LOCAL_MACHINE\SOFTWARE\IBM\MQSeries\CurrentVersion\Configuration\QueueManager\qmname\ where qmname is the name of the queue manager.

For UNIX and Linux systems, service components are started in the order they are defined in the qm.ini file. On Windows, because the Windows registry is used, WebSphere MQ issues a **RegEnumKey** call which returns the values in alphabetic order. Therefore, on Windows the services are called in alphabetic order, as they are defined in the registry.

The ordering of the ServiceComponent definitions is significant. This ordering dictates the order in which components are run for a given service. For example, the AuthorizationService on Windows is configured with the default OAM component named MQSeries.WindowsNT.auth.service. Additional components can be defined for this service in order to override the default OAM. Unless MQCACF_SERVICE_COMPONENT is specified, the first component encountered in alphabetic order is used to process the request, and the name for that component is used.

Return codes

Service components provide return codes to the queue manager to report on various conditions. They report the success or failure of the operation, and indicate whether the queue manager is to proceed to the next service component. A separate *Continuation* parameter carries this indication.

Component data

A single service component might require data to be shared between its various functions. Installable services provide an optional data area to be passed on each invocation of a service component. This data area is for the exclusive use of the service component. It is shared by all the invocations of a particular function, even if they are made from different address spaces or processes. It is guaranteed to be addressable from the service component whenever it is called. You must declare the size of this area in the *ServiceComponent* stanza.

Initialization and termination of components

The use of component initialization and termination options.

When the component initialization routine is invoked, it must call the queue manager **MQZEP** function for each entry-point supported by the component. **MQZEP** defines an entry-point to the service. All the undefined exit points are assumed to be NULL.

A component is always invoked once with the primary initialization option, before it is invoked in any other way.

A component can be invoked with the secondary initialization option on certain platforms. For example, it can be invoked once for each operating system process, thread, or task by which the service is accessed.

If secondary initialization is used:

- The component can be invoked more than once for secondary initialization. For each such call, a matching call for secondary termination is issued when the service is no longer needed.

For naming services this is the MQZ_TERM_NAME call.

For authorization services this is the MQZ_TERM_AUTHORITY call.

- The entry points must be re-specified (by calling MQZEP) each time the component is called for primary and secondary initialization.
- Only one copy of component data is used for the component; there is not a different copy for each secondary initialization.
- The component is not invoked for any other calls to the service (from the operating system process, thread, or task, as appropriate) before secondary initialization has been carried out.
- The component must set the *Version* parameter to the same value for primary and secondary initialization.

The component is always invoked with the primary termination option once, when it is no longer required. No further calls are made to this component.

The component is invoked with the secondary termination option, if it has been invoked for secondary initialization.

Object authority manager (OAM)

The authorization service component supplied with the WebSphere MQ products is called the Object Authority Manager (OAM).

By default, the OAM is active and works with the control commands **dspmqa** (display authority), **dmpmqaut** (dump authority), and **setmqaut** (set or reset authority).

The syntax of these commands and how to use them are described in [The control commands](#).

The OAM works with the *entity* of a principal or group.

- On UNIX and Linux systems:
 - the principal is a user ID, or an ID associated with an application program running on behalf of a user.
 - the group is a UNIX or Linux system-defined collection of principals.
 - Authorizations can be granted or revoked at the group level only. A request to grant or revoke a user's authority updates the primary group for that user.
- On Windows systems:
 - the principal is a Windows user ID, or an ID associated with an application program running on behalf of a user.
 - the group is a Windows group.
 - Authorizations can be granted or revoked at the principal or group level.

When an MQI request is made or a command is issued, the OAM checks the authorization of the entity associated with the operation to see whether it can:

- Perform the requested operation.
- Access the specified queue manager resources.

The authorization service enables you to augment or replace the authority checking provided for queue managers by writing your own authorization service component.

Name service

The name service is an installable service that provides support to the queue manager for looking up the name of the queue manager that owns a specified queue. No other queue attributes can be retrieved from a name service.

The name service enables an application to open remote queues for output as if they were local queues. A name service is not invoked for objects other than queues.

Note: The remote queues **must** have their *Scope* attribute set to CELL.

When an application opens a queue, it looks for the name of the queue first in the queue manager's directory. If it does not find it there, it looks in as many name services as have been configured, until it finds one that recognizes the queue name. If none recognizes the name, the open fails.

The name service returns the owning queue manager for that queue. The queue manager then continues with the MQOPEN request as if the command had specified the queue and queue manager name in the original request.

The name service interface (NSI) is part of the WebSphere MQ framework.

How the name service works

If a queue definition specifies the *Scope* attribute as queue manager, that is, SCOPE(QMGR) in MQSC, the queue definition (along with all the queue attributes) is stored in the queue manager's directory only. This cannot be replaced by an installable service.

If a queue definition specifies the *Scope* attribute as cell, that is, SCOPE(CELL) in MQSC, the queue definition is again stored in the queue manager's directory, along with all the queue attributes. However, the queue and queue-manager name are also stored in a name service. If no service is available that can store this information, a queue with the *Scope* cell cannot be defined.

The directory in which the information is stored can be managed by the service, or the service can use an underlying service, for example, an LDAP directory, for this purpose. In either case, definitions stored in the directory must persist, even after the component and queue manager have terminated, until they are explicitly deleted.

Note:

1. To send a message to a remote host's local queue definition (with a scope of CELL) on a different queue manager within a naming directory cell, you need to define a channel.
2. You cannot get messages directly from the remote queue, even when it has a scope of CELL.
3. No remote queue definition is required when sending to a queue with a scope of CELL.
4. The naming service centrally defines the destination queue, although you still need a transmission queue to the destination queue manager and a pair of channel definitions. In addition, the transmission queue on the local system must have the same name as the queue manager owning the target queue, with the scope of cell, on the remote system.

For example, if the remote queue manager has the name QM01, the transmission queue on the local system must also have the name QM01.

Authorization service interface

The authorization service provides entry points for use by the queue manager.

The entry points are as follows:

MQZ_AUTHENTICATE_USER

Authenticates a user ID and password, and can set identity context fields.

MQZ_CHECK_AUTHORITY

Checks whether an entity has authority to perform one or more operations on a specified object.

MQZ_CHECK_PRIVILEGED

Checks whether a specified user is a privileged user.

MQZ_COPY_ALL_AUTHORITY

Copies all the current authorizations that exist for a referenced object to another object.

MQZ_DELETE_AUTHORITY

Deletes all authorizations associated with a specified object.

MQZ_ENUMERATE_AUTHORITY_DATA

Retrieves all the authority data that matches the selection criteria specified.

MQZ_FREE_USER

Frees associated allocated resources.

MQZ_GET_AUTHORITY

Gets the authority that an entity has to access a specified object.

MQZ_GET_EXPLICIT_AUTHORITY

Gets either the authority that a named group has to access a specified object (but without the additional authority of the **nobody** group) or the authority that the primary group of the named principal has to access a specified object.

MQZ_INIT_AUTHORITY

Initializes authorization service component.

MQZ_INQUIRE

Queries the supported functionality of the authorization service.

MQZ_REFRESH_CACHE

Refresh all authorizations.

MQZ_SET_AUTHORITY

Sets the authority that an entity has to a specified object.

MQZ_TERM_AUTHORITY

Terminates authorization service component.

In addition, on WebSphere MQ for Windows, the authorization service provides the following entry points for use by the queue manager:

- **MQZ_CHECK_AUTHORITY_2**
- **MQZ_GET_AUTHORITY_2**
- **MQZ_GET_EXPLICIT_AUTHORITY_2**
- **MQZ_SET_AUTHORITY_2**

These entry points support the use of the Windows Security Identifier (NT SID).

These names are defined as **typedefs**, in the header file `cmqzc.h`, which can be used to prototype the component functions.

The initialization function (**MQZ_INIT_AUTHORITY**) must be the main entry point for the component. The other functions are invoked through the entry point address that the initialization function has added into the component entry point vector.

Name service interface

A name service provides entry points for use by the queue manager.

The following entry points are provided:

MQZ_INIT_NAME

Initialize the name service component.

MQZ_TERM_NAME

Terminate the name service component.

MQZ_LOOKUP_NAME

Look up the queue-manager name for the specified queue.

MQZ_INSERT_NAME

Insert an entry containing the owning queue-manager name for the specified queue into the directory used by the service.

MQZ_DELETE_NAME

Delete the entry for the specified queue from the directory used by the service.

If there is more than one name service configured:

- For lookup, the `MQZ_LOOKUP_NAME` function is invoked for each service in the list until the queue name is resolved (unless any component indicates that the search should stop).
- For insert, the `MQZ_INSERT_NAME` function is invoked for the first service in the list that supports this function.

- For delete, the MQZ_DELETE_NAME function is invoked for the first service in the list that supports this function.

Do not have more than one component that supports the insert and delete functions. However, a component that only supports lookup is feasible, and could be used, for example, as the last component in the list to resolve any name that is not known by any other name service component to a queue manager at which the name can be defined.

In the C programming language the names are defined as function data types using the typedef statement. These can be used to prototype the service functions, to ensure that the parameters are correct.

The header file that contains all the material specific to installable services is cmqzc.h for the C language.

Apart from the initialization function (MQZ_INIT_NAME), which must be the component's main entry point, functions are invoked by the entry point address that the initialization function has added, using the MQZEP call.

Using multiple service components

You can install more than one component for a service. This allows components to provide only partial implementations of the service, and to rely on other components to provide the remaining functions.

Example of using multiple components

Suppose you create two a name services components called ABC_name_serv and XYZ_name_serv.

ABC_name_serv

This component supports inserting a name in, or deleting a name from, the service directory, but does not support looking up a queue name.

XYZ_name_serv

This component supports looking up a queue name, but does not support inserting a name in, or deleting a name from, the service directory.

Component ABC_name_serv holds a database of queue names, and uses two simple algorithms to either insert, or delete, a name from the service directory.

Component XYZ_name_serv uses a simple algorithm that returns a fixed queue-manager name for any queue name with which it is invoked. It does not hold a database of queue names, and therefore does not support the insert and delete functions.

The components are installed on the same queue manager. The *ServiceComponent* stanzas are ordered so that component ABC_name_serv is invoked first. Any calls to insert or delete a queue in a component directory are handled by component ABC_name_serv; it is the only one that implements these functions. However, a lookup call that component ABC_name_serv cannot resolve is passed on to the lookup-only component, XYZ_name_serv. This component supplies a queue-manager name from its simple algorithm.

Omitting entry points when using multiple components

If you decide to use multiple components to provide a service, you can design a service component that does not implement certain functions. The installable services framework places no restrictions on which you can omit. However, for specific installable services, omission of one or more functions might be logically inconsistent with the purpose of the service.

Example of entry points used with multiple components

Table 53 on page 367 shows an example of the installable name service for which the two components have been installed. Each supports a different set of functions associated with this particular installable service. For insert function, the ABC component entry-point is invoked first. Entry points that have not been defined to the service (using **MQZEP**) are assumed to be NULL. An entry-point for initialization is

provided in the table, but this is not required because initialization is carried out by the main entry-point of the component.

When the queue manager has to use an installable service, it uses the entry-points defined for that service (the columns in [Table 53 on page 367](#)). Taking each component in turn, the queue manager determines the address of the routine that implements the required function. It then calls the routine, if it exists. If the operation is successful, any results and status information are used by the queue manager.

Table 53. Example of entry-points for an installable service		
Function number	ABC name service component	XYZ name service component
MQZID_INIT_NAME (Initialize)	ABC_initialize()	XYZ_initialize()
MQZID_TERM_NAME (Terminate)	ABC_terminate()	XYZ_terminate()
MQZID_INSERT_NAME (Insert)	ABC_Insert()	NULL
MQZID_DELETE_NAME (Delete)	ABC_Delete()	NULL
MQZID_LOOKUP_NAME (Lookup)	NULL	XYZ_Lookup()

If the routine does not exist, the queue manager repeats this process for the next component in the list. In addition, if the routine does exist but returns a code indicating that it could not perform the operation, the attempt continues with the next available component. Routines in service components might return a code that indicates that no further attempts to perform the operation should be made.

Configuring services and components

Configure service components using the queue manager configuration files, except on Windows systems, where each queue manager has its own stanza in the Registry.

1. Add stanzas to the queue manager configuration file to define the service to the queue manager and specify the location of the module.

Each service used must have a *Service* stanza, which defines the service to the queue manager.

For each component within a service, there must be a *ServiceComponent* stanza. This identifies the name and path of the module containing the code for that component.

For more information, see [“Service stanza format” on page 367](#) and [“Service component stanza format” on page 368](#)

The authorization service component, known as the Object Authority Manager (OAM), is supplied with the product. When you create a queue manager, the queue manager configuration file (or the Registry on Windows systems) is automatically updated to include the appropriate stanzas for the authorization service and for the default component (the OAM). For the other components, you must configure the queue manager configuration file manually.

The code for each service component is loaded into the queue manager when the queue manager is started, using dynamic binding, where this is supported on the platform.

2. Stop and restart the queue manager to activate the component.

Service stanza format

The Service stanza contains the name of the service and the number of entry-points defined for the service.

The format of the stanza is as follows:

```
Service:
  Name=<service_name>
  EntryPoints=<entries>
```

where:

<service_name>

The name of the service. This is defined by the service.

<entries>

The number of entry-points defined for the service. This includes the initialization and termination entry points.

Service stanza format for Windows systems

On Windows systems, the *Service* stanza includes a *SecurityPolicy* attribute.

The format of the stanza is:

```
Service:
  Name=<service_name>
  EntryPoints=<entries>
  SecurityPolicy=<policy>
```

where:

<service_name>

The name of the service. This is defined by the service.

<entries>

The number of entry-points defined for the service. This includes the initialization and termination entry points.

<policy>

NTSIDsRequired (the Windows Security Identifier), or Default. If you do not specify NTSIDsRequired, the Default value is used. This attribute is valid only if Name has a value of AuthorizationService.

See also [“Configuring authorization service stanzas: Windows systems” on page 369](#).

Service component stanza format

The format of the service component stanza is:

```
ServiceComponent:
  Service=<service_name>
  Name=<component_name>
  Module=<module_name>
  ComponentDataSize=<size>
```

where:

<service_name>

The name of the service. This must match the Name specified in a service stanza.

<component_name>

A descriptive name of the service component. This must be unique, and contain only the characters that are valid for the names of WebSphere MQ objects (for example, queue names). This name occurs in operator messages generated by the service. We recommend that you use a name starting with a company trademark or similar distinguishing string.

<module_name>

The name of the module to contain the code for this component.

<size>

The size in bytes of the component data area passed to the component on each call. Specify zero if no component data is required.

These two stanzas can occur in any order and the stanza keys under them can also occur in any order. For either of these stanzas, all the stanza keys must be present. If a stanza key is duplicated, the last one is used.

At startup time, the queue manager processes each service component entry in the configuration file in turn. It then loads the specified component module, invoking the entry-point of the component (which must be the entry-point for initialization of the component), passing it a configuration handle.

Configuring authorization service stanzas: UNIX and Linux systems

On UNIX and Linux systems, each queue manager has its own queue manager configuration file.

For example, the default path and file name of the queue manager configuration file for queue manager QMNAME is `/var/mqm/qmgrs/QMNAME/qm.ini`.

The *Service* stanza and the *ServiceComponent* stanza for the default authorization component are added to `qm.ini` automatically, but can be overridden by `mqsnoaut`. Any other *ServiceComponent* stanzas must be added manually.

For example, the following stanzas in the queue manager configuration file define two authorization service components on WebSphere MQ for AIX. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

```
Service:
  Name=AuthorizationService
  EntryPoints=13

ServiceComponent:
  Service=AuthorizationService
  Name=MQSeries.UNIX.auth.service
  Module=MQ_INSTALLATION_PATH/lib/amqzfu
  ComponentDataSize=0

ServiceComponent:
  Service=AuthorizationService
  Name=user.defined.authorization.service
  Module=/usr/bin/udas01
  ComponentDataSize=96
```

Figure 73. UNIX and Linux authorization service stanzas in `qm.ini`

The service component stanza (`MQSeries.UNIX.auth.service`) defines the default authorization service component, the OAM. If you remove this stanza and restart the queue manager, the OAM is disabled and no authorization checks are made.

Configuring authorization service stanzas: Windows systems

On WebSphere MQ for Windows each queue manager has its own stanza in the registry.

The *Service* stanza and the *ServiceComponent* stanza for the default authorization component are added to the Registry automatically, but can be overridden using `mqsnoaut`. Any other *ServiceComponent* stanzas must be added manually.

You can also add the `SecurityPolicy` attribute using the WebSphere MQ services. The `SsecurityPolicy` attribute applies only if the service specified on the *Service* stanza is the authorization service, that is, the default OAM. The `SecurityPolicy` attribute allows you to specify the security policy for each queue manager. The possible values are:

Default

Specify `Default` if you want the default security policy to take effect. If a Windows security identifier (NT SID) is not passed to the OAM for a particular user ID, an attempt is made to obtain the appropriate SID by searching the relevant security databases.

NTSIDsRequired

Requires that an NT SID is passed to the OAM when performing security checks.

For information about the *Service* stanza format, see [“Service stanza format for Windows systems” on page 368](#). For more general information about security, see [Setting up security on Windows, UNIX and Linux systems](#).

The service component stanza, `MQSeries.WindowsNT.auth.service` defines the default authorization service component, the OAM. If you remove this stanza and restart the queue manager, the OAM is disabled and no authorization checks are made.

Configuring name service stanzas: Unix and Linux systems

Put your short description here; used for first paragraph and abstract.

The following examples of UNIX and Linux configuration file stanzas for the name service specify a name service component provided by the (fictitious) ABC company.

```
# Stanza for name service
Service:
  Name=NameService
  EntryPoints=5

# Stanza for name service component, provided by ABC
ServiceComponent:
  Service=NameService
  Name=ABC.Name.Service
  Module=/usr/lib/abcname
  ComponentDataSize=1024
```

Figure 74. Name service stanzas in *qm.ini* (for UNIX and Linux systems)

Note: On Windows systems, name service stanza information is stored in the Registry.

Refreshing the OAM after changing a user's authorization

In WebSphere MQ, you can refresh the OAM's authorization group information immediately after changing a user's authorization group membership, reflecting changes made at the operating system level, without needing to stop and restart the queue manager. In order to do this, issue the **REFRESH SECURITY** command.

Note: When you change authorizations with the `setmqaut` command, the OAM implements such changes immediately.

Queue managers store authorization data on a local queue called `SYSTEM.AUTH.DATA.QUEUE`. This data is managed by `amqzfuma.exe`.

Related reference

[REFRESH SECURITY](#)

Writing and compiling API exits

API exits let you write code that changes the behavior of WebSphere MQ API calls, such as `MQPUT` and `MQGET`, and then insert that code immediately before or immediately after those calls.

Note: Not supported on WebSphere MQ for z/OS.

Why use API exits?

Each of your applications has a specific job to do, and its code should do that task as efficiently as possible. At a higher level, you might want to apply standards or business processes to a particular queue manager for **all** the applications that use that queue manager. It is more efficient to do this above the level of individual applications, and thus without having to change the code of each application affected.

Here are a few suggestions of areas in which API exits might be useful:

- For *security*, you can provide authentication, checking that applications are authorized to access a queue or queue manager. You can also police applications' use of the API, authenticating the individual API calls, or even the parameters they use.
- For *flexibility*, you can respond to rapid changes in your business environment without changing the applications that rely on the data in that environment. You could, for example, have API exits that respond to changes in interest rates, currency exchange rates, or the price of components in a manufacturing environment.
- For *monitoring* use of a queue or queue manager, you can trace the flow of applications and messages, log errors in the API calls, set up audit trails for accounting purposes, or collect usage statistics for planning purposes.

What happens when an API exit runs?

Once you have written an exit program and identified it to WebSphere MQ, the queue manager automatically invokes your exit code at the registered points.

The API exit routines to run are identified in stanzas on IBM i, Windows, UNIX and Linux systems. This topic covers the stanzas in the configuration files mqs.ini and qm.ini.

The definition of the routines can occur in three places:

1. ApiExitCommon, in the mqs.ini file, identifies routines, for the whole of WebSphere MQ, applied when queue managers start. These can be overridden by routines defined for individual queue managers (see item “3” on page 371 in this list).
2. ApiExitTemplate, in the mqs.ini file, identifies routines, for the whole of WebSphere MQ, copied to the ApiExitLocal set (see item “3” on page 371 in this list) when a new queue manager is created.
3. ApiExitLocal, in the qm.ini file, identifies routines that apply to a particular queue manager.

When a new queue manager is created, the ApiExitTemplate definitions in mqs.ini are copied to the ApiExitLocal definitions in qm.ini for the new queue manager. When a queue manager is started, both the ApiExitCommon and ApiExitLocal definitions are used. The ApiExitLocal definitions replace the ApiExitCommon definitions if both identify a routine of the same name. The Sequence attribute, described in [“Configuring API exits” on page 376](#) determines the order in which the routines defined in the stanzas run.

Using API exits across multiple installations of WebSphere MQ

Ensure that the API exits written for the earlier version of WebSphere MQ are used to work with all versions because the changes made to exits in version 7.1 might not work with an earlier version. For more information about the changes made to exits, see [“Writing and compiling exits and installable services” on page 357](#).

The samples provided for API exits amqsaem and amqsaxe reflect the changes required while writing exits. The client application must ensure that the correct WebSphere MQ libraries that correspond to the installation of the queue manager with which the application is associated are linked to it prior to the launch of the application.

Writing API exits

You can write exits for every API call using the C programming language.

Exits are available for every API call, as follows:

- MQCB, to reregister a callback for the specified object handle and control activation and changes to the callback
- MQCTL, to perform controlling actions on the object handles opened for a connection
- MQCONN/MQCONN, to provide a queue manager connection handle for use on subsequent API calls
- MQDISC, to disconnect from a queue manager
- MQBEGIN, to begin a global unit of work (UOW)
- MQBACK, to back out a UOW
- MQCMIT, to commit a UOW
- MQOPEN, to open a WebSphere MQ resource for subsequent access
- MQCLOSE, to close a WebSphere MQ resource that had previously been opened for access
- MQGET, to retrieve a message from a queue that has previously been opened for access
- MQPUT1, to place a message on to a queue
- MQPUT, to place a message on to a queue that has previously been opened for access
- MQINQ, to inquire on the attributes of a WebSphere MQ resource that has previously been opened for access

- MQSET, to set the attributes of a queue that has previously been opened for access
- MQSTAT, to retrieve status information
- MQSUB, to register the applications subscription to a particular topic
- MQSUBRQ, to make a request for a subscription

MQ_CALLBACK_EXIT provides an exit function to perform before and after callback processing. For more information, see [Callback - MQ_CALLBACK_EXIT](#).

Within API exits, the calls take the general form:

```
MQ_call_EXIT (parameters, context, ApiCallParameters)
```

where *call* is the MQI call name without the MQ prefix; for example, PUT, GET. The *parameters* control the function of the exit, primarily providing communication between the exit and the external control blocks MQAXP (the API exit parameter structure) and MQAXC (the API exit context structure). *context* describes the context in which the API exit was called, and *ApiCallParameters* represent the parameters to the MQI call.

To help you write your API exit, a sample exit, amqsaxe0.c, is provided; this exit generates trace entries to a file that you specify. You can use this sample as your starting point when writing exits. For more information about using the sample exit, see [“The API exit sample program” on page 108](#).

For more information about the API exit calls, external control blocks, and associated topics, see [API exit reference](#).

For general information on how to write, compile and configure an exit, see [“Writing and compiling exits and installable services” on page 357](#).

Using message handles in API exits

You can control which message properties an API exit has access to. Properties are associated with an ExitMsgHandle. Properties set in a put exit are set on the message being put, but properties retrieved in a get exit are not returned to the application.

When you register an MQ_INIT_EXIT exit function using the MQXEP MQI call with **Function** set to MQXF_INIT and **ExitReason** set to MQXR_CONNECTION, you pass in an MQXEPO structure as the **ExitOpts** parameter. The MQXEPO structure contains the ExitProperties field, which specifies the set of properties to be made available to the exit. It is specified as a character string representing the prefix of the properties, which corresponds to an MQRFH2 folder name.

Each API exit receives an MQAXP structure, containing an ExitMsgHandle field. This field is set to a value generated by WebSphere MQ and is specific to a connection. The handle is therefore unchanged between API exits of the same or different types on the same connection.

In an MQ_PUT_EXIT or MQ_PUT1_EXIT with an **ExitReason** of MQXR_BEFORE, that is, an API exit performed before putting a message, any properties (other than message descriptor properties) associated with the ExitMsgHandle when the exit completes are set on the message being put. To prevent this happening, set ExitMsgHandle to MQHM_NONE. You can also supply a different message handle.

In an MQ_GET_EXIT, the ExitMsgHandle is cleared of properties and populated with the properties specified in the ExitProperties field when the MQ_INIT_EXIT was registered, other than message descriptor properties. These properties are not made available to the getting application. If the getting application specified a message handle in the MQGMO (Get message options) field, then any properties associated with that handle, including message descriptor properties, are available to the API exit. To prevent the ExitMsgHandle being populated with properties, set it to MQHM_NONE.

A sample program, amqsaem0.c, is provided to illustrate the use of message handles in API exits.

Compiling API exits

After you have written an exit, you compile and link it as follows.

The following examples show the commands used for the sample program described in “The API exit sample program” on page 108. For platforms other than Windows systems, you can find the sample API exit code in `MQ_INSTALLATION_PATH/samp` and the compiled and linked shared library in `MQ_INSTALLATION_PATH/samp/bin`. For Windows systems, you can find the sample API exit code in `MQ_INSTALLATION_PATH\Tools\c\Samples`. `MQ_INSTALLATION_PATH` represents the directory in which WebSphere MQ was installed.

Note to users:

1. Guidance on programming 64 bit applications is listed in [Coding standards on 64-bit platforms](#)

With the introduction of Multicast clients, API exits and data-conversion exits need to be able to run on the client-side because some messages might not go through the queue manager. The following libraries are now part of the client packages as well as the server packages:

Table 54. Libraries that are now in the client and server packages	
Operating system	Libraries
Windows	32 bit & 64 bit: mqm.dll & mqm.pdb
Linux & HP-UX	32 bit & 64 bit: libmqm.so & libmqm_r.so
AIX	32 bit & 64 bit: libmqm.a & libmqm_r.a
Solaris	32 bit & 64 bit: libmqm.so

Compiling API exits on Unix and Linux systems

Examples of how to Compile API exits on UNIX and Linux systems.

On all platforms, the entry point to the module is MQStart.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

On AIX

Compile the API exit source code by issuing one of the following commands:

32 bit applications

Non-threaded

```
cc -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits/amqsaxe \
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits/amqsaxe_r \
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
cc -q64 -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits64/amqsaxe \
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -q64 -e MQStart -bE:amqsaxe.exp -bM:SRE -o /var/mqm/exits64/amqsaxe_r \
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

On HP-UX Itanium platform

32 bit applications

Non-threaded

Compile the API Exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld +b: -b amqsaxe.o +ee MQStart -o /var/mqm/exits/amqsaxe  
rm amqsaxe.o
```

Threaded

Compile the API Exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld +b: -b amqsaxe.o +ee MQStart -o /var/mqm/exits/amqsaxe_r  
rm amqsaxe.o
```

64 bit applications

Non-threaded

Compile the API Exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld -b amqsaxe.o +ee MQStart -o /var/mqm/exits64/amqsaxe  
rm amqsaxe.o
```

Threaded

Compile the API Exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o amqsaxe.o amqsaxe0.c -IMQ_INSTALLATION_PATH/inc
```

Link the API Exit source code

```
ld -b amqsaxe.o +ee MQStart -o /var/mqm/exits64/amqsaxe_r  
rm amqsaxe.o
```

On Linux

Compile the API exit source code by issuing one of the following commands:

31 bit applications

Non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/amqsaxe amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/amqsaxe_r amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

32 bit applications

Non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/amqsaxe amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/amqsaxe_r amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsaxe amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/amqsaxe_r amqsaxe0.c \  
-IMQ_INSTALLATION_PATH/inc
```

On Solaris

Compile the API exit source code by issuing one of the following commands:

32 bit applications

SPARC platform

```
cc -xarch=v8plus -KPIC -mt -G -o /var/mqm/exits/amqsaxe \  
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc \  
-R/usr/lib/32 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=386 -KPIC -mt -G -o /var/mqm/exits/amqsaxe \  
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc \  
-R/usr/lib/32 -lsocket -lnsl -ldl
```

64 bit applications

SPARC platform

```
cc -xarch=v9 -KPIC -mt -G -o /var/mqm/exits64/amqsaxe \  
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc \  
-R/usr/lib/64 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=amd64 -KPIC -mt -G -o /var/mqm/exits64/amqsaxe \  
amqsaxe0.c -IMQ_INSTALLATION_PATH/inc \  
-R/usr/lib/64 -lsocket -lnsl -ldl
```

On Windows systems

Compile and link the sample API exit program, `amqsaxe0.c`, on Windows

A manifest file is an optional XML document containing the version, or any other, information that can be embedded in a compiled application or DLL.

If you have no such document, omit the `-manifest manifest.file` parameter in the `mt` command.

Adapt the commands in the examples in [Figure 75 on page 376](#) or [Figure 76 on page 376](#) to compile and link `amqsaxe0.c` on Windows. The commands work with Microsoft Visual Studio 2005, 2008, or 2010. The examples assume the WebSphere MQ `C:\Program Files\IBM\WebSphere MQ\tools\c\samples` directory is the current directory.

32 bit

```
cl /c /nologo /MD /Foamsaxe0.obj amqsaxe0.c
link /nologo /dll /def:amqsaxe.def
amqsaxe0.obj \
/manifest /out:amqsaxe.dll
mt -nologo -manifest amqsaxe.dll.manifest \
-outputresource:amqsaxe.dll;2
```

Figure 75. Compile and link *amqsaxe0.c* on 32 bit Windows

64 bit

```
cl /c /nologo /MD /Foamsaxe0.obj amqsaxe0.c
link /nologo /dll /def:amqsaxe.def \
/libpath:..\..\lib64 \

amqsaxe0.obj /manifest /out:amqsaxe.dll
mt -nologo -manifest amqsaxe.dll.manifest \
-outputresource:amqsaxe.dll;2
```

Figure 76. Compile and link *amqsaxe0.c* on 64 bit Windows

Related concepts

[“The API exit sample program” on page 108](#)

The sample API exit generates an MQI trace to a user-specified file with a prefix defined in the MQAPI_TRACE_LOGFILE environment variable.

Configuring API exits

You configure IBM WebSphere MQ to enable API exits by changing the configuration information.

To change the configuration information, you must change the stanzas that define the exit routines and the sequence in which they run. This information can be changed in the following ways:

- Using the IBM WebSphere MQ Explorer (On Windows and Linux (x86 and x86-64 platforms))
- Using the **amqmdain** command (On Windows)
- Using the mqs.ini and qm.ini files directly (On Windows, UNIX and Linux systems).

The mqs.ini file contains information relevant to all the queue managers on a particular node. You can find it in the /var/mqm directory on UNIX and Linux and in the WorkPath specified in the HKLM\SOFTWARE\IBM\WebSphere MQ key on Windows systems.

The qm.ini file contains information relevant to a specific queue manager. There is one queue manager configuration file for each queue manager, held in the root of the directory tree occupied by the queue manager. For example, the path and the name for a configuration file for a queue manager called QMNAME is:

On UNIX and Linux systems:

```
/var/mqm/qmgrs/QMNAME/qm.ini
```

On Windows systems:

```
C:\Program Files\IBM\WebSphere MQ\qmgrs\QMNAME\qm.ini
```

Before editing a configuration file, back it up so that you have a copy you can revert to if the need arises.

You can edit configuration files either:

- Automatically, using commands that change the configuration of queue managers on the node
- Manually, using a standard text editor

If you set an incorrect value on a configuration file attribute, the value is ignored and an operator message is issued to indicate the problem. (The effect is the same as missing out the attribute entirely.)

Stanzas to configure

The stanzas that must be changed are the following:

ApiExitCommon

Defined in mqs.ini and in the IBM WebSphere MQ Explorer on the IBM WebSphere MQ properties page, under Exits.

When any queue manager starts, the attributes in this stanza are read, and then overridden by the API exits defined in qm.ini.

ApiExitTemplate

Defined in mqs.ini and in the IBM WebSphere MQ Explorer on the IBM WebSphere MQ properties page, under Exits.

When any queue manager is created, the attributes in this stanza are copied into the newly created qm.ini file under the ApiExitLocal stanza.

ApiExitLocal

Defined in qm.ini and in the IBM WebSphere MQ Explorer on the queue manager properties page, under Exits.

When the queue manager starts, API exits defined here override the defaults defined in mqs.ini.

Attributes for the stanzas

- Name the API exit using the following attribute:

Name=ApiExit_name

The descriptive name of the API exit passed to it in the ExitInfoName field of the MQAXP structure.

This name must be unique, no longer than 48 characters, and contain only valid characters for the names of IBM WebSphere MQ objects (for example, queue names).

- Identify the module and entry point of the API exit code to run using the following attributes:

Function=function_name

The name of the function entry point into the module containing the API exit code. This entry point is the MQ_INIT_EXIT function.

The length of this field is limited to MQ_EXIT_NAME_LENGTH.

Module=module_name

The module containing the API exit code.

If this field contains the full path name of the module it is used as is.

If this field contains just the module name, the module is located using the ExitsDefaultPath attribute in the ExitPath in qm.ini.

On platforms that support separate threaded libraries, you must provide both a non-threaded and a threaded version of the API exit module. The threaded version must have an `_r` suffix. The threaded version of the IBM WebSphere MQ application stub implicitly appends `_r` to the given module name before it is loaded.

The length of this field is limited to the maximum path length the platform supports.

- Optionally pass data with the exit using the following attribute:

Data=data_name

Data to be passed to the API exit in the ExitData field of the MQAXP structure.

If you include this attribute, leading and trailing blanks are removed, the remaining string is truncated to 32 characters, and the result is passed to the exit. If you omit this attribute, the default value of 32 blanks is passed to the exit.

The maximum length of this field is 32 characters.

- Identify the sequence of this exit in relation to other exits using the following attribute:

Sequence=sequence_number

The sequence in which this API exit is called relative to other API exits. An exit with a low sequence number is called before an exit with a higher sequence number. There is no need for the sequence numbering of exits to be contiguous. A sequence of 1, 2, 3 has the same result as a sequence of 7, 42, 1096. If two exits have the same sequence number, the queue manager decides which one to call first. You can tell which was called after the event by putting the time or a marker in ExitChainArea indicated by the ExitChainAreaPtr in MQAXP or by writing your own log file.

This attribute is an unsigned numeric value.

Sample stanzas

The sample mqs.ini file contains the following stanzas:

ApiExitTemplate

This stanza defines an exit with the descriptive name OurPayrollQueueAuditor, module name auditor, and sequence number 2. A data value of 123 is passed to the exit.

ApiExitCommon

This stanza defines an exit with the descriptive name MQPoliceman, module name tmqp, and sequence number 1. The data passed is an instruction (CheckEverything).

```
mqs.ini

ApiExitTemplate:
  Name=OurPayrollQueueAuditor
  Sequence=2
  Function=EntryPoint
  Module=/usr/ABC/auditor
  Data=123
ApiExitCommon:
  Name=MQPoliceman
  Sequence=1
  Function=EntryPoint
  Module=/usr/MQPolice/tmqp
  Data=CheckEverything
```

The following sample qm.ini file contains an ApiExitLocal definition of an exit with the descriptive name ClientApplicationAPIchecker, module name ClientAppChecker, and sequence number 3.

```
qm.ini

ApiExitLocal:
  Name=ClientApplicationAPIchecker
  Sequence=3
  Function=EntryPoint
  Module=/usr/Dev/ClientAppChecker
  Data=9.20.176.20
```

Channel-exit programs for messaging channels

This collection of topics contains information about WebSphere MQ channel-exit programs for messaging channels.

Message channel agents (MCAs) can also call data-conversion exits. For more information about writing data-conversion exits, see [“Writing data-conversion exits” on page 397](#).

Some of this information also applies to exits on MQI channels, which connect WebSphere MQ MQI clients to queue managers. For more information, see [Channel-exit programs for MQI channels](#).

Channel-exit programs are called at defined places in the processing carried out by MCA programs.

Some of these user-exit programs work in complementary pairs. For example, if a user-exit program is called by the sending MCA to encrypt the messages for transmission, the complementary process must be functioning at the receiving end to reverse the process.

Table 55 on page 379 shows the types of channel exit that are available for each channel type.

<i>Table 55. Channel exits available for each channel type</i>						
Channel Type	Message exit	Message-retry exit	Receive exit	Security exit	Send exit	Auto-definition exit
Sender channel	Yes		Yes	Yes	Yes	
Server channel	Yes		Yes	Yes	Yes	
Cluster-sender channel	Yes		Yes	Yes	Yes	Yes
Receiver channel	Yes	Yes	Yes	Yes	Yes	Yes
Requester channel	Yes	Yes	Yes	Yes	Yes	
Cluster-receiver channel	Yes	Yes	Yes	Yes	Yes	Yes
Client-connection channel			Yes	Yes	Yes	
Server-connection channel			Yes	Yes	Yes	Yes

If you are going to run channel exits on a client, you cannot use the MQSERVER environment variable. Instead, create and reference a client channel definition table (CCDT) as described in [Client channel definition table](#).

Processing overview

An overview of how MCAs use channel-exit programs.

On startup, the MCAs exchange a startup dialog to synchronize processing. Then they switch to a data exchange that includes the security exits. These exits must end successfully for the startup phase to complete and to allow messages to be transferred.

The security check phase is a loop, as shown in [Figure 77 on page 380](#).




Figure 77. Security exit loop

During the message transfer phase, the sending MCA gets messages from a transmission queue, calls the message exit, calls the send exit, and then sends the message to the receiving MCA, as shown in [Figure 78 on page 380](#).




Figure 78. Example of a send exit at the sender end of message channel




Figure 79. Example of a receive exit at the receiver end of message channel

The receiving MCA receives a message from the communications link, calls the receive exit, calls the message exit, and then puts the message on the local queue, as shown in [Figure 79 on page 381](#). (The receive exit can be called more than once before the message exit is called.)

Writing channel-exit programs

You can use the following information to help you write channel-exit programs.

User exits and channel-exit programs can use all MQI calls, except as noted in the sections that follow. For MQ V7 and later, the MQCXP structure version 7 and higher contains the connection handle hConn, which can be used instead of issuing MQCONN. For earlier versions, to obtain the connection handle, an MQCONN must be issued, even though an MQRC_ALREADY_CONNECTED warning is returned because the channel itself is connected to the queue manager.

Note that the channel exit must be threadsafe.

For exits on client-connection channels, the queue manager to which the exit tries to connect depends on how the exit was linked. If the exit was linked with MQM.LIB and you do not specify a queue manager name on the MQCONN call, the exit tries to connect to the default queue manager on your system. If the exit was linked with MQM.LIB and you specify the name of the queue manager that was passed to the exit through the QMgrName field of MQCD, the exit tries to connect to that queue manager. If the exit was linked with MQIC.LIB or any other library, the MQCONN call fails whether you specify a queue manager name or not.

You should avoid altering the state of the transaction associated with the passed hConn in a channel exit; you must not use the MQCMIT, MQBACK or MQDISC verbs with the channel hConn, and you cannot use the MQBEGIN verb specifying the channel hConn.

If MQCONN is used specifying MQCNO_HANDLE_SHARE_BLOCK or MQCNO_HANDLE_SHARE_NO_BLOCK to create a new IBM WebSphere MQ connection, then it is your responsibility to ensure that the connection is correctly managed and disconnects from the queue manager correctly. For example, a channel exit that creates a new connection to the queue manager on every invocation without disconnecting, results in connection handles building up and an increase in the number of agent threads.

An exit runs in the same thread as the MCA itself and uses the same connection handle. So, it runs inside the same UOW as the MCA and any calls made under sync point are committed or backed out by the channel at the end of the batch.

Therefore, a channel message exit could send notification messages that are only committed to that queue when the batch containing the original message is committed. So, it is possible to issue sync point MQI calls from a channel message exit.

A channel exit can change fields in the MQCD. However, these changes are not acted on, except in the circumstances listed. If a channel exit program changes a field in the MQCD data structure, the new value is ignored by the IBM WebSphere MQ channel process. However, the new value remains in the MQCD and is passed to any remaining exits in an exit chain and to any conversation sharing the channel instance. For more information, see [Changing MQCD fields in a channel exit](#)

Also, for programs written in C, non-reentrant C library function must not be used in a channel-exit program.

If you use multiple channel exit libraries simultaneously, problems can arise on some UNIX and Linux platforms if the code for two different exits contains identically named functions. When a channel exit is loaded, the dynamic loader resolves function names in the exit library to the addresses where the library is loaded. If two exit libraries define separate functions which happen to have identical names, this resolution process might incorrectly resolve the function names of one library to use the functions of another. If this problem occurs, specify to the linker that it must only export the required exit and MQStart functions, as these functions are unaffected. Other functions must be given local visibility so that they are not used by functions outside their own exit library. Consult the documentation for the linker for more information.

All exits are called with a channel exit parameter structure (MQCXP), a channel definition structure (MQCD), a prepared data buffer, data length parameter, and buffer length parameter. The buffer length must not be exceeded:

- For message exits, you must allow for the largest message required to be sent across the channel, plus the length of the MQXQH structure.
- For send and receive exits, the largest buffer you must allow for is as follows:

LU 6.2

32 KB

TCP:

32 KB

Note: The maximum usable length might be 2 bytes less than this length. Check the value returned in MaxSegmentLength for details. For more information about MaxSegmentLength, see [MaxSegmentLength](#).

NetBIOS:

64 KB

SPX:

64 KB

Note: Receive exits on sender channels and sender exits on receiver channels use 2 KB buffers for TCP.

- For security exits, the distributed queuing facility allocates a buffer of 4000 bytes.

It is permissible for the exit to return an alternative buffer, together with the relevant parameters. See [“Channel-exit programs for messaging channels”](#) on page 378 for call details.

Writing channel-exit programs on Windows, UNIX and Linux systems

You can use the following information to help you write channel-exit programs for Windows, UNIX and Linux systems.

Follow the instructions outlined in [“Writing and compiling exits and installable services”](#) on page 357. Use the following channel exit specific information, where appropriate:

The exit must be written in C, and is a DLL on Windows.

Define a dummy MQStart() routine in the exit and specify MQStart as the entry point in the library. [Figure 80 on page 383](#) shows how to set up an entry to your program:

```
#include <cmqec.h>

void MQStart() {} /* dummy entry point - for consistency only */
void MQENTRY ChannelExit ( PMQCXP  pChannelExitParms,
                           PMQCD   pChannelDefinition,
                           PMQLONG pDataLength,
                           PMQLONG pAgentBufferLength,
                           PMQVOID pAgentBuffer,
                           PMQLONG pExitBufferLength,
                           PMQPTR  pExitBufferAddr)
{
    ... Insert code here
}
```

Figure 80. Sample source code for a channel exit

When writing channel exits for Windows using Visual C++, you must write your own DEF file. An example of how is shown in [Figure 81 on page 383](#). For further information on writing channel exit programs, see “Writing channel-exit programs” on page 381.

```
EXPORTS
ChannelExit
```

Figure 81. Sample DEF file for Windows

Channel security exit programs

You can use security exit programs to verify that the partner at the other end of a channel is genuine. This is known as authentication. To specify that a channel must use a security exit, specify the exit name in the SCYEXIT field of the channel definition.

Note: Authentication can also be achieved with channel authentication records. [Channel authentication records](#) provide great flexibility in preventing access to queue managers from certain users and channels, and in mapping remote users to IBM WebSphere MQ user identifiers. SSL and TLS support is also provided by IBM WebSphere MQ to authenticate your users and to provide encryption and data integrity checks for your data. For more information about SSL and TLS, see [WebSphere MQ support for SSL and TLS](#). However, if you still require more sophisticated (or different) forms of security processing, and other types of checks and security context establishment, consider writing security exits.

For security exits written prior to IBM WebSphere MQ Version 7.1 it is worth noting that earlier versions of IBM WebSphere MQ queried the underlying secure sockets provider (e.g. GSKit) to determine the remote partner's certificate Subject Distinguished Name (SSLPEER) and Issuer Distinguished Name (SSLCERTI). In IBM WebSphere MQ Version 7.1 support was added for a range of new security attributes. In order to access these attributes IBM WebSphere MQ Version 7.1 obtains the DER encoding of the certificate and uses it to determine the Subject and Issuer DN. The Subject and Issuer DN attributes appear in the following channel status attributes:

- SSLPEER (PCF selector MQCACH_SSL_SHORT_PEER_NAME)
- SSLCERTI (PCF selector MQCACH_SSL_CERT_ISSUER_NAME)

These values are returned by channel status commands as well as the data passed to channel security exits listed, as shown:

- MQCD SSLPeerNamePtr
- MQCXP SSLRemCertIssNamePtr

In IBM WebSphere MQ Version 7.1, a SERIALNUMBER attribute is also included in the Subject DN and contains the serial number for the remote partner's certificate. Also some DN attributes are returned in a different sequence from previous releases. Consequently the composition of the SSLPEER and SSLCERTI fields are altered in Version 7.1 from previous releases and it is therefore recommended that any security exits or applications dependant on these fields be examined and updated.

Existing WebSphere MQ peer name filters specified via the SSLPEER field of a channel definition are not affected and will continue to operate in the same manner as in earlier releases. This is because the WebSphere MQ peer name matching algorithm has been updated to process existing SSLPEER filters without any need to alter the channel definitions. This change is most likely to affect security exits and applications which depend upon the Subject DN and Issuer DN values returned by the PCF programming interface.

A security exit can be written in C or Java.

Channel security exit programs are called at the following places in the processing cycle of an MCA:

- At MCA initiation and termination.
- Immediately after the initial data negotiation is finished on channel startup. The receiver or server end of the channel can initiate a security message exchange with the remote end by providing a message to be delivered to the security exit at the remote end. It might also decline to do so. The exit program is started again to process any security message received from the remote end.
- Immediately after the initial data negotiation is finished on channel startup. The sender or requester end of the channel processes a security message received from the remote end, or initiates a security exchange when the remote end cannot. The exit program is started again to process all subsequent security messages that might be received.

A requester channel never gets called with MQXR_INIT_SEC. The channel notifies the server that it has a security exit program, and the server then has the opportunity to initiate a security exit. If it does not have one, it informs the requester and a zero length flow is returned to the exit program.

Note: Avoid sending zero-length security messages.

Examples of the data exchanged by security-exit programs are illustrated in figures [Figure 82 on page 385](#) through [Figure 85 on page 387](#). These examples show the sequence of events that occur involving the security exit of the receiver, and the security exit of the sender. Successive rows in the figures represent the passage of time. In some cases, the events at the receiver and sender are not correlated, and therefore can occur at the same time or at different times. In other cases, an event at one exit program results in a complementary event occurring later at the other exit program. For example, in [Figure 82 on page 385](#):

1. The receiver and sender are each invoked with MQXR_INIT, but these invocations are not correlated and can therefore occur at the same time or at different times.
2. The receiver is next invoked with MQXR_INIT_SEC, but returns MQXCC_OK which requires no complementary event at the sender exit.
3. The sender is next invoked with MQXR_INIT_SEC. This is not correlated with the invocation of the receiver with MQXR_INIT_SEC. The sender returns MQXCC_SEND_SEC_MSG, which causes a complementary event at the receiver exit.
4. The receiver is then invoked with MQXR_SEC_MSG, and returns MQXCC_SEND_SEC_MSG, which causes a complementary event at the sender exit.
5. The sender is then invoked with MQXR_SEC_MSG, and returns MQXCC_OK which requires no complementary event at the receiver exit.

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_OK	
	Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_SEND_SEC_MSG	
<div> <div></div> <div> <i>Message transfer begins</i> </div> </div>	Invoked with MQXR_SEC_MSG Responds with MQXCC_OK

Figure 82. Sender-initiated exchange with agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_OK	
	Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_OK	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_SUPPRESS_FUNCTION
<i>Channel closes</i>	
Invoked with MQXR_TERM Responds with MQXCC_OK	Invoked with MQXR_TERM Responds with MQXCC_OK

Figure 83. Sender-initiated exchange with no agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_SEND_SEC_MSG
Invoked with MQXR_SEC_MSG Responds with MQXCC_OK	
<i>Message transfer begins</i>	
Invoked with MQXR_TERM Responds with MQXCC_OK	Invoked with MQXR_TERM Responds with MQXCC_OK

Figure 84. Receiver-initiated exchange with agreement

Receiver exit	Sender exit
Invoked with MQXR_INIT Responds with MQXCC_OK	Invoked with MQXR_INIT Responds with MQXCC_OK
Invoked with MQXR_INIT_SEC Responds with MQXCC_SEND_SEC_MSG	
	Invoked with MQXR_SEC_MSG Responds with MQXCC_OK
Invoked with MQXR_SEC_MSG Responds with MQXCC_SUPPRESS_FUNCTION	
<i>Channel closes</i>	

Figure 85. Receiver-initiated exchange with no agreement

The channel security exit program is passed an agent buffer containing the security data, excluding any transmission headers, generated by the security exit. This data can be any suitable data so that either end of the channel is able to perform security validation.

The security exit program at both the sending and receiving end of the message channel can return either of two response codes to any call:

- Security exchange ended with no errors
- Suppress the channel and close down

Note:

1. The channel security exits typically work in pairs. When you define the appropriate channels, make sure that compatible exit programs are named for both ends of the channel.
2. In IBM i, security exit programs that have been compiled with "Use adopted authority" (USEADPAUT=*YES) can adopt QMQM or QMQMADM authority. Take care that the exit does not use this feature to pose a security risk to your system.
3. On an SSL channel on which the other end of the channel provides a certificate, the security exit receives the Distinguished Name of the subject of this certificate in the MQCD field accessed by SSLPeerNamePtr and the Distinguished Name of the issuer in the MQCXP field accessed by SSLRemCertIssNamePtr. Uses to which this name can be put are:
 - To restrict access over the SSL channel.
 - To change MQCD.MCAUserIdentifier based on the name.

Related concepts

[Channel authentication records](#)

[Secure Sockets Layer \(SSL\) and Transport Layer Security \(TLS\) concepts](#)

Writing a security exit

You can write a security exit by using the security exit skeleton code.

[Figure 86 on page 388](#) illustrates how to write a security exit.

```
void MQENTRY MQStart() {}
void MQENTRY EntryPoint (PMQVOID pChannelExitParms,
                        PMQVOID pChannelDefinition,
                        PMQLONG pDataLength,
                        PMQLONG pAgentBufferLength,
                        PMQVOID pAgentBuffer,
                        PMQLONG pExitBufferLength,
                        PMQPTR pExitBufferAddr)
{
    PMQCXP pParms = (PMQCXP)pChannelExitParms;
    PMQCD pChDef = (PMQCD)pChannelDefinition;
    /* TODO: Add Security Exit Code Here */
}
```

Figure 86. Security exit skeleton code

The standard WebSphere MQ Entry Point MQStart must exist, but is not required to perform any function. The name of the function (EntryPoint in this example) can be changed, but the function must be exported when the library is compiled and linked. As in the previous example, the pointers pChannelExitParms must be cast to PMQCXP and pChannelDefinition must be cast to PMQCD. For general information about calling channel exits and the use of parameters, see [MQ_CHANNEL_EXIT](#). These parameters are used in a security exit as follows:

PMQVOID pChannelExitParms

input/output

Pointer to MQCXP structure - cast to PMQCXP to access fields. This structure is used to communicate between the Exit and MCA. The following fields in the MQCXP are of particular interest for Security Exits:

ExitReason

Tells the Security Exit the current state in the security exchange and is used when deciding what action to take.

ExitResponse

The response to the MCA which dictates the next stage in the security exchange.

ExitResponse2

Extra control flags to govern how the MCA interprets the response of the Security Exit.

ExitUserArea

16 bytes (maximum) of storage which can be used by the Security Exit to maintain state between calls.

ExitData

Contains the data specified in the SCYDATA field of the channel definition (32 bytes padded to the right with blanks).

PMQVOID pChannelDefinition

input/output

Pointer to MQCD structure - cast to PMQCD to access fields. This parameter contains the definition of the channel. The following fields in the MQCD are of particular interest for Security Exits:

ChannelName

The channel name (20 bytes padded to the right with blanks).

ChannelType

A code defining the channel type.

MCA User Identifier

This group of three fields is initialized to the value of the MCAUSER field specified in the channel definition. Any user identifier specified by the Security Exit in these fields is used for access control (not applicable to SDR, SVR, CLNTCONN, or CLUSSDR channels).

MCAUserIdentifier

First 12 bytes of identifier padded to the right with blanks.

LongMCAUserIdPtr

Pointer to a buffer containing the full length identifier (not guaranteed null terminated) takes priority over MCAUserIdentifier.

LongMCAUserIdLength

Length of string pointed to by LongMCAUserIdPtr - must be set if LongMCAUserIdPtr is set.

Remote User Identifier

Only applies to CLNTCONN/SVRCONN channel pairs. If no CLNTCONN Security Exit is defined then these three fields are initialized by the client MCA, so they might contain a user identifier from the environment of the client which can be used by a SVRCONN Security Exit for authentication and when specifying the MCA User Identifier. If a CLNTCONN Security Exit is defined then these fields are not initialized and can be set by the CLNTCONN Security Exit, or security messages can be used to pass a user identifier from Client to Server.

RemoteUserIdentifier

First 12 Bytes of identifier padded to the right with blanks.

LongRemoteUserIdPtr

Pointer to a buffer containing the full length identifier (not guaranteed null terminated) takes priority over RemoteUserIdentifier.

LongRemoteUserIdLength

Length of string pointed to by LongRemoteUserIdPtr - must be set if LongRemoteUserIdPtr is set.

PMQLONG pDataLength

input/output

Pointer to MQLONG. Contains the length of any Security Exit contained in the AgentBuffer upon invocation of the Security Exit. Must be set by a Security Exit to the length of any message being sent in the AgentBuffer or ExitBuffer.

PMQLONG pAgentBufferLength

input

Pointer to MQLONG. The length of the data contained in the AgentBuffer on invocation of the Security Exit.

PMQVOID pAgentBuffer

input/output

On invocation of the Security Exit, this points to any message sent from the partner exit. If ExitResponse2 in the MQCXP structure has the MQXR2_USE_AGENT_BUFFER flag set (default) then a Security Exit needs to set this parameter to point to any message data being sent.

PMQLONG pExitBufferLength

input/output

Pointer to MQLONG. This parameter is initialized to 0 on the first invocation of a Security Exit and the value returned is maintained between calls to the Security Exit during a security exchange.

PMQPTR pExitBufferAddr

input/output

This parameter is initialized to a null pointer on the first invocation of a Security Exit and the value returned is maintained between calls to the Security Exit during a security exchange. If the MQXR2_USE_EXIT_BUFFER flag is set in the ExitResponse2 in the MQCXP structure then a Security Exit needs to set this parameter to point to any message data being sent.

Differences in behavior between security exits defined on CLNTCONN/SVRCONN channel pairs and other channel pairs

Security exits can be defined on all types of channel. However, the behavior of security exits defined on CLNTCONN/SVRCONN channel pairs is slightly different from security exits defined on other channel pairs.

A Security Exit on a CLNTCONN channel can set the Remote User Identifier in the channel definition for processing by a partner SVRCONN exit, or for OAM authorization if no SVRCONN Security Exit is defined and the MCAUSER field of the SVRCONN is not set.

If no CLNTCONN Security Exit is defined then the Remote User Identifier in the channel definition is set to a user identifier from the client environment (which can be blank) by the client MCA.

A security exchange between Security Exits defined on a CLNTCONN and SVRCONN channel pair completes successfully when the SVRCONN Security Exit returns an ExitResponse of MQXCC_OK. A security exchange between other channel pairs completes successfully when the Security Exit which initiated the exchange returns an ExitResponse of MQXCC_OK.

However, the MQXCC_SEND_AND_REQUEST_SEC_MSG ExitResponse code can be used to force continuation of the security exchange: If an ExitResponse of MQXCC_SEND_AND_REQUEST_SEC_MSG is returned by a CLNTCONN or SVRCONN Security Exit then the partner exit must respond by sending a security message (not MQXCC_OK or a null response) or the channel terminates. For Security Exits defined on other types of channel, an ExitResponse of MQXCC_OK returned in response to a MQXCC_SEND_AND_REQUEST_SEC_MSG from the partner Security Exit results in continuation of the security exchange as if a null response was returned and not in termination of the channel.

SSPI security exit

WebSphere MQ for Windows supplies a security exit that provides authentication for WebSphere MQ channels by using the Security Services Programming Interface (SSPI). The SSPI provides the integrated security facilities of Windows.

This security exit is for both the WebSphere MQ client and the WebSphere MQ server.

The security packages are loaded from either security.dll or secur32.dll. These DLLs are supplied with your operating system.

One-way authentication is provided on Windows, using NTLM authentication services. Two-way authentication is provided on Windows 2000, using Kerberos authentication services.

The security exit program is supplied in source and object format. You can use the object code as it is, or you can use the source code as a starting point to create your own user-exit programs. For more information about using the object or source code of the SSPI security exit, see [“Using the SSPI security exit on Windows systems” on page 162](#)

Channel send and receive exit programs

You can use the send and receive exits to perform tasks such as data compression and decompression. You can specify a list of send and receive exit programs to be run in succession.

Channel send and receive exit programs are called at the following places in the processing cycle of an MCA:

- The send and receive exit programs are called for initialization at MCA initiation and for termination at MCA termination.
- The send exit program is invoked at one or other end of the channel, depending on the end at which a transmission for one message transfer is sent, immediately before a transmission is sent over the link. Note 4 explains why exits are available in both directions even though message channels send messages in one direction only.
- The receive exit program is invoked at one or other end of the channel, depending on the end at which a transmission for one message transfer is received, immediately after a transmission has been taken from the link. Note 4 explains why exits are available in both directions even though message channels send messages in one direction only.

There might be many transmissions for one message transfer, and there could be many iterations of the send and receive exit programs before a message reaches the message exit at the receiving end.

The channel send and receive exit programs are passed an agent buffer containing the transmission data as sent or received from the communications link. For send exit programs, the first 8 bytes of the buffer are reserved for use by the MCA, and must not be changed. If the program returns a different buffer, then these first 8 bytes must exist in the new buffer. The format of data presented to the exit programs is not defined.

A good response code must be returned by send and receive exit programs. Any other response causes an MCA abnormal end (abend).

Note: Do not issue an MQGET, MQPUT, or MQPUT1 call within sync point from a send or receive exit.

Note:

1. Send and receive exits typically work in pairs. For example a send exit might compress the data and a receive exit decompress it, or a send exit might encrypt the data and a receive exit decrypt it. When you define the appropriate channels, make sure that compatible exit programs are named for both ends of the channel.
2. If compression is turned on for the channel, the exits are passed compressed data.
3. Channel send and receive exits might be called for message segments other than for application data, for example, status messages. They are not called during the startup dialog, nor the security check phase.
4. Although message channels send messages in one direction only, channel-control data, such as heart beats and end of batch processing, flows in both directions, and these exits are available in both directions, also. However, some of the initial channel startup data flows are exempt from processing by any of the exits.
5. There are circumstances in which send and receive exits could be invoked out of sequence; for example, if you are running a series of exit programs or if you are also running security exits. Then, when the receive exit is first called upon to process data, it might receive data that has not passed

through the corresponding send exit. If the receive exit just performed the operation, for example decompression, without first checking that it was required, the results would be unexpected.

You need to code your send and receive exits in such a way that the receive exit can check that the data it is receiving has been processed by the corresponding send exit. The recommended way to do so is to code your exit programs so that:

- The send exit sets the value of the ninth byte of data to 0 and shifts all the data along 1 byte, before performing the operation. (The first 8 bytes are reserved for use by the MCA.)
- If the receive exit receives data that has a 0 in byte 9, it knows that the data has come from the send exit. It removes the 0, performs the complementary operation, and shifts the resulting data back by 1 byte.
- If the receive exit receives data that has something other than 0 in byte 9, it assumes that the send exit has not run, and sends the data back to the caller unchanged.

When using security exits, if the channel is ended by the security exit it is possible that a send exit might be called without the corresponding receive exit. One way to prevent this problem is to code the security exit to set a flag, in MQCD.SecurityUserData or MQCD.SendUserData, for example, when the exit decides to end the channel. Then the send exit needs to check this field, and process the data only if the flag is not set. This check prevents the send exit from unnecessarily altering the data, and thus prevents any conversion errors that could occur if the security exit received altered data.

Channel send exit programs - reserving space

You can use send and receive exits to transform the data before transmission. Channel send exit programs can add their own data about the transformation by reserving space in the transmission buffer.

This data is processed by the receive exit program and then removed from the buffer. For example, you might want to encrypt the data and add a security key for decryption.

How you reserve space and use it

When the send exit program is called for initialization, set the *ExitSpace* field of MQXCP to the number of bytes to be reserved. See [MQXCP](#) for details. *ExitSpace* can be set only during initialization, that is when *ExitReason* has the value MQXR_INIT. When the send exit is invoked immediately before transmission, with *ExitReason* set to MQXR_XMIT, *ExitSpace* bytes are reserved in the transmission buffer. *ExitSpace* is not supported on z/OS.

The send exit need not use all the reserved space. It can use less than *ExitSpace* bytes or, if the transmission buffer is not full, the exit can use more than the amount reserved. When setting the value of *ExitSpace*, you must leave at least 1 KB for message data in the transmission buffer. Channel performance can be affected if reserved space is used for large amounts of data.

What happens at the receiving end of the channel

Channel receive exit programs must be set up to be compatible with the corresponding send exits. Receive exits must know the number of bytes in the reserved space and must remove the data in that space.

Multiple send exits

You can specify a list of send and receive exit programs to be run in succession. WebSphere MQ maintains a total for the space reserved by all the send exits. This total space must leave at least 1 KB for message data in the transmission buffer.

The following example shows how space is allocated for three send exits, called in succession:

1. When called for initialization:

- Send exit A reserves 1 KB.
- Send exit B reserves 2 KB.

- Send exit C reserves 3 KB.
2. The maximum transmission size is 32 KB and the user data is 5 KB long.
 3. Exit A is called with 5 KB of data; up to 27 KB are available, because 5 KB is reserved for exits B and C. Exit A adds 1 KB, the amount it reserved.
 4. Exit B is called with 6 KB of data; up to 29 KB are available, because 3 KB is reserved for exit C. Exit B adds 1 KB, less than the 2 KB it reserved.
 5. Exit C is called with 7 KB of data; up to 32 KB are available. Exit C adds 10K, more than the 3 KB it reserved. This amount is valid, because the total amount of data, 17 KB, is less than the 32 KB maximum.

Channel message exit programs

You can use the channel message exit to perform tasks such as encryption on the link, validation or substitution of incoming user IDs, message data conversion, journaling, and reference message handling. You can specify a list of message exit programs to be run in succession.

Channel message exit programs are called at the following places in the processing cycle of the MCA:

- At MCA initiation and termination
- Immediately after a sending MCA has issued an MQGET call
- Before the receiving MCA issues an MQPUT call

The message exit is passed an agent buffer containing the transmission queue header, MQXQH, and the application message text as retrieved from the queue. (The format of MQXQH is given in [MQXQH](#).) If you use reference messages; that is, messages that contain only a header which points to some other object that is to be sent, the message exit recognizes the header, MQRMH. It identifies the object, retrieves it in whatever way is appropriate appends it to the header, and passes it to the MCA for transmission to the receiving MCA. At the receiving MCA, another message exit recognizes that this message is a reference message, extracts the object, and passes the header on to the destination queue. See [“Reference messages” on page 254](#) and [“Running the Reference Message samples” on page 135](#) for more information about reference messages and some sample message exits that handle them.

Message exits can return the following responses:

- Send the message (GET exit). The message might have been changed by the exit. (This returns MQXCC_OK.)
- Put the message on the queue (PUT exit). The message might have been changed by the exit. (This returns MQXCC_OK.)
- Do not process the message. The message is placed on the dead-letter queue (undelivered message queue) by the MCA.
- Close the channel.
- Bad return code, which causes the MCA to abnormally end.

Note:

1. Message exits are called once for every complete message transferred, even when the message is split into parts.
2. In UNIX systems, if you provide a message exit for any reason the automatic conversion of user IDs to lowercase characters does not operate. See [Security of objects on UNIX and Linux systems](#).
3. An exit runs in the same thread as the MCA itself. It also runs inside the same unit of work (UOW) as the MCA because it uses the same connection handle. Therefore, any calls made under sync point are committed or backed out by the channel at the end of the batch. For example, one channel message exit program can send notification messages to another and these messages are only committed to the queue when the batch containing the original message is committed.

Therefore, it is possible to issue sync point MQI calls from a channel message exit program.

Message conversion outside the message exit

Before calling the message exit, the receiving MCA performs some conversions on the message. This topic describes the algorithms used to perform the conversions.

Which headers are processed

A conversion routine runs in the MCA of the receiver before the message exit is called. The conversion routine begins with the MQXQH header at the beginning of the message. The conversion routine then processes through the chained headers that follow the MQXQH, performing conversion where necessary. The chained headers can extend beyond the offset contained in the HeaderLength parameter of the MQCXP data that is passed to the message exit of the receiver. The following headers are converted in-place:

- MQXQH (format name "MQXMIT ")
- MQMD (this header is part of the MQXQH and has no format name)
- MQMDE (format name "MQHMDE ")
- MQDH (format name "MQHDIST ")
- MQWIH (format name "MQHWIH ")

The following headers are not converted, but are stepped over as the MCA continues to process the chained headers:

- MQDLH (format name "MQDEAD ")
- any headers with format names beginning with the three characters 'MQH' (for example "MQHRF ") that are not otherwise mentioned

How the headers are processed

The Format parameter of each WebSphere MQ header is read by the MCA. The Format parameter is 8 bytes within the header, which are 8 single-byte characters containing a name.

The MCA then interprets the data following each header as being of the named type. If the Format is the name of a header type eligible for WebSphere MQ data conversion, it is converted. If it is another name indicating non-MQ data (for example MQFMT_NONE or MQFMT_STRING) then the MCA stops processing the headers.

What is the MQCXP HeaderLength?

The HeaderLength parameter in the MQCXP data supplied to a message exit is the total length of the MQXQH (which includes the MQMD), MQMDE and MQDH headers at the start of the message. These headers are chained using the 'Format' names and lengths.

MQWIH

Chained headers can extend beyond the HeaderLength into the user data area. The MQWIH header, if it is present, is one of those headers that appear beyond the HeaderLength.

If there is an MQWIH header in the chained headers, it is converted in-place before the message exit of the receiver is called.

Channel message retry exit program

The channel message-retry exit is called when an attempt to open the target queue is unsuccessful. You can use the exit to determine under which circumstances to retry, how many times to retry, and how frequently.

This exit is also called at the receiving end of the channel at MCA initiation and termination.

The channel message-retry exit is passed an agent buffer containing the transmission queue header, MQXQH, and the application message text as retrieved from the queue. The format of MQXQH is given in [Overview for MQXQH](#).

The exit is invoked for all reason codes; the exit determines for which reason codes it wants the MCA to retry, for how many times, and at what intervals. (The value of the message-retry count set when the channel was defined is passed to the exit in the MQCD, but the exit can ignore this value.)

The MsgRetryCount field in MQCXP is incremented by the MCA each time the exit is invoked, and the exit returns either MQXCC_OK with the wait time contained in the MsgRetryInterval field of MQCXP, or MQXCC_SUPPRESS_FUNCTION. Retries continue indefinitely until the exit returns MQXCC_SUPPRESS_FUNCTION in the ExitResponse field of MQCXP. See [MQCXP](#) for information about the action taken by the MCA for these completion codes.

If all the retries are unsuccessful, the message is written to the dead-letter queue. If there is no dead-letter queue available, the channel stops.

If you do not define a message-retry exit for a channel and a failure occurs that is likely to be temporary, for example MQRC_Q_FULL, the MCA uses the message-retry count and message-retry intervals set when the channel was defined. If the failure is of a more permanent nature and you have not defined an exit program to handle it, the message is written to the dead-letter queue.

Channel auto-definition exit program

The channel auto-definition exit can be used when a request is received to start a receiver or server-connection channel but no definition for that channel exists (not for WebSphere MQ for z/OS). It can also be called on all platforms for cluster-sender and cluster-receiver channels to allow definition modification for an instance of the channel.

The channel auto-definition exit can be called on all platforms except z/OS when a request is received to start a receiver or server-connection channel but no channel definition exists. You can use it to modify the supplied default definition for an automatically defined receiver or server-connection channel, SYSTEM.AUTO.RECEIVER, or SYSTEM.AUTO.SVRCON. See [Preparing channels](#) for a description of how channel definitions can be created automatically.

The channel auto-definition exit can also be called when a request is received to start a cluster-sender channel. It can be called for cluster-sender and cluster-receiver channels to allow definition modification for this instance of the channel. In this case, the exit also applies to WebSphere MQ for z/OS. A common use of the channel auto-definition exit is to change the names of message exits (MSGEXIT, RCVEEXIT, SCYEXIT, and SENDEXIT) because exit names have different formats on different platforms. If no channel auto-definition exit is specified, the default behavior on z/OS is to examine a distributed exit name of the form *[path]/libraryname(function)* and take up to eight chars of function, if present, or libraryname. On z/OS, a channel auto-definition exit program must alter the fields addressed by MsgExitPtr, MsgUserDataPtr, SendExitPtr, SendUserDataPtr, ReceiveExitPtr, and ReceiveUserDataPtr, rather than the MsgExit, MsgUserData, SendExit, SendUserData, ReceiveExit and ReceiveUserData fields themselves.

For more information, see [Auto-definition of channels](#).

As with other channel exits, the parameter list is:

```
MQ_CHANNEL_AUTO_DEF_EXIT (ChannelExitParms, ChannelDefinition)
```

ChannelExitParms are described in [MQCXP](#). ChannelDefinition is described in [MQCD](#).

MQCD contains the values that are used in the default channel definition if they are not altered by the exit. The exit can modify only a subset of the fields; see [MQ_CHANNEL_AUTO_DEF_EXIT](#). However, attempting to change other fields does not cause an error.

The channel auto-definition exit returns a response of either MQXCC_OK or MQXCC_SUPPRESS_FUNCTION. If neither of these responses is returned, the MCA continues processing as though MQXCC_SUPPRESS_FUNCTION were returned. That is, the auto-definition is abandoned, no new channel definition is created, and the channel cannot start.

Compiling channel exit programs on Windows, UNIX and Linux systems

Use the following examples to help you compile channel-exit programs for Windows, UNIX and Linux systems.

Windows

Windows

The compiler and linker command for channel-exit programs on Windows:

```
cl.exe /Ic:\mqm\tools\c\include /nologo /c myexit.c
link.exe /nologo /dll myexit.obj /def:myexit.def /out:myexit.dll
```

UNIX and Linux systems

Linux

UNIX

In these examples `exit` is the library name and `ChannelExit` is the function name. On AIX the export file is called `exit.exp`. These names are used by the channel definition to reference the exit program using the format described in [MQCD- channel definition](#). See also the `MSGEXIT` parameter of the [DEFINE CHANNEL](#) command.

Sample compiler and linker commands for channel exits on AIX:

```
$ xlc_r -q64 -e MQStart -bE:exit.exp -bM:SRE -o /var/mqm/exits64/exit
exit.c -I/usr/mqm/inc
```

Sample compiler and linker commands for channel exits on HP-UX

```
$ c89 +DD64 +z -c -D_HPUX_SOURCE -o exit.o exit.c -I/opt/mqm/inc
$ ld -b exit.o +ee MQStart +ee ChannelExit -o
/var/mqm/exits64/exit -L/usr/lib/pa20_64 -lpthread
$ rm exit.o
```

Sample compiler and linker commands for channel-exits on Linux platforms where the queue manager is 32 bit:

```
$ gcc -shared -fPIC -o /var/mqm/exits/exit exit.c -I/opt/mqm/inc
```

Sample compiler and linker commands for channel-exits on Linux platforms where the queue manager is 64-bit:

```
$ gcc -m64 -shared -fPIC -o /var/mqm/exits64/exit exit.c -I/opt/mqm/inc
```

Sample compiler and linker commands for channel exits on Solaris:

```
$ cc -xarch=v9 -mt -G -o /var/mqm/exits64/exit exit.c -I/opt/mqm/inc
-R/usr/lib/64 -lsocket -lnsl -ldl
```

On the client, a 32 bit or 64 bit exit can be used. This exit must be linked to `mqic_r`.

On AIX, all functions that are called by IBM WebSphere MQ must be exported. A sample export file for this make file:

```
#!/
channelExit
MQStart
```

Configuring channel exits

To call the channel exit, you must name it in the channel definition.

Channel exits must be named in the channel definition. You can do this naming when you first define the channels, or you can add the information later using, for example, the `MQSC` command `ALTER CHANNEL`.

You can also give the channel exit names in the MQCD channel data structure. The format of the exit name depends on your IBM WebSphere MQ platform; see [MQCD](#) or [Script \(MQSC\) Commands](#) for information.

If the channel definition does not contain a user-exit program name, the user exit is not called.

The channel auto-definition exit is the property of the queue manager, not the individual channel. In order for this exit to be called, it must be named in the queue manager definition. To alter a queue manager definition, use the MQSC command ALTER QMGR.

Writing data-conversion exits

This collection of topics contains information about how to write data-conversion exits.

Note: Not supported in MQSeries for VSE/ESA.

When you do an MQPUT, your application creates the message descriptor (MQMD) of the message. Because WebSphere MQ needs to be able to understand the contents of the MQMD regardless of the platform it is created on, it is converted automatically by the system.

Application data, however, is not converted automatically. If character data is being exchanged between platforms where the *CodedCharSetId* and *Encoding* fields differ, for example, between ASCII and EBCDIC, the application must arrange for conversion of the message. Application data conversion can be performed by the queue manager itself or by a user exit program, referred to as a *data-conversion exit*. The queue manager can perform data conversion itself, using one of its built-in conversion routines, if the application data is in one of the built-in formats (such as MQFMT_STRING). This topic contains information about the data-conversion exit facility that WebSphere MQ provides for when the application data is not in a built-in format.

Control can be passed to the data-conversion exit during an MQGET call. This avoids converting across different platforms before reaching the final destination. However, if the final destination is a platform that does not support data conversion on the MQGET, you must specify CONVERT(YES) on the sender channel that sends the data to its final destination. This ensures that WebSphere MQ converts the data during transmission. In this case, your data-conversion exit must reside on the system where the sender channel is defined.

The MQGET call is issued directly by the application. Set the *CodedCharSetId* and *Encoding* fields in the MQMD to the character set and encoding required. If your application uses the same character set and encoding as the queue manager, set *CodedCharSetId* to MQCCSI_Q_MGR, and *Encoding* to MQENC_NATIVE. After the MQGET call completes, these fields have the values appropriate to the message data returned. These might differ from the values required if the conversion was not successful. Your application should reset these fields to the values required before each MQGET call.

The conditions required for the data-conversion exit to be called are defined for the MQGET call in [MQGET](#).

For a description of the parameters that are passed to the data-conversion exit, and detailed usage notes, see [Data conversion](#) for the MQ_DATA_CONV_EXIT call and the MQDXP structure.

Programs that convert application data between different machine encodings and CCSIDs must conform to the WebSphere MQ data conversion interface (DCI).

With the introduction of Multicast clients, API exits and data-conversion exits need to be able to run on the client-side because some messages might not go through the queue manager. The following libraries are now part of the client packages as well as the server packages:

Table 56. Libraries that are now in the client and server packages	
Operating system	Libraries
Windows	32 bit & 64 bit: mqm.dll & mqm.pdb
Linux & HP-UX	32 bit & 64 bit: libmqm.so & libmqm_r.so

Table 56. Libraries that are now in the client and server packages (continued)	
Operating system	Libraries
AIX	32 bit & 64 bit: libmqm.a & libmqm_r.a
Solaris	32 bit & 64 bit: libmqm.so

Invoking the data-conversion exit

A data-conversion exit is a user-written exit that receives control during the processing of an MQGET call.

The exit is invoked if the following are true:

- The MQGMO_CONVERT option is specified on the MQGET call.
- Some or all of the message data is not in the requested character set or encoding.
- The *Format* field in the MQMD structure associated with the message is not MQFMT_NONE.
- The *BufferLength* specified on the MQGET call is not zero.
- The message data length is not zero.
- The message contains data that has a user-defined format. The user-defined format can occupy the entire message, or be preceded by one or more built-in formats. For example, the user-defined format might be preceded by an MQFMT_DEAD_LETTER_HEADER format. The exit is invoked to convert only the user-defined format; the queue manager converts any built-in formats that precede the user-defined format.

A user-written exit can also be invoked to convert a built-in format, but this happens only if the built-in conversion routines cannot convert the built-in format successfully.

There are some other conditions, described fully in the usage notes of the MQ_DATA_CONV_EXIT call in [MQ_DATA_CONV_EXIT](#).

See [MQGET](#) for details of the MQGET call. Data-conversion exits cannot use MQI calls, other than MQXCNCV.

A new copy of the exit is loaded when an application attempts to retrieve the first message that uses that *Format* since the application connected to the queue manager. A new copy might also be loaded at other times if the queue manager has discarded a previously loaded copy.

The data-conversion exit runs in an environment like that of the program that issued the MQGET call. As well as user applications, the program can be an MCA (message channel agent) sending messages to a destination queue manager that does not support message conversion. The environment includes address space and user profile, where applicable. The exit cannot compromise the integrity of the queue manager, because it does not run in the queue manager's environment.

Writing a data-conversion exit for WebSphere MQ on UNIX and Linux systems

Information about steps to consider when writing data-conversion exit programs for WebSphere MQ on UNIX and Linux systems.

Follow these steps:

1. Name your message format. The name must fit in the *Format* field of the MQMD, and be in uppercase, for example, MYFORMAT. The *Format* name must not have leading blanks. Trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long. Remember to use this name each time that you send a message.

If the data conversion exit is used in a threaded environment, the loadable object must be followed by `_r` to indicate that it is a threaded version.

2. Create a structure to represent your message. See [Valid syntax](#) for an example.
3. Run this structure through the `crtmqcvx` command to create a code fragment for your data-conversion exit.

The functions generated by the `crtmqcvx` command use macros that assume that all structures are packed; amend them if this is not the case.

4. Copy the supplied skeleton source file, renaming it to the name of your message format that you set in step “1” on page 398. The skeleton source file, and the copy, are read-only.

The skeleton source file is called `amqsvfc0.c`.

5. On WebSphere MQ for AIX, a skeleton export file called `amqsvfc.exp` is also supplied. Copy this file, renaming it to `MYFORMAT.EXP`.
6. The skeleton includes a sample header file, `amqsvmha.h`, in the directory `MQ_INSTALLATION_PATH/inc`, where `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.. Make sure that your include path points to this directory to pick up this file.

The `amqsvmha.h` file contains macros that are used by the code generated by the `crtmqcvx` command. If the structure to be converted contains character data, these macros call `MQXCNCVC`.

7. Find the following comment boxes in the source file and insert code as described:
 - a. Toward the end of the source file, a comment box starts with:

```
/* Insert the functions produced by the data-conversion exit */
```

Here, insert the code fragment generated in step “3” on page 398.

- b. Near the middle of the source file, a comment box starts with:

```
/* Insert calls to the code fragments to convert the format's */
```

This is followed by a commented-out call to the function `ConverttagSTRUCT`.

Change the name of the function to the name of the function that you added in step “7.a” on page 399. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

- c. Near the beginning of the source file, a comment box starts with:

```
/* Insert the function prototypes for the functions produced by */
```

Here, insert the function prototype statements for the functions added in step “3” on page 398 above.

8. Compile your exit as a shared library, using `MQStart` as the entry point. To do this, see “[Compiling data-conversion exits on UNIX and Linux systems](#)” on page 399.
9. Place the output in the exit directory. The default exit directory is `/var/mqm/exits` for 32 bit systems and `/var/mqm/exits64`, for 64 bit systems. You can change these directories in the `qm.ini` or `mqclient.ini` file. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Note:

1. If `crtmqcvx` uses packed structures, all WebSphere MQ applications must be compiled in this way.
2. Data-conversion exit programs must be reentrant.
3. `MQXCNCVC` is the *only* MQI call that can be issued from a data-conversion exit.

Compiling data-conversion exits on UNIX and Linux systems

Examples of how to compile a data conversion exit on UNIX and Linux systems.

On all platforms, the entry point to the module is `MQStart`.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

AIX

Compile the exit source code by issuing one of the following commands:

32 bit applications

Non-threaded

```
cc -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits/MYFORMAT \
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits/MYFORMAT_r \
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
cc -q64 -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits64/MYFORMAT \
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Threaded

```
xlc_r -q64 -e MQStart -bE:MYFORMAT.exp -bM:SRE -o /var/mqm/exits64/MYFORMAT_r \
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

HP-UX Itanium platform

Compile and link the exit source code by issuing one of the following sets of commands:

32-bit applications

Non-threaded

Compile the exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld +b: -b MYFORMAT.o +ee MQStart -o \
/var/mqm/exits/MYFORMAT -L/usr/lib/hpux32
rm MYFORMAT.o
```

Threaded

Compile the exit source code:

```
c89 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld +b: -b MYFORMAT.o +ee MQStart -o \
/var/mqm/exits/MYFORMAT_r -L/usr/lib/hpux32 \
-lpthread
rm MYFORMAT.o
```

64-bit applications

Non-threaded

Compile the exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld -b MYFORMAT.o +ee MQStart \  
-o /var/mqm/exits64/MYFORMAT \  
-L/usr/lib/hpux64 \  
rm MYFORMAT.o
```

Threaded

Compile the exit source code:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o MYFORMAT.o MYFORMAT.c -IMQ_INSTALLATION_PATH/inc
```

Link the exit object:

```
ld -b MYFORMAT.o +ee MQStart \  
-o /var/mqm/exits64/MYFORMAT_r \  
-L/usr/lib/hpux64 -lpthread \  
rm MYFORMAT.o
```

Linux

Compile the exit source code by issuing one of the following commands:

31 bit applications

Non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/MYFORMAT MYFORMAT.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/MYFORMAT_r MYFORMAT.c \  
-IMQ_INSTALLATION_PATH/inc
```

32 bit applications

Non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/MYFORMAT MYFORMAT.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/MYFORMAT_r MYFORMAT.c \  
-IMQ_INSTALLATION_PATH/inc
```

64 bit applications

Non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/MYFORMAT MYFORMAT.c \  
-IMQ_INSTALLATION_PATH/inc
```

Threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/MYFORMAT_r MYFORMAT.c  
-IMQ_INSTALLATION_PATH/inc
```

Solaris

Compile the exit source code by issuing one of the following commands:

32-bit applications

SPARC platform

```
cc -xarch=v8plus -KPIC -mt -G -o /var/mqm/exits/MYFORMAT \  
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc -R/usr/lib/32 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=386 -KPIC -mt -G -o /var/mqm/exits/MYFORMAT \  
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc -R/usr/lib/32 -lsocket -lnsl -ldl
```

64-bit applications

SPARC platform

```
cc -xarch=v9 -KPIC -mt -G -o /var/mqm/exits64/MYFORMAT \  
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc -R/usr/lib/64 -lsocket -lnsl -ldl
```

x86-64 platform

```
cc -xarch=amd64 -KPIC -mt -G -o /var/mqm/exits64/MYFORMAT \  
MYFORMAT.c -IMQ_INSTALLATION_PATH/inc -R/usr/lib/64 -lsocket -lnsl -ldl
```

Writing a data-conversion exit for WebSphere MQ for Windows

Information about steps to consider when writing data-conversion exit programs for WebSphere MQ for Windows.

Follow these steps:

1. Name your message format. The name must fit in the *Format* field of the MQMD. The *Format* name must not have leading blanks. Trailing blanks are ignored. The object's name must have no more than eight non-blank characters, because the *Format* is only eight characters long.

A .DEF file called amqsvfcn.def is also supplied in the samples directory, `MQ_INSTALLATION_PATH\Tools\C\Samples`. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed. Take a copy of this file and rename it, for example, to MYFORMAT.DEF. Make sure that the name of the DLL being created and the name specified in MYFORMAT.DEF are the same. Overwrite the name FORMAT1 in MYFORMAT.DEF with the new format name.

Remember to use this name each time that you send a message.

2. Create a structure to represent your message. See [Valid syntax](#) for an example.
3. Run this structure through the `crtmqcvx` command to create a code fragment for your data-conversion exit.

The functions generated by the CRTMQCVX command use macros that are written assuming that all structures are packed; amend them if this is not the case.

4. Copy the supplied skeleton source file, `amqsvfc0.c`, renaming it to the name of your message format that you set in step “1” on page 402.

`amqsvfc0.c` is in `MQ_INSTALLATION_PATH\Tools\C\Samples` where `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed. (The default installation directory is `C:\Program Files\IBM\WebSphere MQ`.)

The skeleton includes a sample header file `amqsvmha.h` in the `MQ_INSTALLATION_PATH\Tools\C\include` directory. Make sure that your include path points to this directory to pick up this file.

The `amqsvmha.h` file contains macros that are used by the code generated by the `CRTMQCVX` command. If the structure to be converted contains character data, these macros call `MQXCNVC`.

5. Find the following comment boxes in the source file and insert code as described:

- a. Toward the end of the source file, a comment box starts with:

```
/* Insert the functions produced by the data-conversion exit */
```

Here, insert the code fragment generated in step “3” on page 402.

- b. Near the middle of the source file, a comment box starts with:

```
/* Insert calls to the code fragments to convert the format's */
```

This is followed by a commented-out call to the function `ConverttagSTRUCT`.

Change the name of the function to the name of the function that you added in step “5.a” on page 403. Remove the comment characters to activate the function. If there are several functions, create calls for each of them.

- c. Near the beginning of the source file, a comment box starts with:

```
/* Insert the function prototypes for the functions produced by */
```

Here, insert the function prototype statements for the functions added in step “3” on page 402.

6. Create the following command file:

```
c1 -I MQ_INSTALLATION_PATH\Tools\C\Include -Tp \
MYFORMAT.C

MYFORMAT.DEF
```

where `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.

7. Issue the command file to compile your exit as a DLL file.
8. Place the output in the exit subdirectory below the WebSphere MQ data directory. The default directory for installing your exits on 32 bit systems is `MQ_DATA_PATH\Exits` and for 64 bit systems is `MQ_DATA_PATH\Exits64`

The path used to look for the data-conversion exits is given in the registry. The registry folder is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\IBM\MQSeries\CurrentVersion\Configuration\ClientExitPa
th\
```

and the registry key is: `ExitsDefaultPath`. This path can be set for each queue manager and the exit is only looked for in that path or paths.

Note:

1. If `CRTMQCVX` uses packed structures, all WebSphere MQ applications must be compiled in this way.
2. Data-conversion exit programs must be reentrant.
3. `MQXCNVC` is the *only* MQI call that can be issued from a data-conversion exit.

Exit and switch load files on Windows operating systems

The IBM WebSphere MQ for Windows Version 7.5 queue manager processes are 32-bit. As a result, when using 64-bit applications, some types of exit and XA switch load files also need to have a 32-bit version

available for use by the queue manager. If the 32-bit version of the exit or XA switch load file is required and is not available, then the relevant API call or command fails.

Two attributes are supported in the `qm.ini` file for `ExitPath`. These are `ExitsDefaultPath=MQ_INSTALLATION_PATH\exits` and `ExitsDefaultPath64=MQ_INSTALLATION_PATH\exits64`. `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed. Using these ensures that the appropriate library can be found. If an exit is used in a WebSphere MQ cluster, this also ensures that the appropriate library on a remote system can be found.

The following table lists the different types of Exit and Switch load files and notes whether 32-bit or 64-bit versions, or both, are required, according to whether 32-bit or 64-bit applications are being used:

File types	32-bit applications	64-bit applications
API-crossing exit	32-bit	32-bit and 64-bit
Data conversion exit	32-bit	64-bit
Server Channel exits (all types)	32-bit	32-bit
Client Channel exits (all types)	32-bit	64-bit
Installable service exit	32-bit	32-bit
Service trace module	32-bit	32-bit and 64-bit
Cluster WLM exit	32-bit	32-bit
Pub/Sub routing exit	32-bit	32-bit
Database switch load files	32-bit	32-bit and 64-bit
External Transaction Manager AX libraries	32-bit	64-bit

Referencing connection definitions using a pre-connect exit from a repository

WebSphere MQ MQI clients can be configured to look up a repository to obtain connection definitions using a pre-connect exit library.

Introduction

A client application can connect to a queue manager using client channel definition tables (CCDT). Generally, the CCDT file is located on a central network file server, and have clients referencing it. Since it is difficult to manage and administer various client applications referencing the CCDT file, a flexible approach is to store the client definitions in a global repository like an LDAP directory, a WebSphere Registry and Repository or any other repository. Storing the client connection definitions in a repository makes managing client connection definitions easier, and applications can access the correct and most current client connection definitions.

During the MQCONN/X call execution, the IBM WebSphere MQ MQI client loads an application specified pre-connect exit library, and invokes an exit function to retrieve connection definitions. The retrieved connection definitions are then used to establish connection to a queue manager. The details of exit library and function to invoke are specified in the `mqclient.ini` configuration file.

Syntax

```
void MQ_PRECONNECT_EXIT (pExitParms, pQMGrName, ppConnectOpts, pCompCode, pReason);
```

Parameters

pExitParms

Type: PMQNXP input /output

The **PreConnection** exit parameter structure.

The structure is allocated and maintained by the caller of the exit.

pQMgrName

Type: PMQCHAR input/output

Name of the queue manager.

On input, this parameter is the filter string supplied to the MQCONN API call through the **QMgrName** parameter. This field might be blank, explicit, or contain certain wildcard characters. The field is changed by the exit. The parameter is NULL when the exit is called with MQXR_TERM.

ppConnectOpts

Type: ppConnectOpts input/output

Options that control the action of MQCONN.

This is a pointer to an MQCNO connection options structure that controls the action of the MQCONN API call. The parameter is NULL when the exit is called with MQXR_TERM. The MQI client always provides an MQCNO structure to the exit, even if it was not originally provided by the application. If an application provides an MQCNO structure, the client makes a duplicate to pass it to the exit where it is modified. The client retains the ownership of the MQCNO.

An MQCD referenced through the MQCNO takes precedence over any connection definition provided through the array. The client uses the MQCNO structure to connect to the queue manager and the others are ignored.

pCompCode

Type: PMQLONG input/output

Completion code.

Pointer to an MQLONG that receives the exits completion code. It must be one of the following values:

- MQCC_OK - Successful completion
- MQCC_WARNING - Warning (partial completion)
- MQCC_FAILED - Call failed

pReason

Type: PMQLONG input/output

Reason qualifying pCompCode.

Pointer to an MQLONG that receives the exit reason code. If the completion code is MQCC_OK, the only valid value is:

- MQRC_NONE - (0, x'000') No reason to report.

If the completion code is MQCC_FAILED or MQCC_WARNING, the exit function can set the reason code field to any valid MQRC_* value.

C Invocation

```
void MQ_PRECONNECT_EXIT (&ExitParms, &QMgrName, &pConnectOpts, &CompCode, &Reason);
```

Parameter

```
PMQNX  pExitParms    /*PreConnect exit parameter structure*/
PMQCHAR pQMgrName    /*Name of the queue manager*/
PPMQCNO ppConnectOpts/*Options controlling the action of MQCONN*/
PMQLONG pCompCode    /*Completion code*/
PMQLONG pReason      /*Reason qualifying pCompCode*/
```

PreConnect stanza of the client configuration file

Use the PreConnect stanza to configure the PreConnect exit in the `mqclient.ini` file.

The following attributes can be included in the PreConnect stanza:

Data=<URL>

URL of the repository where connection definitions are stored. For example, when using an LDAP server:

Data = ldap://myLDAPServer.com:389/cn=wmq,ou=ibm,ou=com

Function=<myFunc>

Name of the functional entry point into the library that contains the PreConnect exit code.

The function definition adheres to the PreConnect exit prototype [MQ_PRECONNECT_EXIT](#).

The maximum length of this field is `MQ_EXIT_NAME_LENGTH`.

Module=<amqldapi>

The name of the module containing the API exit code.

If this field contains the full path name of the module, it is used as is.

Sequence=<sequence_number>

The sequence in which this exit is called relative to other exits. An exit with a low sequence number is called before an exit with a higher sequence number. There is no need for the sequence numbering of exits to be continuous; a sequence of 1, 2, 3 has the same result as a sequence of 7, 42, 1096. This attribute is an unsigned numeric value.

Multiple PreConnect stanzas can be defined within the `mqclient.ini` file. The processing order of each exit is determined by the Sequence attribute of the stanza.

Writing and compiling publish exits

You can configure a publish exit at the queue manager to change the contents of a published message before it is received by subscribers. You can also change the message header, or not deliver the message to a subscription.

Publish exits are not supported on z/OS.

You can use the publish exit to inspect and alter messages delivered to subscribers:

- Examine the contents of a message published to each subscriber
- Modify the contents of a message published to each subscriber
- Alter the queue to which a message is put
- Stop the delivery of a message to a subscriber

Writing a publish exit

Use the steps in [“Writing and compiling exits and installable services” on page 357](#), to help you write and compile your exit.

The provider of the publish exit defines what the exit does. The exit, however, must conform to the rules defined in [MQPSXP](#).

WebSphere MQ does not provide an implementation of the `MQ_PUBLISH_EXIT` entry point. It does provide a C language typedef declaration. Use the typedef to declare the parameters to a user-written exit correctly. The following example illustrates how to use the typedef declaration:

```
#include "cmqec.h"

MQ_PUBLISH_EXIT MyPublishExit;

void MQENTRY MyPublishExit( PMQPSXP pExitParms,
                           PMQPBC  pPubContext,
                           PMQSBC  pSubContext )
{
```

```
/* C language statements to perform the function of the exit */  
}
```

The publish exit runs within the queue manager process, as a result of the following operations:

- A Publish operation where a message is delivered to one or more subscribers
- A Subscribe operation where one or more retained messages are delivered
- A Subscription Request operation where one or more retained messages are delivered

If the publish exit is called for a connection, the first time that it is called an *ExitReason* code of MQXR_INIT is set. Before the connection disconnects after using a publish exit, the exit is called with an *ExitReason* code of MQXR_TERM.

If the publish exit is configured, but cannot be loaded when the queue manager is started, publish/subscribe message operations are inhibited for the queue manager. You must fix the problem or restart the queue manager before publish/subscribe messaging is re-enabled.

Each WebSphere MQ connection that requires the publish exit might fail to load or initialize the exit. If the exit fails to load or initialize, publish/subscribe operations that require the publish exit are disabled for that connection. The operations fail with the WebSphere MQ reason code MQRC_PUBLISH_EXIT_ERROR.

The context in which the publish exit is called is the connection by an application to the queue manager. A user data area is maintained by queue manager for each connection that is performing publish operations. The exit can retain information in the user data area for each connection.

A publish exit can use some MQI calls. It can only use those MQI calls that manipulate message properties. The calls are:

- MQBUFMH
- MQCRTMH
- MQDLTMH
- MQDLTMP
- MQMHBUF
- MQINQMP
- MQSETMP

If the publish exit changes the destination queue manager or queue name, no new authority check is carried out.

Compiling a publish exit

The publish exit is a dynamically loaded library; it can be thought of as a channel-exit. For information about compiling exits, see [“Writing and compiling exits and installable services”](#) on page 357.

Sample publish exit

The sample exit program is called `amqspse0.c`. It writes a different message to a log file depending on whether the exit was called for initialize, publish, or terminate operations. It also demonstrates the use of the exit user area field to allocate and free storage appropriately.

Configuring publish exits

You must define certain attributes to configure a publish exit.

On Windows and Linux you can use the WebSphere MQ explorer to define the attributes. The attributes are defined on the queue manager properties page, under Publish/Subscribe.

To configure the publish exit in the `qm.ini` file on UNIX and Linux systems, create a stanza called `PublishSubscribe`. The `PublishSubscribe` stanza has the following attributes:

PublishExitPath=[path] | module_name

Module name and path containing the publish exit code. The maximum length of this field is MQ_EXIT_NAME_LENGTH. The default is no publish exit.

PublishExitFunction=function_name

Name of the function entry point into the module that contains the publish exit code. The maximum length of this field is MQ_EXIT_NAME_LENGTH.

PublishExitData=string

If the queue manager is calling a publish exit, it passes an MQPSXP structure as input. The data specified using the *PublishExitData* attribute is provided in the *ExitData* field of the structure. The string can be up to MQ_EXIT_DATA_LENGTH characters in length. The default is 32 blank characters.

Writing and compiling cluster workload exits

Write a cluster workload exit program to customize the workload management of clusters. You might take the cost of using a channel at different times of day, or message content, into account when routing messages. These are factors that are not considered by the standard workload management algorithm.

In most cases the workload management algorithm is sufficient for your needs. However, so that you can provide your own user-exit program to tailor workload management, WebSphere MQ includes a user exit, the cluster workload exit.

You might have some particular information about your network or messages that you could use to influence workload balancing. You might know which are the high-capacity channels or the cheap network routes, or you might want to route messages depending upon their content. You could decide to write a cluster workload exit program, or use one supplied by a third party.

The cluster workload exit is called when accessing a cluster queue. It is called by MQOPEN, MQPUT1 and MQPUT.

The target queue manager selected at MQOPEN time is fixed if MQOO_BIND_ON_OPEN is specified. In this case the exit is run only once.

If the target queue manager is not fixed at MQOPEN time, the target queue manager is chosen at the time of the MQPUT call. If the target queue manager is not available, or fails while the message is still on the transmission queue, the exit is called again. A new target queue manager is selected. If the message channel fails whilst the message is being transferred, and the message is backed out, a new target queue manager is selected.

On platforms other than z/OS, the queue manager will load the new cluster workload exit the next time the queue manager is started.

If the queue manager definition does not contain a cluster workload exit program name, the cluster workload exit is not called.

Various data are passed to a cluster workload exit in the exit parameter structure, MQWXP:

- The message definition structure, MQMD.
- The message length parameter.
- A copy of the message, or part of the message.

On non-z/OS platforms, if you use CLWLMode=FAST, each operating system process loads its own copy of the exit. Different connections to the queue manager can cause different copies of the exit to be invoked. If the exit is run in the default safe mode, CLWLMode=SAFE, a single copy of the exit runs in its own separate process.

Writing cluster workload exits

For platforms other than z/OS, cluster workload exits must not use MQI calls. In other respects, the rules for writing and compiling cluster workload exit programs are like the rules that apply to channel exit programs. Follow the steps in [“Writing and compiling exits and installable services” on page 357](#), and use the sample program, [“Sample cluster workload exit” on page 409](#) to help write and compile your exit.

For more information about channel exits, see [“Writing channel-exit programs” on page 381](#).

Configuring cluster workload exits

You name cluster workload exits in the queue manager definition by specifying the cluster workload exit attribute on the ALTER QMGR command. For example:

```
ALTER QMGR CLWLEXIT(myexit)
```

Sample cluster workload exit

WebSphere MQ includes a sample cluster workload exit program. You can copy the sample and use it as a basis for your own programs.

On platforms other than z/OS

The sample cluster workload exit program is supplied in C and is called `amqswlm0.c`. It can be found in:

Table 57. Sample cluster workload exit program location (not z/OS)	
Platform	Filepath
AIX, HP-UX, Sun Solaris	<code>MQ_INSTALLATION_PATH/samp</code>
Windows	<code>MQ_INSTALLATION_PATH\Tools\c\Samples</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

This sample exit routes all messages to a particular queue manager, unless that queue manager becomes unavailable. It reacts to the failure of the queue manager by routing messages to another queue manager.

Indicate which queue manager you want messages to be sent to. Supply the name of the cluster-receiver channel in the CLWLDATA attribute on the queue-manager definition. For example:

```
ALTER QMGR CLWLDATA('my-cluster-name.my-queue-manager')
```

To enable the exit, supply its full path and name in the CLWLEXIT attribute:

On UNIX and Linux systems:

```
ALTER QMGR CLWLEXIT('path/amqswlm(cwlFunction)')
```

On Windows:

```
ALTER QMGR CLWLEXIT('path\amqswlm(cwlFunction)')
```

Now, instead of using the supplied workload management algorithm, WebSphere MQ calls this exit to route all messages to your chosen queue manager.

Building an IBM WebSphere MQ application

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

Building your application on AIX

The AIX publications describe how to build executable applications from the programs that you write.

This topic describes the additional tasks, and the changes to the standard tasks, that you must perform when building WebSphere MQ for AIX applications to run under AIX. C, C++, and COBOL are supported. For information about preparing your C++ programs, see [Using C++](#).

The tasks that you must perform to create an executable application using WebSphere MQ for AIX vary with the programming language that your source code is written in. In addition to coding the MQI calls in

your source code, you must add the appropriate language statements to include the WebSphere MQ for AIX include files for the language that you are using. Make yourself familiar with the contents of these files. See [“IBM WebSphere MQ data definition files”](#) on page 77 for a full description.

When you run threaded server or threaded client applications, set the environment variable `AIXTHREAD_SCOPE=S`.

Preparing C programs in AIX

This topic contains information about linking libraries necessary to prepare C programs on AIX.

Precompiled C programs are supplied in the `MQ_INSTALLATION_PATH/samp/bin` directory. Use the ANSI compiler and run the following commands. For further information about programming 64 bit applications, see [Coding standards on 64-bit platforms](#).

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

For 32-bit applications:

```
$ xlc_r -o amqsput_32 amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -lmqm
```

where `amqsput0` is a sample program.

For 64-bit applications:

```
$ xlc_r -q64 -o amqsput_64 amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -lmqm
```

where `amqsput0` is a sample program.

If you are using the VisualAge C/C++ compiler for C++ programs you must include the option `-qnamemangling=v5` to get all the WebSphere MQ symbols resolved when linking the libraries.

If you want to use the programs on a machine that has only the WebSphere MQ MQI client for AIX installed, recompile the programs to link them with the client library (`-lmqic`) instead.

Linking libraries

You need the following libraries:

- Link your programs with the appropriate library provided by WebSphere MQ.

In a non-threaded environment, link to one of the following libraries:

Library file	Program/exit type
<code>libmqm.a</code>	Server for C
<code>libmqic.a</code> & <code>libmqm.a</code>	Client for C

In a threaded environment, link to one of the following libraries:

Library file	Program/exit type
<code>libmqm_r.a</code>	Server for C
<code>libmqic_r.a</code> & <code>libmqm_r.a</code>	Client for C

For example, to build a simple threaded WebSphere MQ application from a single compilation unit run the following commands.

For 32-bit applications:

```
$ xlc_r -o amqsputc_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -lmqm_r
```

where amqspu0 is a sample program.

For 64-bit applications:

```
$ xlc_r -q64 -o amqsputc_64_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -lmqm_r
```

where amqspu0 is a sample program.

If you want to use the programs on a machine that has only the WebSphere MQ MQI client for AIX installed, recompile the programs to link them with the client library (-lmqic) instead.

Note:

1. If you are writing an installable service (see the [Administering](#) for further information), you need to link to the libmqmzf.a library in a non-threaded application and to the libmqmzf_r.a library in a threaded application.
2. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries, Encina, or BEA Tuxedo, you need to link to the libmqmxa.a (or libmqmxa64.a if your transaction manager treats the 'long' type as 64 bit) and libmqz.a libraries in a non-threaded application and to the libmqmxa_r.a (or libmqmxa64_r.a) and libmqz_r.a libraries in a threaded application.
3. You need to link trusted applications to the threaded WebSphere MQ libraries. However, only one thread in a trusted application on WebSphere MQ on UNIX and Linux systems can be connected at a time.
4. You must link WebSphere MQ libraries before any other product libraries.

Preparing COBOL programs in AIX

Use this information when preparing COBOL programs in AIX using IBM COBOL Set and Micro Focus COBOL.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

- 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

- 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

In the following examples set the **COBCPY** environment variable to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

for 64 bit applications.

You need to link your program with one of the following library files:

Library file	Program/exit type
libmqmcb.a	Server for COBOL (unthreaded application)
libmqmcb_r.a	Server for COBOL (threaded application)
libmqicb.a	Client for COBOL (unthreaded application)
libmqicb_r.a	Client for COBOL (threaded application)

You can use the IBM COBOL Set compiler or Micro Focus COBOL compiler depending on the program:

- Programs beginning amqm are suitable for the Micro Focus COBOL compiler, and
- Programs beginning amq0 are suitable for either compiler.

Preparing COBOL programs using IBM COBOL Set for AIX

Sample COBOL programs are supplied with IBM WebSphere MQ. To compile such a program, enter the appropriate command from the following list:

32 bit non-threaded server application

```
$ cob2 -o amq0put0 amq0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqmc -qLIB \
-I<COBCPY>
```

32 bit non-threaded client application

```
$ cob2 -o amq0put0 amq0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb -qLIB \
-I<COBCPY>
```

32 bit threaded server application

```
$ cob2_r -o amq0put0 amq0put0.cbl -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqmcb_r -qLIB -I<COBCPY>
```

32 bit threaded client application

```
$ cob2_r -o amq0put0 amq0put0.cbl -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqicb_r -qLIB -I<COBCPY>
```

64 bit non-threaded server application

```
$ cob2 -o amq0put0 amq0put0.cbl -q64 -L MQ_INSTALLATION_PATH/lib -lmqmc \
-qLIB -I<COBCPY>
```

64 bit non-threaded client application

```
$ cob2 -o amq0put0 amq0put0.cbl -q64 -L MQ_INSTALLATION_PATH/lib -lmqicb \
-qLIB -I<COBCPY>
```

64 bit threaded server application

```
$ cob2_r -o amq0put0 amq0put0.cbl -q64 -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqmcb_r -qLIB -I<COBCPY>
```

64 bit threaded client application

```
$ cob2_r -o amq0put0 amq0put0.cbl -q64 -qTHREAD -L MQ_INSTALLATION_PATH/lib \
-lmqicb_r -qLIB -I<COBCPY>
```

Preparing COBOL programs using Micro Focus COBOL

Set environment variables before compiling your program as follows:

```
export COBCPY=<COBCPY>
export LIBPATH=MQ_INSTALLATION_PATH/lib:$LIBPATH
```

To compile a 32 bit COBOL program using Micro Focus COBOL, enter:

- Server for COBOL

```
$ cob32 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcbl
```

- Client for COBOL

```
$ cob32 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqicbl
```

- Threaded Server for COBOL

```
$ cob32 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcbl_r
```

- Threaded Client for COBOL

```
$ cob32 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib -lmqicbl_r
```

To compile a 64 bit COBOL program using Micro Focus COBOL, enter:

- Server for COBOL

```
$ cob64 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcbl
```

- Client for COBOL

```
$ cob64 -xvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicbl
```

- Threaded Server for COBOL

```
$ cob64 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcbl_r
```

- Threaded Client for COBOL

```
$ cob64 -xtvP amqminqx.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicbl_r
```

where amqminqx is a sample program

See the Micro Focus COBOL documentation for a description of the environment variables that you need to set up.

Preparing CICS application programs in AIX

Use this information when preparing CICS programs in AIX.

XA switch modules are provided to enable you to link CICS with IBM WebSphere MQ:

Table 58. Essential code for CICS application programs on AIX: XA initialization routine		
Description	C (source)	C (exec) - add to your XAD.Stanza
XA initialization routine	amqzscix.c	amqzsc - CICS for AIX

Use the prebuilt version of the IBM WebSphere MQ switch load file *amqzsc*, which is provided with the product.

Always link your C transactions with the threadsafe IBM WebSphere MQ library *libmqm_r.a.*, and your COBOL transactions with the COBOL library *libmqmcb_r.a.*

You can find more information about supporting CICS transactions in the [Administering](#).

TXSeries CICS support

IBM WebSphere MQ on AIX supports TXSeries CICS using the XA interface. Ensure that CICS applications are linked to the threaded version of the IBM WebSphere MQ libraries.

You can run CICS programs using IBM COBOL Set for AIX or Micro Focus COBOL. The following sections describe the difference between running CICS programs on IBM COBOL Set for AIX and Micro Focus COBOL.

Write WebSphere MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of MQRC_HOBJ_ERROR.

Preparing CICS COBOL programs using IBM COBOL Set for AIX

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

To use IBM COBOL, follow these steps:

1. Export the following environment variable:

```
export LDFlags="-qLIB -bI:/usr/lpp/cics/lib/cicsprIBMCOB.exp \  
-IMQ_INSTALLATION_PATH/inc -I/usr/lpp/cics/include \  
-e _iwz_cobol_main \  

```

where LIB is a compiler directive.

2. Translate, compile, and link the program by typing:

```
cicstcl -l IBMCOB <yourprog>.cpp
```

Preparing CICS COBOL programs using Micro Focus COBOL

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the IBM WebSphere MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -LMQ_INSTALLATION_PATH/lib \
MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe_r
```

Note: With `cicsmkcobol`, IBM WebSphere MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, you are recommended to move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving the functions, do not include the IBM WebSphere MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call IBM WebSphere MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
- New version or release of CICS for AIX
- New version or release of any supported database product (for COBOL transactions only)
- New version or release of IBM WebSphere MQ

2. Export the following environment variable:

```
COBCPY=MQ_INSTALLATION_PATH/inc export COBCPY
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e <yourprog>.ccp
```

Preparing CICS C programs

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

Build CICS C programs using the standard CICS facilities:

1. Export **one** of the following environment variables:
 - `LDFLAGS = "-L/MQ_INSTALLATION_PATHlib -lmqm_r" export LDFLAGS`
 - `USERLIB = "-LMQ_INSTALLATION_PATHlib -lmqm_r" export USERLIB`
2. Translate, compile, and link the program by typing:

```
cicstcl -l C amqscic0.ccs
```

CICS C sample transaction

Sample C source for a AIX IBM WebSphere MQ transaction is provided by `AMQSCIC0.CCS`. The transaction reads messages from the transmission queue `SYSTEM.SAMPLE.CICS.WORKQUEUE` on the default queue manager and places them onto the local queue with a queue name that is contained in the transmission header of the message. Any failures are sent to the queue `SYSTEM.SAMPLE.CICS.DLQ`. Use the sample MQSC script `AMQSCIC0.TST` to create these queues and sample input queues.

Building your application on HP Integrity NonStop Server

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when you build IBM WebSphere MQ client for HP Integrity NonStop Server applications to run under HP Integrity NonStop Server.

C, COBOL, and pTAL are supported.

OSS and Guardian headers and public libraries

Provides lists of OSS and Guardian headers and public libraries. Listed are OSS headers, OSS public executable and public import libraries, Guardian headers, and Guardian public executable and public import libraries.

[“OSS headers” on page 416](#)

[“OSS public executable and public import libraries” on page 417](#)

[“Guardian headers” on page 417](#)

[“Guardian public executable and public import libraries” on page 418](#)

OSS headers

Table 59. OSS headers		
Object	Location	Description
cmqbc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqfc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqec.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqpzc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqxc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqzc.h	<mqinstall>/inc	IBM WebSphere MQ C-language header (OSS)
cmqcbol.cpy	<mqinstall>/inc	IBM WebSphere MQ COBOL copybook (OSS)
cmqbt.tal	<mqinstall>/inc	IBM WebSphere MQ pTAL header (OSS)
cmqcft.tal	<mqinstall>/inc	IBM WebSphere MQ pTAL header (OSS)
cmqpst.tal	<mqinstall>/inc	IBM WebSphere MQ pTAL header (OSS)
cmqt.tal	<mqinstall>/inc	IBM WebSphere MQ pTAL header (OSS)
cmqxt.tal	<mqinstall>/inc	IBM WebSphere MQ pTAL header (OSS)

OSS public executable and public import libraries

Table 60. OSS public executable and public import libraries		
Object	Location	Description
libmqic.so	<mqinstall>/bin	IBM WebSphere MQ public executable library (OSS unthreaded)
libmqic_r.so	<mqinstall>/bin	IBM WebSphere MQ public executable library (OSS multi-threaded)
libmqic.so	<mqinstall>/lib	IBM WebSphere MQ public import library (OSS unthreaded)
libmqic_r.so	<mqinstall>/lib	IBM WebSphere MQ public import library (OSS multi-threaded)
mqicb	<mqinstall>/lib	IBM WebSphere MQ public import library for COBOL (OSS)

Guardian headers

Table 61. Guardian headers		
Object	Location	Description
cmqbcch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqfch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqech	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqpsch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqxch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqzch	<mqinstall>/inc/G	IBM WebSphere MQ C-language header (Guardian)
cmqcbol	<mqinstall>/inc/G	IBM WebSphere MQ COBOL copybook (Guardian)
cmqbt	<mqinstall>/inc/G	IBM WebSphere MQ pTAL header (Guardian)
cmqcft	<mqinstall>/inc/G	IBM WebSphere MQ pTAL header (Guardian)
cmqpst	<mqinstall>/inc/G	IBM WebSphere MQ pTAL header (Guardian)
cmqt	<mqinstall>/inc/G	IBM WebSphere MQ pTAL header (Guardian)

Table 61. Guardian headers (continued)		
Object	Location	Description
cmqxt	<mqinstall>/inc/G	IBM WebSphere MQ pTAL header (Guardian)

Guardian public executable and public import libraries

Table 62. Guardian public executable and public import libraries		
Object	Location	Description
mqic	<mqinstall>/bin/G	IBM WebSphere MQ public executable library (Guardian)
mqicb	<mqinstall>/lib/G	IBM WebSphere MQ public import library for COBOL (Guardian)

Preparing C programs in HP Integrity NonStop Server

This topic contains information to consider when you are preparing C programs in HP Integrity NonStop Server together with examples of the commands you use when you are building applications when you are using the OSS C compiler and when you are using the Guardian C compiler.

Precompiled C programs are supplied in the MQ_INSTALLATION_PATH/opt/mqm/samp/bin directory. To build a sample from source code, use the c89 compiler.

You must link your programs with the appropriate library provided by IBM WebSphere MQ. The following table lists the libraries that you must link to when you are preparing C programs on HP Integrity NonStop Server.

Table 63. . HP Integrity NonStop Server link libraries	
Library	Description
libmqic.so	OSS unthreaded
libmqic_r.so	OSS multi-threaded
mqic	Guardian

Multi-threaded native IBM WebSphere MQ applications must use the Posix User Threads (PUT) feature. There is no support for Standard Posix Threads (SPT) in this product.

Building applications using the OSS C compiler

This section contains examples of the commands that are used to build programs that are targeted for either OSS or Guardian when you are using the OSS compiler.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

The following example builds an unthreaded C client OSS application:

```
c89 -Wstype=oss -o amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic
```

The following example builds a multi-threaded C client OSS application:

```
c89 -Wstype=oss -D_PUT_MODEL_ -o amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic_r -lput
```

The following example builds a Guardian C client application:

```
c89 -wsystype=guardian -o /G/vol/subvol/amqsputc amqsput0.c -IMQ_INSTALLATION_PATH/opt/mqm/inc
-LMQ_INSTALLATION_PATH/opt/mqm/lib/G -lmqic
```

Building applications using the Guardian C compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian compiler.

MQ_INSTALLATION_PATH represents the Guardian volume and subvolume in which IBM WebSphere MQ is installed.

The following example builds a Guardian C client application:

```
CCOMP /in AMQSPUT0/ AMQSPUTC;&
runnable,systype guardian,nolist,&
ssv0 "$system.system",&
ssv1 "MQINSTALLATION_SUBVOL",&
ld(-LMQINSTALLATION_SUBVOL -lmqic)
```

Preparing COBOL programs

This topic contains information to consider when you are preparing C programs for the IBM WebSphere MQ client for HP Integrity NonStop Server. It contains examples of the commands you use when you are building applications when you are using the OSS ECOBOL compiler and when you are using the Guardian ECOBOL compiler.

To build a COBOL sample from source code, use the ECOBOL compiler.

The following table lists the libraries that are needed when you are preparing COBOL programs on HP Integrity NonStop Server. You must link your programs with the appropriate library provided by IBM WebSphere MQ.

Table 64. . HP Integrity NonStop Server link libraries	
Library	Description
libmqic.so	OSS unthreaded
mqic	Guardian

When you run a COBOL application that connects to a queue manager, you must first set the *SAVE-ENVIRONMENT* variable to ON. To set the *SAVE-ENVIRONMENT* variable to ON:

- For OSS, enter the following command:

```
export SAVE-ENVIRONMENT=ON
```

- For Guardian, enter the following command:

```
param SAVE-ENVIRONMENT ON
```

If you do not set the *SAVE-ENVIRONMENT* variable to ON, when the application attempts to connect to a queue manager, it fails with reason code 2058 (080A) (RC2058): [MQRC_Q_MGR_NAME_ERROR](#).

Building applications using the OSS ECOBOL compiler

This section contains examples of the commands that are used to build programs that are targeted for either OSS or Guardian when you are using the OSS ECOBOL compiler.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

The following example builds a COBOL client OSS application:

```
ecobol -wsystype=oss
-wcobol="ansi;port"
-wcobol="consult MQ_INSTALLATION_PATH/opt/mqm/lib/mqicb"
-wcopylib=MQ_INSTALLATION_PATH/opt/mqm/inc/cmqcobol.cpy
-LMQ_INSTALLATION_PATH/opt/mqm/lib -lmqic
-o amq0put0
MQ_INSTALLATION_PATH/opt/mqm/samp/amq0put0.cbl
```

The following example builds a COBOL client Guardian application:

```
ecobol -wsystype=guardian
-wcobol="ansi;port;save_all"
-wcobol="consult MQ_INSTALLATION_PATH/opt/mqm/lib/mqicb"
-wcopylib=MQ_INSTALLATION_PATH/opt/mqm/inc/cmqcobol.cpy
-LMQ_INSTALLATION_PATH/opt/mqm/lib/G -lmqic
-o amq0put0
MQ_INSTALLATION_PATH/opt/mqm/samp/amq0put0.cbl
```

Building applications using the Guardian ECOBOL compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian ECOBOL compiler.

MQ_INSTALLATION_SUBVOL represents the Guardian volume and subvolume in which IBM WebSphere MQ is installed.

The following example builds a COBOL client Guardian application:

```
ECOBOL /in MQSPUTL/ MQSPUT, MQINSTALLATION_SUBVOL.cmqcobol;
call-shared;ansi;port;save_all;nolist;runnable;
consult MQINSTALLATION_SUBVOL.mqicb;
eld(-LMQINSTALLATION_SUBVOL -lmqic)
```

Preparing pTAL programs

Learn to build pTAL programs for the IBM WebSphere MQ client on the HP Integrity NonStop Server platform.

To build a pTAL sample from source code, use the EPTAL compiler.

Note:

- pTAL IBM WebSphere MQ applications must use a main routine that is written in either the C or COBOL languages.
- pTAL applications can be built only in Guardian.

The following table lists the library that is needed when you are preparing pTAL programs on HP Integrity NonStop Server. You must link your programs with the appropriate library provided by IBM WebSphere MQ.

Table 65. . HP Integrity NonStop Server link library	
Library	Description
mqic	Guardian

Building applications using the Guardian EPTAL compiler

This section contains examples of the commands that are used to build programs that are targeted for Guardian when you are using the Guardian EPTAL compiler.

MQINSTALLATION_SUBVOL represents the Guardian volume and subvolume in which IBM WebSphere MQ is installed.

pTAL IBM WebSphere MQ applications must use a main routine that is written in either the C or COBOL languages.

The following example builds a pTAL client Guardian application:

```
ASSIGN SSV0, $SYSTEM.SYSTEM
ASSIGN SSV1, MQINSTALLATION_SUBVOL

EPTAL /in MQINSTALLATION_SUBVOL.MQSPUTT/ MQSPUTO;nolist

CCOMP /in MQINSTALLATION_SUBVOL.MQSPTMC/ MQSPUT;
runnable,systype guardian,extensions,nolist,
ssv0 "$system.system",
ssv1 "MQINSTALLATION_SUBVOL",
eid(MQSPUTO -LMQINSTALLATION_SUBVOL -lmqic)
```

Building your application on HP-UX

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building WebSphere MQ for HP-UX applications to run under HP-UX.

C, C++, and COBOL are supported. For information about preparing your C++ programs, see [Using C++](#).

The tasks that you must perform to create an executable application using WebSphere MQ for HP-UX vary with the programming language that your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the WebSphere MQ for HP-UX include files for the language that you are using. Make yourself familiar with the contents of these files. See “IBM WebSphere MQ data definition files” on page 77 for a full description.

Throughout this topic, we use the backslash (\) character to split long commands over more than one line. Do not enter this character; enter each command as a single line.

Preparing C programs in HP-UX

This topic contains information to consider when preparing C programs in HP-UX; with examples for the IA64 (IPF) platform.

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Work in your normal environment. Precompiled C programs are supplied in the *MQ_INSTALLATION_PATH*/samp/bin directory.

For further information about programming 64 bit applications see [Coding standards on 64-bit platforms](#).

To use SSL, WebSphere MQ MQI clients on HP-UX must be built using POSIX threads.

Some examples to consider are:

- “IA64 (IPF) platform” on page 421
- “Linking libraries” on page 423

IA64 (IPF) platform

Build examples of amqsput0, cliexit, and srvexit on IA64(IPF) platform.

The following example builds the sample program amqsput0 as a client application in a non-threaded 32 bit environment:

```
c89 -Wl,+b,: +e -D_HPUX_SOURCE -o amqsputc_32 amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqic
```

The following example builds the sample program amqsput0 as a client application in a threaded 32 bit environment:

```
c89 -mt -Wl,+b,: +e -D_HPUX_SOURCE -o amqsputc_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqic_r -lpthread
```

The following example builds the sample program amqspu0 as a client application in a non-threaded 64 bit environment:

```
c89 +DD64 +e -D_HPUX_SOURCE -o amqspu0c_64 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic
```

The following example builds the sample program amqspu0 as a client application in a threaded 64 bit environment:

```
c89 -mt +DD64 +e -D_HPUX_SOURCE -o amqspu0c_64_r amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic_r -lpthread
```

The following example builds the sample program amqspu0 as a server application in a non-threaded 32 bit environment:

```
c89 -Wl,+b,: +e -D_HPUX_SOURCE -o amqspu0_32 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqm
```

The following example builds the sample program amqspu0 as a server application in a threaded 32 bit environment:

```
c89 -mt -Wl,+b,: +e -D_HPUX_SOURCE -o amqspu0_32_r amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -lmqm_r -lpthread
```

The following example builds the sample program amqspu0 as a server application in a non-threaded 64 bit environment:

```
c89 +DD64 +e -D_HPUX_SOURCE -o amqspu0_64 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm
```

The following example builds the sample program amqspu0 as a server application in a threaded 64 bit environment:

```
c89 -mt +DD64 +e -D_HPUX_SOURCE -o amqspu0_64_r amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm_r -lpthread
```

The following example builds a client exit cliexit in a non-threaded 32 bit environment:

```
c89 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -IMQ_INSTALLATION_PATH/inc  
ld +b: -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_32 -LMQ_INSTALLATION_PATH/lib \  
-L/usr/lib/hpux32 -lmqic
```

The following example builds a client exit cliexit in a threaded 32 bit environment:

```
c89 -mt +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -IMQ_INSTALLATION_PATH/inc  
ld +b: -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_32_r -LMQ_INSTALLATION_PATH/lib \  
-L/usr/lib/hpux32 -lmqic_r -lpthread
```

The following example builds a client exit cliexit in a non-threaded 64 bit environment:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -IMQ_INSTALLATION_PATH/inc  
ld -b cliexit.o +ee MQStart -o /var/mqm/exits64/cliexit_64 \  
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic
```

The following example builds a client exit cliexit in a threaded 64 bit environment:

```
c89 -mt +DD64 +e +z -c -D_HPUX_SOURCE -o cliexit.o cliexit.c -IMQ_INSTALLATION_PATH/inc
```

```
ld -b cliexit.o +ee MQStart -o /var/mqm/exits/cliexit_64_r \
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqic_r -lpthread
```

The following example builds a server exit srvexit in a non-threaded 32 bit environment:

```
c89 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -IMQ_INSTALLATION_PATH/inc
ld +b: -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_32 -LMQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqm
```

The following example builds a server exit srvexit in a threaded 32 bit environment:

```
c89 -mt +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -IMQ_INSTALLATION_PATH/inc
ld +b: -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_32_r -LMQ_INSTALLATION_PATH/lib \
-L/usr/lib/hpux32 -lmqm_r -lpthread
```

The following example builds a server exit srvexit in a non-threaded 64 bit environment:

```
c89 +DD64 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c
-IMQ_INSTALLATION_PATHMQ_INSTALLATION_PATH/inc
ld -b srvexit.o +ee MQStart -o /var/mqm/exits64/srvexit_64 \
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm
```

The following example builds a server exit srvexit in a threaded 64 bit environment:

```
c89 -mt +DD64 +e +z -c -D_HPUX_SOURCE -o srvexit.o srvexit.c -IMQ_INSTALLATION_PATH/inc
ld -b srvexit.o +ee MQStart -o /var/mqm/exits/srvexit_64_r \
-LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -lmqm_r -lpthread
```

Linking libraries

You need to link your programs with the appropriate library provided by WebSphere MQ.

The following table shows which library to use in different environments

Hardware platform	Threaded or non-threaded environment	Program/exit type	Library file
IA64 (IPF)	Threaded	Server & Client for C	libmqm_r.so
IA64 (IPF)	Threaded	Client for C	libmqic_r.so
IA64 (IPF)	Non-threaded	Server & Client for C	libmqm.so
IA64 (IPF)	Non-threaded	Client for C	libmqic.so

Note:

1. If you are writing an installable service (see the [Administering](#) for further information), you need to link to the libmqmzf.s1 library.
2. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the libmqmxa.s1 (or libmqmxa64.s1 if your transaction manager treats the 'long' type as 64 bit) and libmqz.s1 libraries in a non-threaded application and to the libmqmxa_r.s1 (or libmqmxa64_r.s1) and libmqz_r.s1 libraries in a threaded application.
3. You must link WebSphere MQ libraries before any other product libraries.

Preparing COBOL programs in HP-UX

Learn about preparing COBOL programs in HP-UX, using Micro Focus Server Express with WebSphere MQ on the IA64 (IPF) platform, and running programs in the WebSphere MQ MQI client environment.

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Notes to users

1. 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

2. 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

3. In the following examples set COBCPY to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

for 64 bit applications.

Compile the programs using the Micro Focus compiler. The copy files that declare the structures are in `MQ_INSTALLATION_PATH/inc`:

```
$ export LIB=MQ_INSTALLATION_PATH/lib:$LIB
$ export COBCPY="<COBCPY>"
```

Compiling 32 bit programs:

```
$ cob32 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqmbc Server for COBOL
$ cob32 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb Client for COBOL
$ cob32 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqmbc_r Threaded Server for COBOL
$ cob32 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r Threaded Client for COBOL
```

Compiling 64 bit programs:

```
$ cob64 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmbc Server for COBOL
$ cob64 -xv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb Client for COBOL
$ cob64 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmbc_r Threaded Server for COBOL
$ cob64 -xtv amqsput.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r Threaded Client for COBOL
```

where `amqsput` is a sample program

Ensure that you have specified adequate runtime stack sizes; 16 KB is the recommended minimum.

You need to link your programs with the appropriate library provided by WebSphere MQ. The following table shows which library to use in different environments

Hardware platform	Program/exit type	Library file
IA64 (IPF)	Server for COBOL	libmqmbc.so
IA64 (IPF)	Client for COBOL	libmqicb.so
IA64 (IPF)	Threaded applications	libmqmbc_r.so

Using Micro Focus Server Express with WebSphere MQ on the IA64 (IPF) platform

See “Address Space models supported by WebSphere MQ for HP-UX on IA64 (IPF)” on page 426 for details on using Micro Focus Server Express in conjunction with WebSphere MQ on the HP/IPF platform.

Programs to run in the WebSphere MQ MQI client environment

If you are using LU 6.2 to connect your MQI client to a server, link your application to `libsna.a`, part of the SNAplusAPI product. Use the `-lv3` and `-lstx` options on your compile and link command.

- The `-lv3` option gives your program access to the AT&T signaling library (the SNAplusAPI uses AT&T signals)
- The `-lstx` option links your program to the streams component

Preparing CICS programs in HP-UX

Learn to build CICS transaction programs in HP-UX.

To build the sample CICS transaction, `amqscic0.ccs`, run the following command:

```
$ export USERLIB="-lmqm_r"
$ cicstcl -l C amqscic0.ccs
```

An XA switch module is provided to enable you to link CICS with WebSphere MQ:

Table 66. Essential code for CICS applications (HP-UX)		
Description	C (source)	C (exec)
XA initialization routine	<code>amqzscix.c</code>	<code>amqzsc</code>

You can find more information about supporting CICS transactions in the [Administering](#).

TXSeries CICS support

WebSphere MQ on HP-UX supports TXSeries CICS using the XA interface. Ensure that CICS applications are linked to the threaded version of the MQ libraries.

Write WebSphere MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of `MQRC_HOBJ_ERROR`.

CICS C sample transaction

Sample C source for a CICS WebSphere MQ transaction is provided by `AMQSCIC0.CCS`. The transaction reads messages from the transmission queue `SYSTEM.SAMPLE.CICS.WORKQUEUE` on the default queue manager and places them onto the local queue with the queue name that is contained in the transmission header of the message. Any failures are sent to the queue `SYSTEM.SAMPLE.CICS.DLQ`. Use the sample MQSC script `AMQSCIC0.TST` to create these queues and sample input queues.

Preparing CICS COBOL programs using Micro Focus COBOL

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the WebSphere MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -LMQ_INSTALLATION_PATH/lib \
MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe_r
```

Note: With `cicsmkcobol`, WebSphere MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, you are recommended to move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving these functions do not include the WebSphere MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call WebSphere MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
- New version or release of CICS for HP-UX
- New version or release of any supported database product (for COBOL transactions only)
- New version or release of WebSphere MQ

2. Export the following environment variable:

```
COBCPY=MQ_INSTALLATION_PATH/inc export COBCPY
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e <yourprog>.ccp
```

Address Space models supported by WebSphere MQ for HP-UX on IA64 (IPF)

The HP-UX provides several address space models that can be exploited by WebSphere MQ applications.

HP-UX supports two Address Space models:

- MGAS - Mostly Global Address space (this is the default and is used by WebSphere MQ)
- MPAS - Mostly Private Address space

Applications which connect to WebSphere MQ can use either the MGAS or MPAS address space models. Applications built using the MPAS model that connect to WebSphere MQ using shared memory might incur a minor performance cost due to the inefficiency in mapping the shared memory pages used by WebSphere MQ into the virtual address space of the MPAS program.

COBOL applications built using Micro Focus Server Express use the MPAS model by default.

You can use the **chatx** program to check and change the addressing model used by a program.

If you encounter problems connecting to WebSphere MQ from 32-bit MPAS programs, consider using the MGAS addressing model, or building your application as a 64-bit MPAS application rather than a 32-bit MPAS application.

More details on the MGAS and MPAS address space models can be found in the HP-UX documentation.

Building your application on Linux

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building WebSphere MQ for Linux applications to run.

C and C++ are supported. For information about preparing your C++ programs, see [Using C++](#).

Preparing C programs in Linux

Precompiled C programs are supplied in the `MQ_INSTALLATION_PATH/samp/bin` directory. To build a sample from source code, use the `gcc` compiler.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Work in your normal environment. For further information about programming 64 bit applications, see [Coding standards on 64-bit platforms](#).

Linking libraries

The following tables lists the libraries that are needed when preparing C programs on Linux.

- You need to link your programs with the appropriate library provided by WebSphere MQ.

In a non-threaded environment, link to one of the following libraries:

Library file	Program/exit type
libmqm.so	Server for C
libmqic.so & libmqm.so	Client for C

In a threaded environment, link to one of the following libraries:

Library file	Program/exit type
libmqm_r.so	Server for C
libmqic_r.so & libmqm_r.so	Client for C

Note:

- If you are writing an installable service (see the [Administering](#) for further information), you need to link to the `libmqmzf.so` library.
- If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the `libmqmxa.so` (or `libmqmxa64.so` if your transaction manager treats the 'long' type as 64 bit) and `libmqz.so` libraries in a non-threaded application and to the `libmqmxa_r.so` (or `libmqmxa64_r.so`) and `libmqz_r.so` libraries in a threaded application.
- You must link WebSphere MQ libraries before any other product libraries.

Building 31-bit applications

This topic contains examples of the commands used to build 31-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

C client application, 31-bit, non-threaded

```
gcc -m31 -o famqspu32 amqspu0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic
```

C client application, 31-bit, threaded

```
gcc -m31 -o amqsputc_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server application, 31-bit, non-threaded

```
gcc -m31 -o amqsput_32 amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm
```

C server application, 31-bit, threaded

```
gcc -m31 -o amqsput_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

C++ client application, 31-bit, non-threaded

```
g++ -m31 -fsigned-char -o imqsputc_32 imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl  
-limqb23gl -lmqic
```

C++ client application, 31-bit, threaded

```
g++ -m31 -fsigned-char -o imqsputc_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl_r  
-limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 31-bit, non-threaded

```
g++ -m31 -fsigned-char -o imqsput_32 imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl  
-limqb23gl -lmqm
```

C++ server application, 31-bit, threaded

```
g++ -m31 -fsigned-char -o imqsput_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl_r  
-limqb23gl_r -lmqm_r -lpthread
```

C client exit, 31-bit, non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib  
-lmqic
```

C client exit, 31-bit, threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/cliexit_32_r cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib  
-lmqic_r -lpthread
```

C server exit, 31-bit, non-threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/srvexit_32 srvexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib
```

```
-Wl,-rpath=/usr/lib  
-lmqm
```

C server exit, 31-bit, threaded

```
gcc -m31 -shared -fPIC -o /var/mqm/exits/srvexit_32_r srvexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib  
-Wl,-rpath=/usr/lib  
-lmqm_r -lpthread
```

Building 32-bit applications

This topic contains examples of the commands used to build 32-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

C client application, 32-bit, non-threaded

```
gcc -m32 -o amqsputc_32 amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic
```

C client application, 32-bit, threaded

```
gcc -m32 -o amqsputc_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqic_r -lpthread
```

C server application, 32-bit, non-threaded

```
gcc -m32 -o amqsput_32 amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm
```

C server application, 32-bit, threaded

```
gcc -m32 -o amqsput_32_r amqsput0.c -IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib  
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm_r -lpthread
```

C++ client application, 32-bit, non-threaded

```
g++ -m32 -fsigned-char -o imqsputc_32 imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl -limqb23gl -lmqic
```

C++ client application, 32-bit, threaded

```
g++ -m32 -fsigned-char -o imqsputc_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 32-bit, non-threaded

```
g++ -m32 -fsigned-char -o imqsput_32 imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl -limqb23gl -lmqm
```

C++ server application, 32-bit, threaded

```
g++ -m32 -fsigned-char -o imqsput_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl_r -limqb23gl_r -lmqm_r -lpthread
```

C client exit, 32-bit, non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/cliexit_32 cliexit.c
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib
-Wl,-rpath=/usr/lib
-lmqic
```

C client exit, 32-bit, threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/cliexit_32_r cliexit.c
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib
-Wl,-rpath=/usr/lib
-lmqic_r -lpthread
```

C server exit, 32-bit, non-threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/srvexit_32 srvexit.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -lmqm
```

C server exit, 32-bit, threaded

```
gcc -m32 -shared -fPIC -o /var/mqm/exits/srvexit_32_r srvexit.c
IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib
-Wl,-rpath=/usr/lib
lmqm_r -lpthread
```

Building 64-bit applications

This topic contains examples of the commands used to build 64-bit programs in various environments.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

C client application, 64-bit, non-threaded

```
gcc -m64 -o amqsputc_64 amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqic
```

C client application, 64-bit, threaded

```
gcc -m64 -o amqsputc_64_r amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqic_r
-lpthread
```

C server application, 64-bit, non-threaded

```
gcc -m64 -o amqsput_64 amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm
```

C server application, 64-bit, threaded

```
gcc -m64 -o amqsput_64_r amqsput0.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -lmqm_r
-lpthread
```

C++ client application, 64-bit, non-threaded

```
g++ -m64 -fsigned-char -o imqsputc_64 imqsput.cpp
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
```

```
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl -limqb23gl -lmqic
```

C++ client application, 64-bit, threaded

```
g++ -m64 -fsigned-char -o imqspu64_r imqspu64.cpp
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

C++ server application, 64-bit, non-threaded

```
g++ -m64 -fsigned-char -o imqspu64 imqspu64.cpp
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -limq23gl -limqb23gl -lmqm
```

C++ server application, 64-bit, threaded

```
g++ -m64 -fsigned-char -o imqspu64_r imqspu64.cpp
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -limq23gl_r -limqb23gl_r -lmqm_r -lpthread
```

C client exit, 64-bit, non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/cliexit_64 cliexit.c
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqic
```

C client exit, 64-bit, threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/cliexit_64_r cliexit.c
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqic_r -lpthread
```

C server exit, 64-bit, non-threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/srvexit_64 srvexit.c
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqm
```

C server exit, 64-bit, threaded

```
gcc -m64 -shared -fPIC -o /var/mqm/exits64/srvexit_64_r srvexit.c
-IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64
-Wl,-rpath=/usr/lib64 -lmqm_r -lpthread
```

Preparing COBOL programs in Linux

Learn about preparing COBOL programs in Linux and preparing COBOL programs using Micro Focus COBOL.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

1. 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

2. On 64 bit platforms, 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

3. In the following examples set COBCPY to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

for 64 bit applications.

You need to link your program with one of the following:

Library file	Program/exit type
libmqmcb.so	Server for COBOL
libmqicb.so	Client for COBOL
libmqmcb_r.so	Server for COBOL (threaded application)
libmqicb_r.so	Client for COBOL (threaded application)

Preparing COBOL programs using Micro Focus COBOL

Set environment variables before compiling your program as follows:

```
export COBCPY=<COBCPY>
export LIB=MQ_INSTALLATION_PATH/lib:$LIB
```

To compile a 32 bit COBOL program, where supported, using Micro Focus COBOL, enter:

```
$ cob32 -xvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcb Server for COBOL
$ cob32 -xvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb Client for COBOL
$ cob32 -xtvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib -lmqmcb_r Threaded Server for COBOL
$ cob32 -xtvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r Threaded Client for COBOL
```

To compile a 64 bit COBOL program using Micro Focus COBOL, enter:

```
$ cob64 -xvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcb Server for COBOL
$ cob64 -xvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb Client for COBOL
$ cob64 -xtvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmcb_r Threaded Server for COBOL
$ cob64 -xtvP amqspout.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r Threaded Client for COBOL
```

where amqspout is a sample program

See the Micro Focus COBOL documentation for a description of the environment variables that you need.

Building your application on Solaris

This information describes the additional tasks, and the changes to the standard tasks, that you must perform when building WebSphere MQ for Solaris applications to run under Solaris.

COBOL, C, and C++ programming languages are supported. For information about preparing your C++ programs, see [Using C++](#).

In addition to coding the MQI calls in your source code, you must add the appropriate include files. Make yourself familiar with the contents of these files. See [“IBM WebSphere MQ data definition files”](#) on page 77 for a full description.

Throughout this topic, the backslash (\) character is used to split long commands over more than one line. Do not enter this character, enter each command as a single line.

Preparing C programs in Solaris

Precompiled C programs are supplied in the `MQ_INSTALLATION_PATH/samp/bin` directory.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

For further information about programming 64 bit applications, see [Coding standards on 64-bit platforms](#).

If you want to use the programs on a machine that has only the WebSphere MQ MQI client for Solaris installed, compile the programs to link them with the client library (`-lmqic`).

If you use the unsupported compiler `?usr?ucb?cc`, your application might compile and link successfully. However, when you run the application, it fails when it attempts to connect to the queue manager.

Note: 32 bit Solaris x86 SSL and TLS clients configured for FIPS 140-2 compliant operation fail when running on Intel systems. This failure occurs because the FIPS 140-2 compliant GSKit-Crypto Solaris x86 32 bit library file does not load on the Intel chipset. On affected systems, error AMQ9655 is reported in the client error log. To resolve this issue, disable FIPS 140-2 compliance or recompile the client application 64 bit, because 64 bit code is not affected.

Linking libraries

You must link with the WebSphere MQ libraries that are appropriate for your application type:

Library files	Program/exit type
<code>libmqm.so</code>	Server for C
<code>libmqic.so</code> & <code>libmqm.so</code>	Client for C

Note:

1. If you are writing an installable service (for further information, see the [Administering](#)), link to the `libmqmzf.so` library.
2. If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you must link to the `libmqmx.a.so` (or `libmqmx.a64.so` if your transaction manager treats the 'long' type as 64 bit) and `libmqz.so` libraries.
3. You must link WebSphere MQ libraries before any other product libraries.

Building applications on x86-64

This topic contains examples of the commands used to build programs in various environments on the x86-64 platform.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

C client application, 32-bit

```
cc -xarch=386 -mt -o amqsputc_32 amqsputc0.c -IMQ_INSTALLATION_PATH/inc
```

```
-LMQ_INSTALLATION_PATH/lib  
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqic -lsocket -lnsl -ldl
```

C client application, 64-bit

```
cc -xarch=amd64 -mt -o amqspu64 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqic -lsocket  
-lnsl -ldl
```

C server application, 32-bit

```
cc -xarch=386 -mt -o amqspu32 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib  
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqm -lsocket -lnsl -ldl
```

C server application, 64-bit

```
cc -xarch=amd64 -mt -o amqspu64 amqspu0.c -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm -lsocket  
-lnsl -ldl
```

C++ client application, 32-bit

```
CC -xarch=386 -mt -o imqspu32 imqspu.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib  
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as -lmqic -lsocket -lnsl -ldl
```

C++ client application, 64-bit

```
CC -xarch=amd64 -mt -o imqspu64 imqspu.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as  
-limqb23as  
-lmqic -lsocket -lnsl -ldl
```

C++ server application, 32-bit

```
CC -xarch=386 -mt -o imqspu32 imqspu.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib  
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C++ server application, 64-bit

```
CC -xarch=amd64 -mt -o imqspu64 imqspu.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as  
-limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C client exit, 32-bit

```
cc -xarch=386 -mt -G -KPIC -o /var/mqm/exits/cliexit32 cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib  
-R/usr/lib/32 -lmqic  
-lsocket -lnsl -ldl
```

C client exit, 64-bit

```
cc -xarch=amd64 -mt -G -KPIC -o /var/mqm/exits64/cliexit64 cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64  
-R/usr/lib/64 -lmqic  
-lsocket -lnsl -ldl
```

C server exit, 32-bit

```
cc -xarch=386 -mt -G -KPIC -o /var/mqm/exits/srvexit_32 srvexit.c
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib
-R/usr/lib/32 -lmqm
-lsocket -lnsl -ldl
```

C server exit, 64-bit

```
cc -xarch=amd64 -mt -G -KPIC -o /var/mqm/exits64/srvexit_64 srvexit.c
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64
-R/usr/lib/64 -lmqm
-lsocket -lnsl -ldl
```

Building applications on SPARC

This topic contains examples of the commands used to build programs in various environments on the SPARC platform.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

C client application, 32-bit

```
cc -xarch=v8plus -mt -o amqspu32c_32 amqspu32c.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqic -lsocket -lnsl -ldl
```

C client application, 64-bit

```
cc -xarch=v9 -mt -o amqspu32c_64 amqspu32c.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqic
-lsocket -lnsl -ldl
```

C server application, 32-bit

```
cc -xarch=v8plus -mt -o amqspu32c_32 amqspu32c.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -lmqm -lsocket -lnsl -ldl
```

C server application, 64-bit

```
cc -xarch=v9 -mt -o amqspu32c_64 amqspu32c.c -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -lmqm
-lsocket -lnsl -ldl
```

C++ client application, 32-bit

```
CC -xarch=v8plus -mt -o imqspu32c_32 imqspu32c.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as
-lmqic
-lsocket -lnsl -ldl
```

C++ client application, 64-bit

```
CC -xarch=v9 -mt -o imqspu32c_64 imqspu32c.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as
-lmqic -lsocket -lnsl -ldl
```

C++ server application, 32-bit

```
CC -xarch=v8plus -mt -o imqspu32c_32 imqspu32c.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib
```

```
-RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C++ server application, 64-bit

```
CC -xarch=v9 -mt -o imqspu64 imqspu64.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as  
-limqb23as -lmqm  
-lsocket -lnsl -ldl
```

C client exit, 32-bit

```
cc -xarch=v8plus -mt -G -KPIC -o /var/mqm/exits/cliexit_32 cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib  
-R/usr/lib/32 -lmqic  
-lsocket -lnsl -ldl
```

C client exit, 64-bit

```
cc -xarch=v9 -mt -G -KPIC -o /var/mqm/exits64/cliexit_64 cliexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64  
-R/usr/lib/64 -lmqic  
-lsocket -lnsl -ldl
```

C server exit, 32-bit

```
cc -xarch=v8plus -mt -G -KPIC -o /var/mqm/exits/srvexit_32 srvexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib  
-R/usr/lib/32 -lmqm  
-lsocket -lnsl -ldl
```

C server exit, 64-bit

```
cc -xarch=v9 -mt -G -KPIC -o /var/mqm/exits64/srvexit_64 srvexit.c  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64  
-R/usr/lib/64 -lmqm  
-lsocket -lnsl -ldl
```

Preparing COBOL programs in Solaris

Learn about preparing COBOL programs in Solaris.

MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.

1. 32 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

and symbolic links are created in:

```
MQ_INSTALLATION_PATH/inc
```

2. 64 bit COBOL copy books are installed in the following directory:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

3. In the following examples set COBCPY to:

```
MQ_INSTALLATION_PATH/inc/cobcpy32
```

for 32 bit applications, and:

```
MQ_INSTALLATION_PATH/inc/cobcpy64
```

for 64 bit applications.

Compile the programs using Micro Focus compiler. The copy files that declare the structures are in `MQ_INSTALLATION_PATH/inc`:

```
$ export LIB=MQ_INSTALLATION_PATH/lib:$LIB
$ export COBCPY="<COBCPY>"
```

Compiling 32 bit programs:

- `$ cob32 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqmb`
Server for COBOL
- `$ cob32 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb`
Client for COBOL
- `$ cob32 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqmb_r`
Threaded Server for COBOL
- `$ cob32 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib -lmqicb_r`
Threaded Client for COBOL

Compiling 64-bit programs:

- `$ cob64 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmb`
Server for COBOL
- `$ cob64 -xv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb`
Client for COBOL
- `$ cob64 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqmb_r`
Threaded Server for COBOL
- `$ cob64 -xtv amqs0put0.cbl -L MQ_INSTALLATION_PATH/lib64 -lmqicb_r`
Threaded Client for COBOL

where `amqs0put0.cbl` is a sample program.

You must link your program with one of the following:

- `libmqmb.so`
Server for COBOL
- `libmqicb.so`
Client for COBOL

Preparing CICS programs in Solaris

Learn about preparing CICS programs in Solaris.

An XA switch module is provided to enable you to link CICS with WebSphere MQ:

Table 67. Essential code for CICS applications (Solaris)		
Description	C (source)	C (exec)
XA initialization routine	amqzscix.c	amqzsc - TXSeries for Solaris

Always link your transactions with the thread safe WebSphere MQ library libmqm.so.

You can find more information about supporting CICS transactions in the [Administering](#).

TXSeries CICS support

WebSphere MQ for Solaris supports TXSeries CICS using the XA interface.

Write WebSphere MQ programs that are loaded into the same CICS region in either C or COBOL. You cannot make a combination of C and COBOL MQI calls into the same CICS region. Most MQI calls in the second language used fail with a reason code of MQRC_HOBY_ERROR.

Preparing CICS COBOL programs using Micro Focus COBOL

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

To use Micro Focus COBOL, follow these steps:

1. Add the WebSphere MQ COBOL runtime library module to the runtime library using the following command:

```
cicsmkcobol -L/usr/lib/dce -LMQ_INSTALLATION_PATH/lib \
MQ_INSTALLATION_PATH/lib/libmqmcbt.o -lmqe
```

Note: With `cicsmkcobol`, WebSphere MQ does not allow you to make MQI calls in the C programming language from your COBOL application.

If your existing applications have any such calls, move these functions from the COBOL applications to your own library, for example, `myMQ.so`. After moving these functions do not include the WebSphere MQ library `libmqmcbt.o` when building the COBOL application for CICS.

Additionally, if your COBOL application does not make any COBOL MQI call, do not link `libmqmz_r` with `cicsmkcobol`.

This creates the Micro Focus COBOL language method file and enables the CICS runtime COBOL library to call WebSphere MQ on UNIX and Linux systems.

Note: Run `cicsmkcobol` only when you install one of the following products:

- New version or release of Micro Focus COBOL
- New version or release of TXSeries for Solaris
- New version or release of any supported database product (for COBOL transactions only)
- New version or release of WebSphere MQ

2. Export the following environment variable:

```
COBCPY=MQ_INSTALLATION_PATH/inc export COBCPY
```

3. Translate, compile, and link the program by typing:

```
cicstcl -l COBOL -e <yourprog>.ccp
```

Preparing CICS C programs

Build CICS C programs using the standard CICS facilities:

1. Export **one** of the following environment variables:
 - `LDFLAGS = "-LMQ_INSTALLATION_PATH/lib -lmqm_r"` export `LDFLAGS`
 - `USERLIB = "-LMQ_INSTALLATION_PATH/lib -lmqm_r"` export `USERLIB`
2. Translate, compile, and link the program by typing:

```
cicstcl -l C amqscic0.ccs
```

CICS C sample transaction

Sample C source for a CICS WebSphere MQ transaction is provided by AMQSCIC0.CCS. The transaction reads messages from the transmission queue SYSTEM.SAMPLE.CICS.WORKQUEUE on the default queue manager and places them onto the local queue with a queue name that is contained in the transmission header of the message. Any failures are sent to the queue SYSTEM.SAMPLE.CICS.DLQ. Use the sample MQSC script AMQSCIC0.TST to create these queues and sample input queues.

Building your application on Windows systems

The Windows systems publications describe how to build executable applications from the programs that you write.

This topic describes the additional tasks, and the changes to the standard tasks, that you must perform when building WebSphere MQ for Windows applications to run under Windows systems. ActiveX, C, C++, COBOL, and Visual Basic programming languages are supported. For information about preparing your ActiveX programs, see [Using the Component Object Model Interface \(WebSphere MQ Automation Classes for ActiveX\)](#). For information about preparing your C++ programs, see [Using C++](#).

The tasks that you must perform to create an executable application using WebSphere MQ for Windows vary with the programming language that your source code is written in. In addition to coding the MQI calls in your source code, you must add the appropriate language statements to include the WebSphere MQ for Windows include files for the language that you are using. Make yourself familiar with the contents of these files. See [“IBM WebSphere MQ data definition files”](#) on page 77 for a full description.

Building 64-bit applications on Windows

Both 32-bit and 64-bit applications are supported on IBM WebSphere MQ for Windows Version 7.5. The IBM WebSphere MQ executable and library files are supplied in both 32-bit and 64-bit forms, use the appropriate version depending on the application you are working with.

Executable files and libraries

Both 32-bit and 64-bit versions of the IBM WebSphere MQ libraries are supplied in the following locations:

Table 68. Location of IBM WebSphere MQ libraries	
Library version	Directory containing library files
32-bit	<code>MQ_INSTALLATION_PATH\Tools\Lib</code>
64-bit	<code>MQ_INSTALLATION_PATH\Tools\Lib64</code>

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

32-bit applications continue to work normally after migration. The 32-bit files exist in the same directory as in previous versions of the product.

If you want to create 64-bit version you must ensure that your environment is configured to use the library files in `MQ_INSTALLATION_PATH\Tools\Lib64`. Ensure that the LIB environment variable is not set to look in the folder containing the 32-bit libraries.

Preparing C programs in Windows

Work in your typical Windows environment; WebSphere MQ for Windows requires nothing special.

For further information about programming 64-bit applications see [Coding standards on 64-bit platforms](#).

- Link your programs with the appropriate libraries provided by WebSphere MQ:

Library file	Program/exit type
<i>MQ_INSTALLATION_PATH</i> TH\Tools\Lib\mqm.lib	server for 32-bit C
<i>MQ_INSTALLATION_PATH</i> TH \Tools\Lib\mqic.lib	client for 32-bit C
<i>MQ_INSTALLATION_PATH</i> TH \Tools\Lib\mqicxa.lib	client for 32-bit C with transaction co-ordination
<i>MQ_INSTALLATION_PATH</i> TH\Tools\Lib64\mqm.lib	server for 64-bit C
<i>MQ_INSTALLATION_PATH</i> TH\Tools\Lib64\mqic.lib	client for 64-bit C
<i>MQ_INSTALLATION_PATH</i> TH\Tools\Lib64\mqicxa.lib	client for 64-bit C with transaction co-ordination

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

The following command gives an example of compiling the sample program amqsget0 (using the Microsoft Visual C++ compiler).

For 32-bit applications:

```
cl -MD amqsget0.c -Feamqsget.exe MQ_INSTALLATION_PATH\Tools\Lib\mqm.lib
```

For 64-bit applications:

```
cl -MD amqsget0.c -Feamqsget.exe MQ_INSTALLATION_PATH\Tools\Lib64\mqm.lib
```

Note:

- If you are writing an installable service (see the [Administering](#) for further information), you need to link to the mqmzf.lib library.
- If you are producing an application for external coordination by an XA-compliant transaction manager such as IBM TXSeries Encina, or BEA Tuxedo, you need to link to the mqmxa.lib or mqmxa.lib library.
- If you are writing a CICS exit, link to the mqmcics4.lib library.
- You must link WebSphere MQ libraries before any other product libraries.
- The DLLs must be in the path (PATH) that you have specified.
- If you use lowercase characters whenever possible, you can move from WebSphere MQ for Windows to WebSphere MQ on UNIX and Linux systems, where use of lowercase is necessary.

Preparing CICS and Transaction Server programs

Sample C source for a CICS WebSphere MQ transaction is provided by AMQSCIC0.CCS. You build it using the standard CICS facilities. For example, for TXSeries for Windows 2000:

1. Set the environment variable (enter the following code on one line):


```
set CICS_IBMC_FLAGS=-IMQ_INSTALLATION_PATH\Tools\C\Include;
%CICS_IBMC_FLAGS%
```

2. Set the USERLIB environment variable:

```
set USERLIB=MQM.LIB;%USERLIB%
```

3. Translate, compile, and link the sample program:

```
cicstcl -l IBMC amqscic0.ccs
```

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed. This is described in the *Transaction Server for Windows NT Application Programming Guide (CICS) V4*. You can find more information about supporting CICS transactions in the [Administering](#).

Preparing COBOL programs in Windows

Use this information to learn to prepare COBOL programs in Windows, and preparing CICS and Transaction Server programs.

1. The 32 bit COBOL copy books are installed in the following directory:
`MQ_INSTALLATION_PATH\Tools\cobol\CopyBook`.
2. The 64 bit COBOL copy books are installed in the following directory:
`MQ_INSTALLATION_PATH\Tools\cobol\CopyBook64`
3. In the following examples set CopyBook to:

```
CopyBook
```

for 32 bit applications, and:

```
CopyBook64
```

for 64 bit applications.

`MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

To prepare COBOL programs on Windows systems, link your program to one of the following libraries provided by IBM WebSphere MQ:

Library file	Program or exit type
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcbb</code>	32 bit server for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqmcb</code>	32 bit server for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicbb</code>	32 bit client for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib\mqicb</code>	32 bit client for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcbb</code>	64 bit server for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqmcb</code>	64 bit server for Micro Focus COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicbb</code>	64 bit client for IBM COBOL
<code>MQ_INSTALLATION_PATH\Tools\Lib64\mqicb</code>	64 bit client for Micro Focus COBOL

When you are running a program in the MQI client environment, ensure that the DOSCALLS library appears before any COBOL or IBM WebSphere MQ library.

You can use the IBM COBOL Set compiler or Micro Focus COBOL compiler depending on the program:

- Programs beginning amqi are suitable for the IBM COBOL Set compiler,
- Programs beginning amqm are suitable for the Micro Focus COBOL compiler, and
- Programs beginning amq0 are suitable for either compiler.

IBM and Micro Focus COBOL

Relink any existing 32 bit IBM WebSphere MQ Micro Focus COBOL programs using either mqmcb.lib or mqiccb.lib, rather than the mqmccb and mqicccb libraries.

To compile, for example, the sample program amq0put0, using IBM VisualAge COBOL:

1. Set the SYSLIB environment variable to include the path to the IBM WebSphere MQ VisualAge COBOL copybooks (enter the following code on one line):

```
set SYSLIB=MQ_INSTALLATION_PATH\
Tools\Cobol\Copybook\VAcobol;%SYSLIB%
```

2. For use on the IBM WebSphere MQ server:

```
cob2 amq0put0.cbl -qlib "MQ_INSTALLATION_PATH\
Tools\Lib\mqmccb.lib"
```

3. For use on the IBM WebSphere MQ client:

```
cob2 amq0put0.cbl -qlib "MQ_INSTALLATION_PATH\
Tools\Lib\mqicccb.lib"
```

Note: Although you must use the compiler option CALLINT(SYSTEM), this is the default for cob2.

To compile, for example, the sample program amq0put0, using Micro Focus COBOL:

1. Set the COBCPY environment variable to point to the IBM WebSphere MQ COBOL copybooks (enter the following code on one line):

```
set COBCPY=MQ_INSTALLATION_PATH\
Tools\Cobol\Copybook
```

2. Compile the program to give you an object file:

```
cobol amq0put0 LITLINK
```

3. Link the object file to the run time system.

- Set the LIB environment variable to point to the compiler COBOL libraries.
- Link the object file for use on the IBM WebSphere MQ server:

```
cbllink amq0put0.obj mqmcb.lib
```

- Or link the object file for use on the IBM WebSphere MQ client:

```
cbllink amq0put0.obj mqiccb.lib
```

Preparing CICS and Transaction Server programs

To compile and link a TXSeries for Windows NT, V5.1 program using IBM VisualAge COBOL:

1. Set the environment variable (enter the following code on one line):

```
set CICS_IBMCOB_FLAGS=MQ_INSTALLATION_PATH\
Cobol\Copybook\VAcobol;%CICS_IBMCOB_FLAGS%
```

2. Set the USERLIB environment variable:

```
set USERLIB=MQMCBB.LIB
```

3. Translate, compile, and link your program:

```
cicstcl -l IBMCOB myprog.ccp
```

This is described in the *Transaction Server for Windows NT, V4 Application Programming Guide*.

To compile and link a CICS for Windows V5 program using Micro Focus COBOL:

- Set the INCLUDE variable:

```
set
INCLUDE=<drive>:\<programname>\ibm\websphere\tools\c\include;
<drive>:\opt\cics\include;%INCLUDE%
```

- Set the COBCPY environment variable:

```
setCOBCPY=<drive>:\<programname>\ibm\websphere\tools\cobol\copybook;
<drive>:\opt\cics\include
```

- Set the COBOL options:

- set
- COBOPTS=/LITLINK /NOTRUNC

and run the following code:

```
cicstran cicsmq00.ccp
cobol cicsmq00.cbl /LITLINK /NOTRUNC
cbllink -D -Mcicsmq00 -Ocicsmq00.cbmfnt cicsmq00.obj
%CICSLIB%\cicsprCBMFNT.lib user32.lib msvcrt.lib kernel32.lib mqmcb.lib
```

Preparing Visual Basic programs in Windows

Use this information when considering using Visual Basic programs on Windows.

Note: 64-bit versions of the Visual Basic module files are not supplied.

To prepare Visual Basic programs on Windows:

1. Create a new project.
2. Add the supplied module file, CMQB.BAS, to the project.
3. Add other supplied module files if you need them:

CMQBB.BAS	MQAI support
CMQCFB.BAS	PCF support
CMQXB.BAS	Channel exits support
CMQPSB.BAS	Publish/subscribe

See [“Coding in Visual Basic” on page 83](#) for information about using the MQCONNXAny call from within Visual Basic.

Call the procedure MQ_SETDEFAULTS before making any MQI calls in the project code. This procedure sets up default structures that the MQI calls require.

Specify whether you are creating a WebSphere MQ server or client, before you compile or run the project, by setting the conditional compilation variable *MqType*. Set *MqType* in a Visual Basic project to 1 for a server or 2 for a client as follows:

1. Select the Project menu.
2. Select *Name* Properties (where *Name* is the name of the current project).
3. Select the Make tab in the dialog box.
4. In the Conditional Compilation Arguments field, enter this for a server:

```
MqType=1
```

or this for a client:

```
MqType=2
```

SSPI security exit

WebSphere MQ for Windows supplies a security exit for both the WebSphere MQ MQI client and the WebSphere MQ server. This is a channel-exit program that provides authentication for WebSphere MQ channels by using the Security Services Programming Interface (SSPI). The SSPI provides the integrated security facilities of Windows systems.

The security packages are loaded from either security.dll or secur32.dll. These DLLs are supplied with your operating system.

One-way authentication is provided using NTLM authentication services. Two-way authentication is provided using Kerberos authentication services.

The security exit program is supplied in source and object format. You can use the object code as it is, or you can use the source code as a starting point to create your own user-exit programs.

See also [“Using the SSPI security exit on Windows systems” on page 162](#).

Introduction to security exits

A security exit forms a secure connection between two security exit programs, where one program is for the sending message channel agent (MCA), and one is for the receiving MCA.

The program that initiates the secure connection, that is, the first program to get control after the MCA session is established, is known as the *context initiator*. The partner program is known as the *context acceptor*.

The following table shows some of the channel types that are context initiators and their associated context acceptors.

Table 69. Context initiators and their associated context acceptors	
Context Initiator	Context Acceptor
MQCHT_CLNTCONN	MQCHT_SVRCONN
MQCHT_RECEIVER	MQCHT_SENDER
MQCHT_CLUSRCVR	MQCHT_CLUSSDR

The security exit program has two entry points:

- **SCY_NTLM**

This uses NTLM authentication services, which provide one-way authentication. NTLM allows servers to verify the identities of their clients. It does not allow clients to verify a server's identity, or one server to verify the identity of another. NTLM authentication was designed for a network environment in which servers are assumed to be genuine.

- **SCY_KERBEROS**

This uses Kerberos mutual authentication services. The Kerberos protocol does not assume that servers in a network environment are genuine. Parties at both ends of a network connection can verify the identity of the other party. That is, servers can verify the identity of clients and other servers, and clients can verify the identity of a server.

What the security exit does

This topic describes what the SSPI channel-exit programs do.

The supplied channel-exit programs provide either one-way or two-way (mutual) authentication of a partner system when a session is being established. For a particular channel, each exit program has an associated *principal* (similar to a user ID, see [“WebSphere MQ access control and Windows principals” on page 445](#)). A connection between two exit programs is an association between the two principals.

After the underlying session is established, a secure connection between two security exit programs (one for the sending MCA and one for the receiving MCA), is established. The sequence of operations is as follows:

1. Each program is associated with a particular principal, for example as a result of an explicit login operation.
2. The context initiator requests a secure connection with the partner from the security package (for Kerberos, the named partner) and receives a token (called token1). The token is sent, using the underlying session that is already established, to the partner program.
3. The partner program (the context acceptor) passes token1 to the security package, which verifies that the context initiator is authentic. For NTLM, the connection is now established.
4. For the Kerberos-supplied security exit (that is, for mutual authentication), the security package also generates a second token (called token2), which the context acceptor returns to the context initiator by using the underlying session.
5. The context initiator uses token2 to verify that the context acceptor is authentic.
6. At this stage, if both applications are satisfied with the authenticity of the partner's token, the secure (authenticated) connection is established.

WebSphere MQ access control and Windows principals

The access control that WebSphere MQ provides is based on the user and group. The authentication that Windows provides is based on principals, such as user and servicePrincipalName (SPN). In the case of servicePrincipalName, there might be many of these associated with a single user.

The SSPI security exit uses the relevant Windows principals for authentication. If Windows authentication is successful, the exit passes the user ID that is associated with the Windows principal to WebSphere MQ for access control.

The Windows principals that are relevant for authentication vary, depending on the type of authentication used.

- For NTLM authentication, the Windows principal for Context Initiator is the user ID associated with the process that is running. Because this authentication is one way, the principal associated with the Context Acceptor is irrelevant.

- For Kerberos authentication, on CLNTCONN channels, the Windows principal is the user ID associated with the process that is running. Otherwise, the Windows principal is the servicePrincipalName that is formed by adding the following prefix to the QueueManagerName.

```
ibmMQSeries/
```

Using lightweight directory access protocol services with WebSphere MQ for Windows

This topic explains what a directory service is and the part played by a directory access protocol (DAP). It also explains how WebSphere MQ applications can use a lightweight directory access protocol (LDAP) directory using a sample program as a guide.

Note: The sample program is designed for someone who is already familiar with LDAP.

The following topics give more information about directory services, LDAP, and using LDAP with WebSphere MQ.

- [“Directory Service” on page 446](#)
- [“Lightweight Directory Access Protocol \(LDAP\)” on page 446](#)
- [“Using LDAP with WebSphere MQ” on page 447](#)

Directory Service

A directory is a repository of information about objects, which is organized in such a way that it is easy to find the information about a specific object.

A common example is a telephone directory, where information (address and telephone number) is stored about people and companies. Another example is an address book for an email system, where email addresses, and optionally other information such as telephone numbers, are stored for people.

On computer systems, directories can store information about computer resources, such as printers or shared disks. For example, you could use a directory to find out where the nearest color printer is located. In a WebSphere MQ application a directory can be used to provide the association between an application service (such as accounts-receivable processing) and the queue to be used for messages requiring that service (possibly identified through the queue name and its host queue manager name).

Directories are implemented as client-server systems, where the directory server holds all the information and answers requests from clients. The clients could be user-interface programs, which provide the information directly to the user, or application programs which need to locate resources to complete their work. A Directory Service comprises the directory server, administrative programs, and the client libraries and programs that are needed to configure, update, and read the directory.

Lightweight Directory Access Protocol (LDAP)

Many directory services exist, such as Novell Directory Services, DCE Cell Directory Service, Banyan StreetTalk, Windows Directory Services, X.500, and the address book services associated with email products. X.500 was proposed as a standard for global directory services by the International Standards Organization (ISO). It requires an OSI protocol stack for its communications, and largely because of this, its use has been restricted to large organizations and academic institutions. An X.500 directory server communicates with its clients using the Directory Access Protocol (DAP).

LDAP (Lightweight Directory Access Protocol) was created as a simplified version of DAP. It is easier to implement, omits some of the lesser-used features of DAP, and runs over TCP/IP. As a result of these changes it is rapidly being adopted as the directory access protocol for most purposes, replacing the multitude of proprietary protocols previously used. LDAP clients can still access an X.500 server through a gateway (X.500 still requires the OSI protocol stack), or increasingly X.500 implementations typically include native support for LDAP as well as DAP access.

LDAP directories can be distributed and can use replication to enable efficient access to their contents.

For a more complete description of LDAP, see *Understanding LDAP*, an IBM Redbooks® publication.

Using LDAP with WebSphere MQ

In WebSphere MQ configurations, the information that defines message and transmission queues is stored locally. This means that in a WebSphere MQ network the various definitions are distributed, with no central directory of this information being available for browsing. Remote messaging between WebSphere MQ applications is commonly achieved through the use of local definitions of remote queues. The application first issues an MQOPEN call using the name specified in the local definition of the remote queue. To put the message on the remote queue, the application then issues MQPUT, specifying the handle returned from the MQOPEN call. The remote queue definition supplies the name of the destination queue, the destination queue manager, and optionally, a transmission queue. In this technique the application has to know at run time the name specified in the local queue definition.

A variation on the previous avoids the use of local definitions of remote queues. The application can specify the full destination queue name, which includes the remote queue manager name as part of the MQOPEN. The application therefore has to know these two names at run time. The local queue manager must be correctly configured with the local queue definition, and with a suitably named (or default) transmission queue and an associated channel that delivers to the target.

In the case where both the source and target queue managers are defined as being members of the same cluster, the transmission queue and channel aspects of the previous two scenarios can be ignored. If the target transmission queue is a cluster queue, a local definition of a remote queue is also not required. However, similarly to the previous cases described, the application must still know the name of the destination queue.

A directory service can be used to remove this application dependency on queue names (or the combination of queue and queue manager names). The mapping between application criteria and WebSphere MQ object names can be held in a directory and updated dynamically, and independently of applications. At run time the WebSphere MQ application that wants to send a message first queries the directory using application-based criteria, for example where: `service_name = "accounts receivable"`, retrieves the relevant WebSphere MQ object names, and then uses these returned values in the MQOPEN call.

Another example of the use of a directory is for a company that has many small depots or offices, WebSphere MQ MQI clients can be used to send messages to WebSphere MQ servers located in the larger offices. The clients need to know the name of the host machine, MQI channel, and queue name for each server that they send messages to. Occasionally, it might be necessary to move a WebSphere MQ server to another machine; every client that communicates with the server would need to know about the change. An LDAP directory service could be used to store the names of the host machines (and the channel and queue names) and the client programs could retrieve the information from the directory whenever they want to send a message to a server. In this case only the directory needs to be updated if a host name (or channel or queue name) changed.

Multiple destinations for an application message could be stored in a directory, with the one chosen being dependent on availability or load-sharing considerations.

WebSphere MQ can also use an LDAP directory to store authentication information for use with Secure Sockets Layer (SSL). WebSphere MQ classes for Java can also store information in an LDAP directory.

LDAP sample program

The sample program is designed for someone who is familiar with LDAP and probably already uses it. It is intended to show how WebSphere MQ applications can use an LDAP directory.

Building the sample program

This program has been built and tested only on Windows using TCP/IP. As well as the general considerations mentioned in [“Preparing C programs in Windows” on page 439](#), note the following points:

- This program is designed to run as a client program, so it should be linked with the MQIC.LIB library.

- As well as the WebSphere MQ header files and libraries, this program must be built using LDAP client header files and libraries.

For example, using the IBM eNetwork client, link the program with the LIBLDAPSTATICE.LIB and LIBLBERSTATICSSL.LIB libraries.

Configuring the directory

Before the sample program can be run, an LDAP Directory Server must be configured with sample data.

The file MQuser.ldif, in the tools\c\samples directory, contains some sample data in LDIF (LDAP Data Interchange Format). You can edit this file to suit your needs. It contains data for a fictitious company called MQuser that has a Transport Department comprising three offices. Each of these offices has a machine that runs a WebSphere MQ server.

As a minimum, you must edit the three lines that contain the host names of the machines running the WebSphere MQ servers: lines 18, 27, and 36:

```
host: LondonHost
...
host: SydneyHost
...
host: WashingtonHost
```

You must change LondonHost, SydneyHost, and WashingtonHost to the names of three of your machines that run WebSphere MQ servers. You can also change the channel and queue names if you want (the sample uses names of the system defaults). You might also want to increase or decrease the number of offices in the sample data.

Configuring the IBM Tivoli directory server

Refer to the IBM Tivoli Directory Server (ITDS) Administrator's Guide for information about installing the directory. In the topic *Installing and Configuring Server*, work through the sections *Installing Server* and *Basic Server Configuration*. If necessary, read through the topic *Administrator Interface* to familiarize yourself with how the interface works.

In the topic *Configuring - How Do I*, follow the instructions to start the administrator, then work through the section *Configure Database* and create a default database. Skip the section *Configure replica* and using the section *Work with Suffixes*, add a suffix `o=MQuser`.

Before adding any entries to the database, you must extend the directory schema by adding some attribute definitions and an objectclass definition. This is described in the IBM Tivoli Directory Server Administrator's Guide in the chapter *Reference Information* under the section *Directory Schema*. Two sample files are included to help you with this. The file `mq.at.conf` includes the attribute definitions that you must add to the file `?etc?slapd.at.conf`. Do this by including the sample file by editing `slapd.at.conf` and adding a line:

```
include <pathname>/mq.at.conf
```

Alternatively you can edit the file `slapd.at.conf` and add the contents of the sample file directly to it, that is, add the lines:

```
# MQ attribute definitions
attribute mqChannel          ces    mqChannel    1000    normal
attribute mqQueueManager    ces    mqQueueManager 1000    normal
attribute mqQueue            ces    mqQueue      1000    normal
attribute mqPort             cis    mqPort        64     normal
```

Similarly for the objectclass definition, you can either include the sample file by editing `etc?slapd.oc.conf` and add the line:

```
include <pathname>/mq.oc.conf
```


or you can add the contents of the sample file directly to `slapd.oc.conf`, that is, add the lines:

```
# MQ object classdefinition
objectclass mqApplication
    requires
        objectClass,
        cn,
        host,
        mqChannel,
        mqQueue
    allows
        mqQueueManager,
        mqPort,
        description,
        l,
        ou,
        seeAlso
```

You can now start the directory server (Administration, Server, Startup) and add the sample entries to it. To add the sample entries, go to the Administration, Add Entries page of the administrator, type in the full path name of the sample file `MQuser.ldif` and click Submit.

The directory server is now running and loaded with data suitable for running the sample program.

Configuring the Netscape directory server

Using the Netscape Server Administration page, click **Create New Netscape Directory Server**.

You should now be presented with a form containing configuration information. Change the Directory Suffix to **o=MQuser** and add a password for the Unrestricted User. You can also change any other information to suit your installation. Click **OK**, and the directory should be created successfully. Click **Return to Server Administration** and start the directory server. Click the directory name to start the Directory Server Administration server for the new directory.

Before adding any entries to the database, extend the directory schema by adding some attribute definitions and an objectclass definition. Click the **Schema** tab of the Directory Server page. You are now presented with a form that allows you to add new attributes. Add the following attributes (leave the Attribute OID blank for all of them):

Attribute Name	Syntax
mqChannel	Case Exact String
mqQueueManager	Case Exact String
mqQueue	Case Exact String
mqPort	Integer

Add a new objectClass by clicking **Create ObjectClass** in the side panel. Enter **mqApplication** as the ObjectClass Name, select **applicationProcess** as the parent ObjectClass and leave the **ObjectClass OID** blank. Now add some attributes to the objectClass. Select **host**, **mqChannel**, and **mqQueue** as Required Attributes, and select **mqQueueManager** and **mqPort** as Allowed attributes. Press the **Create New ObjectClass** button to create the objectClass.

To add the sample data, click the **Database Management** tab and select **Add Entries** from the side panel. Enter the path name of the sample data file `<pathname>\MQuser.ldif`, enter the password, and click **OK**.

The sample program runs as an unauthorized user, and by default the Netscape Directory does not allow unauthorized users to search the directory. Change this by clicking the **Access Control** tab. Enter the password for the Unrestricted User and click **OK** to load in the access control entries for the directory. These should currently be empty. Press the **New ACI** button to create a new access control entry. In the entry box that is displayed, click **Deny** (which is underlined) and in the resultant dialog box, change it to **Allow**. Add a name, for example, **MQuser-access**, and click **choose a suffix** to select **o=MQuser**. Enter **o=MQuser** as the target, enter the password for the Unrestricted User, and click **Submit**.

The directory server is now running and loaded with data suitable for running the sample program.

Running the sample program

You should now have an LDAP Directory Server running and populated with the sample data. The data specifies three host machines, all of which should be running WebSphere MQ servers. Ensure that the default queue manager is running on each machine (unless you changed the sample data to specify a different queue manager).

Also, start the WebSphere MQ listener program on each machine; the sample uses TCP/IP with the default WebSphere MQ port number, so you can start the listener with the command:

```
runmqclsr -t tcp
```

To test the sample, you might also want to run a program to read the messages arriving at each WebSphere MQ server, for example you could use the amqstrg sample program:

```
amqstrg SYSTEM.DEFAULT.LOCAL.QUEUE
```

The sample program uses three environment variables, one required and two optional. The required variable is LDAP_BASEDN, which specifies the base Distinguished Name for the directory search. To work with the sample data, set this to ou=Transport, o=MQuser, for example, at a command prompt on Windows systems type:

```
set LDAP_BASEDN=ou=Transport, o=MQuser
```

The optional variables are LDAP_HOST and LDAP_VERSION. The LDAP_HOST variable specifies the name of the host where the LDAP server is running; it defaults to the local host if it is not specified. The LDAP_VERSION variable specifies the version of the LDAP protocol to be used, and can be either 2 or 3. Most LDAP servers now support version 3 of the protocol; they all support the older version 2. This sample works equally well with either version of the protocol, and if it is not specified it defaults to version 2.

You can now run the sample by typing the program name followed by the name of the WebSphere MQ application that you want to send messages to, in the case of the sample data the application names are London, Sydney, and Washington. For example, to send messages to the London application:

```
amqslipc London
```

If the program fails to connect to the WebSphere MQ server, an appropriate error message appears. If it connects successfully you can start typing messages, each line that you type (terminated by <return> or <enter>) is sent as a separate message, an empty line ends the program.

Program design

The program has two distinct parts: the first part uses the environment variables and command line value to query an LDAP directory server; the second part establishes the WebSphere MQ connection using the information returned from the directory and sends the messages.

The LDAP calls used in the first part of the program differ slightly depending on whether LDAP version 2 or 3 is being used, and they are described in detail by the documentation that comes with the LDAP client libraries. This section gives a brief description.

The first part of the program checks that it has been called correctly and reads the environment variables. It then establishes a connection with the LDAP directory server at the specified host:

```
if (ldapVersion == LDAP_VERSION3)
{
    if ((ld = ldap_init(ldapHost, LDAP_PORT)) == NULL)
        ...
}
else
{
    if ((ld = ldap_open(ldapHost, LDAP_PORT)) == NULL )
```

```
} ...
```

When a connection has been established, the program sets some options on the server with the "ldap_set_option" call, and then authenticates itself to the server by binding to it:

```
if (ldapVersion == LDAP_VERSION3)
{
    if (ldap_simple_bind_s(ld, bindDN, password) != LDAP_SUCCESS)
        ...
}
else
{
    if (ldap_bind_s(ld, bindDN, password, LDAP_AUTH_SIMPLE) !=
        LDAP_SUCCESS)
        ...
}
```

In the sample program bindDN and password are set to NULL, which means that the program authenticates itself as an anonymous user, that is, it does not have any special access rights and can access only information that is publicly available. In practice, most organizations restrict access to the information that they store in directories so that only authorized users can access it.

The first parameter to the bind call ld is a handle that is used to identify this particular LDAP session throughout the rest of the program. After authenticating, the program searches the directory for entries that match the application name:

```
rc = ldap_search_s(ld,                /* LDAP Handle */
                  baseDN,             /* base distinguished name */
                  LDAP_SCOPE_ONELEVEL, /* one-level search */
                  filterPattern,      /* filter search pattern */
                  attrs,              /* attributes required */
                  FALSE,              /* NOT attributes only */
                  &ldapResult);      /* search result */
```

This is a simple synchronous call to the server that returns the results directly. There are other types of search that are more appropriate for complex queries or when a large number of results is expected. The first parameter to the search is the handle ld that identifies the session. The second parameter is the base distinguished name, which specifies where in the directory the search is to begin, and the third parameter is the scope of the search, that is, which entries relative to the starting point are searched. These two parameters together define which entries in the directory are searched. The next parameter, filterPattern specifies what we are searching for. The attrs parameter lists the attributes that we want to get back from the object when we have found it. The next attribute says whether we want just the attributes or their values as well; setting this to FALSE means that we want the attribute values. The final parameter is used to return the result.

The result could contain many directory entries, each with the specified attributes and their values. We have to extract the values that we want from the result. In this sample program we only expect one entry to be found, so we only look at the first entry in the result:

```
ldapEntry = ldap_first_entry(ld, ldapResult);
```

This call returns a handle that represents the first entry, and we set up a for loop to extract all the attributes from the entry:

```
for (attribute = ldap_first_attribute(ld, ldapEntry, &ber);
     attribute != NULL;
     attribute = ldap_next_attribute(ld, ldapEntry, ber ))
{
```

For each of these attributes, we extract the values associated with it. Again we only expect one value per attribute, so we only use the first value; we determine which attribute we have and store the value in the appropriate program variable:

```

values = ldap_get_values(ld, ldapEntry, attribute);
if (values != NULL && values[0] != NULL)
{
    if (strcmp(attribute, MQ_HOST_ATTR) == 0)
    {
        mqHost = strdup(values[0]);
        ...
    }
}

```

Finally we tidy up by freeing memory (`ldap_value_free`, `ldap_memfree`, `ldap_msgfree`) and close the session by *unbinding* from the server:

```

ldap_unbind(ld);

```

We check that we have found all the WebSphere MQ values that we need from the directory, and if so we call `sendMessages()` to connect to the WebSphere MQ server and send the WebSphere MQ messages.

The second part of the sample program is the `sendMessages()` routine that contains all the WebSphere MQ calls. This is modelled on the `amqspu0` sample program, the differences being that the parameters to the program have been extended and `MQCONN` is used instead of the `MQCONN` call.

Developing applications for IBM WebSphere MQ Telemetry

Telemetry applications integrate sense and control devices with other sources of information available on the internet and in enterprises.

Develop applications for IBM WebSphere MQ Telemetry using design patterns, worked examples, sample programs, programming concepts, and reference information. Use the IBM WebSphere MQ Telemetry daemon for devices to simplify connecting many small devices to IBM WebSphere MQ.

Related concepts

[WebSphere MQ Telemetry](#)

[Telemetry concepts and scenarios for monitoring and control](#)

Related tasks

[Installing WebSphere MQ Telemetry](#)

[Administering WebSphere MQ Telemetry](#)

[Troubleshooting for WebSphere MQ Telemetry](#)

Related reference

[WebSphere MQ Telemetry Reference](#)

IBM WebSphere MQ Telemetry sample programs

Sample scripts are provided to demonstrate the basic usage of the MQ Telemetry Transport v3 Client application. Use the scripts to publish a message and subscribe to a topic.

Before you begin

Start the telemetry (MQXR) service to run the sample programs.

The user ID must be a member of the `mqm` user group.

Run the `SampleMQM` script first, followed by the `MQTTV3Sample` script to perform a publish and subscribe. Run the `CleanupMQM` sample script to delete the queue manager created by the `SampleMQM` script.

As the `SampleMQM` script creates and uses a queue manager called `QM1`, do not run unaltered on a system with a `QM1` queue manager. Any changes made might have implications for the configuration of the existing queue manager.

About this task

- The SampleMQM application creates and starts a telemetry enabled queue manager called QM1. The script also sets up a default transmission queue for QM1, and creates and starts a default channel listening on port 1883. This channel performs no authentication of clients connected to it. The channel has the message channel agent user identifier (MCAUSER) attribute, set to 'guest' on Windows systems or 'nobody' on Linux systems. Clients connected to the channel are treated like the user 'guest' or user 'nobody', depending on the system it is running on. The script authorizes 'guest' on Windows systems and 'nobody' on Linux systems to be able to publish and subscribe to any topic on QM1
- The MQTTV3Sample application is held in the following location;
 - On Windows `MQ_INSTALLATION_PATH\mqxr\samples`
where `MQ_INSTALLATION_PATH` is the location in which IBM WebSphere MQ is installed.
 - On Linux `MQ_INSTALLATION_PATH/mqxr/samples`The MQTTV3Sample application works as a publisher, sending a single message to a topic on the server. It also acts as a subscriber, listening for messages from the server.
- The CleanupMQM sample script ends and deletes QM1 that was created by the SampleMQM script. Use the CleanupMQM sample script if you want to rerun the SampleMQM script or remove QM1.

Procedure

1. Type the following command on a command line to run the SampleMQM script

- On Windows, the command to run the SampleMQM script is as follows:

```
MQ_INSTALLATION_PATH\mqxr\samples\SampleMQM.bat
```

- On AIX and Linux, the command to run the SampleMQM script is as follows:

```
MQ_INSTALLATION_PATH/mqxr/samples/SampleMQM.sh
```

where `MQ_INSTALLATION_PATH` is the location in which IBM WebSphere MQ is installed.

A queue manager called MQXR_SAMPLE_QM is created.

2. Type the following command to run the first part of the MQTTV3Sample script;

- On Windows, on one command line, type the following command;

```
MQ_INSTALLATION_PATH\mqxr\samples\RunMQTTV3Sample.bat -a subscribe
```

- On AIX and Linux, in one shell window, type the following command;

```
MQ_INSTALLATION_PATH/mqxr/samples/RunMQTTV3Sample.sh -a subscribe
```

3. Type the following command to run second part of the MQTTV3Sample script;

- On Windows, on another command line, type the following command;

```
MQ_INSTALLATION_PATH\mqxr\samples\RunMQTTV3Sample.bat -m "Hello from an MQTT v3 application"
```

- On AIX and Linux, in another shell window, type the following command;

```
MQ_INSTALLATION_PATH/mqxr/samples/RunMQTTV3Sample.sh -m "Hello from an MQTT v3 application"
```

4. To remove the queue manager created by the SampleMQM script, you can run the CleanupMQM script using the following command;

- On Windows, type the following command;

```
MQ_INSTALLATION_PATH\mqxr\samples\CleanupMQM.bat
```

- On AIX and Linux in another shell window, type the following command;

```
MQ_INSTALLATION_PATH/mqxr/samples/CleanupMQM.sh
```

Results

The Hello from an MQTT v3 application message, that you typed in the second window will be published by that application, and received by the application in the first window. The application in the first window will show it on the screen.

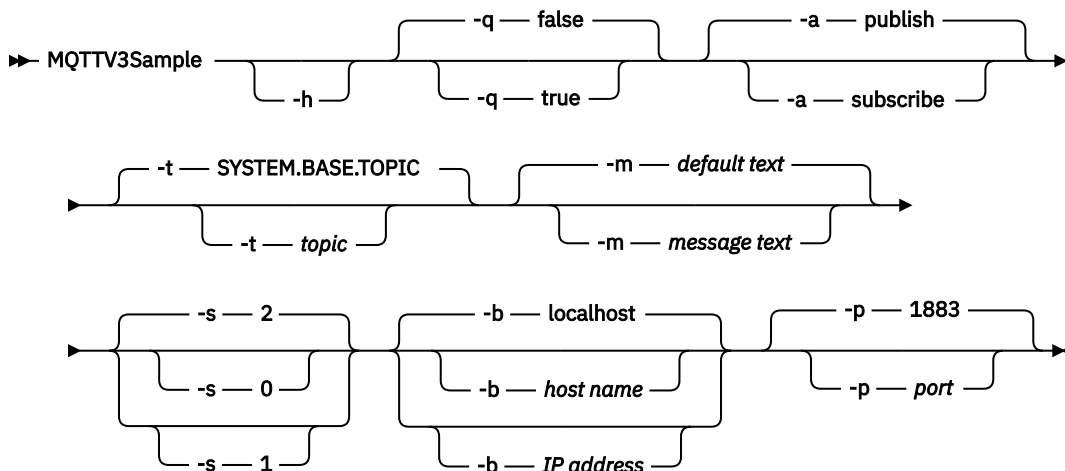
MQTTV3Sample program

Reference information about sample syntax and parameters for the MQTTV3Sample program.

Purpose

The MQTTV3Sample program can be used to publish a message and subscribe to a topic.

MQTTV3Sample syntax



Parameters

- h**
Print this help text and quit
- q**
Set quiet mode, instead of using the default mode of false.
- a**
Set publish or subscribe, instead of assuming the default action of publishing.
- t**
Publish or subscribe to topic, instead of publishing or subscribing to the default topic
- m**
Publish message text instead of sending the default publication text, "Hello from an MQTT v3 application".
- s**
Set QoS instead of using the default QoS, 2.
- b**
Connect to this host name or IP address instead of connecting to the default host name, localhost.
- p**
Use this port instead of using the default, 1883.

Run the MQTTV3Sample program

To subscribe to a topic on Windows, use the command:

```
runMQTTV3Sample -a subscribe
```

To publish a message on Windows, use the command:

```
runMQTTV3Sample
```

For further information about running the sample scripts provided, see [“IBM WebSphere MQ Telemetry sample programs” on page 452](#).

Creating your first MQ Telemetry Transport publisher application using Java

The steps to create an MQTT client application are described in tutorial fashion. Each line of code is explained. At the end of the task, you will have created an MQTT publisher. You can browse the publications using WebSphere MQ Explorer.

Before you begin

Install the WebSphere MQ Telemetry feature on a server that has IBM WebSphere MQ Version 7.1 or later installed.

The client application uses the `com.ibm.mq.micro.client.mqttv3` package in the IBM WebSphere MQ Telemetry Software Development Toolkit (SDK). The SDK is part of the IBM WebSphere MQ Telemetry installation. The client connects to the IBM WebSphere MQ Telemetry feature to exchange messages with IBM WebSphere MQ.

You must also install the telemetry updates for IBM WebSphere MQ Explorer Version 7.1 to administer IBM WebSphere MQ Telemetry. The updates are part of the IBM WebSphere MQ Telemetry installation.

An MQTT client, running on Java SE, requires version 6.0 of Java SE, or later. IBM Java SE v6.0 is part of the IBM WebSphere MQ Version 7.1 installation. It is located at *WebSphere MQ installation directory\java\jre*

About this task

The example is a publish application, PubSync. PubSync publishes Hello World on the topic MQTT Examples, and waits for confirmation that the publication has been delivered to the queue manager.

By setting up a durable subscription to MQTT Examples you can check that the application works.

The procedure uses Eclipse to develop, build, and run the client. You can download Eclipse from the Eclipse project website at www.eclipse.org.

To create the application, you can create the Java files, and compile and run them by using the command line.

In a new directory, create the directory path `.\com\ibm\mq\id`. Create two Java files, `Example.java` and `PubSync.java`. Copy the code from [“Example code” on page 459](#) into the Java files.

Compile the Java code using the command,

```
javac -cp jar_dir\com.mq.micro.client.mqttv3.jar  
com.ibm.mq.id.PubSync.java com.ibm.mq.id.Example.java
```

Run PubSync using the command,

```
java -cp jar_dir\com.mq.micro.client.mqttv3.jar  
com.ibm.mq.id.PubSync
```


Procedure

1. Create a Java project in Eclipse.
 - a) **File > New > Java project** and type a project name. Click **Next**.
Check the JRE is at the correct or later version. Java SE must be at 6.0 or later.
 - b) On the Java Settings page, click **Libraries > Add External Jars...**
 - c) Browse to the directory where you have installed the WebSphere MQ Telemetry SDK folder to. Locate the SDK\clients\java folder and select the all the .jar files > **Open > Finish**.
2. Install the MQTT client Javadoc.

With the MQTT client Javadoc installed, the Java editor provides assistance with the MQTT v3 classes.

- a) In your Java project, open **Package Explorer > Referenced Libraries**. Right click `com.ibm.micro.client.mqttv3.jar` > **Properties**.
 - b) In the Properties navigator click **Javadoc Location**.
 - c) In the Javadoc Location page click **Javadoc URL > Browse...** and find the *WMQ Installation directory*\mqxr\SDK\clients\java\doc\javadoc folder > **OK**.
 - d) Click **Validate... > OK**
- You are prompted to open a browser to view the documentation.
3. Create the class, PubSync, using the Java Class wizard.
 - a) Right-click the Java project you have created > **New > Class**.
 - b) Type the package name, `com.ibm.mq.id`
 - c) Type the class name, PubSync
 - d) Check the method stub box, **public static void main(String [] args)**
 4. Create a file, Example.java in the package `com.ibm.mq.id`. Copy the code in [Figure 89 on page 460](#) into the file.

All the parameters used in the examples are set as properties. You can override the values by changing the defaults in Example.java, or by supplying the properties as options on the Java command line using the -D parameter:



The client identifier used in this example, and the “[Creating an asynchronous publisher for MQ Telemetry Transport using Java](#)” on page 460 examples, is a user name suffixed by a random string.

5. Follow the steps to create the code, or copy the code from [Figure 88 on page 459](#).

The steps that follow explain the code in Pubsync.java.

6. Create a try-catch block.

```
try { ...  
} catch (Exception e) {  
    e.printStackTrace();  
}
```

The MQTT client throws `MqttException`, `MqttPersistenceException` or `MqttSecurityException`. `MqttPersistenceException` and `MqttSecurityException` are subclasses of `MqttException`.

Use the `MqttException.getReasonCode` method to find out the reason for the exception. If a `MqttPersistenceException` or `MqttSecurityException` is thrown, use the `getCause` method to return the underlying throwable exception.

7. Create a new `MqttClient` instance.

```
MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);
```

Provide the client with a server address, which is used later to connect to WebSphere MQ. Set the client identifier to name the client.

- Optionally, you can provide an implementation of the `MqttClientPersistence` interface to replace the default implementation. The default `MqttPersistence` implementation stores QoS 1 and 2 messages awaiting delivery as files; see [“Message persistence in MQTT clients” on page 510](#).
- The default IBM WebSphere MQ TCP/IP port for MQTT is 1883. For SSL, it is 8883. In the example, the default address is set to `tcp://localhost:1883`.
- Typically, it is important to be able to identify a specific physical client using the client identifier. The client identifier must be unique across all clients connect to a server; see [“Client identifier” on page 506](#). Using the same client identifier as a previous instance indicates the present instance is an instance of the same client. If you duplicate a client identifier in two running clients, an exception is thrown in both clients, and one client terminates.
- The length of the client identifier is limited to 23 bytes. An exception is thrown, if the length is exceeded. The client identifier must contain only characters permitted in a queue manager name; for example, no hyphens or spaces.
- Until you call the `MqttClient.connect` method, no message processing takes place.

Use the client object to publish and subscribe topics and recover information about publications that have not been delivered yet.

8. Create a topic to publish on.

```
MqttTopic topic = client.getTopic(Example.topicString);
```

A topic string is limited to 64 K bytes, which exceeds the maximum length of an IBM WebSphere MQ topic string. Otherwise, a topic string follows the same rules as WebSphere MQ topic strings; see [Topic strings](#). The example sets the topic string `MQTT Examples`.

9. Create a publication message.

```
MqttMessage message = new MqttMessage(Example.publication.getBytes());
```

The string `"Hello World"` is converted to a byte array and used to create an `MqttMessage`.

- An MQTT message payload is always a byte array. The `getBytes` method converts a string object to UTF-8. The `MqttMessage` has a convenience `toString` method to return the message payload as a string. It is equivalent to `new String(message.getPayload)`
- A publication message is sent to the queue manager with an RFH2 header, and the message data is sent as a `jms-bytes` message.
- The message object has quality of service and retained attributes. The quality of service (QoS) determines how reliably the message is transferred between the MQTT client and the queue manager; see [“Qualities of service provided by an MQTT client” on page 514](#). The retained attribute controls if a publication is stored by the queue manager for future subscribers. If a publication is not retained, it is sent only to current subscribers; see [“Retained publications and MQTT clients” on page 516](#). The default `MqttMessage` settings are, "Messages are delivered at least once, and are not retained."

10. Connect to the server.

```
client.connect();
```

The example connects to the server using the default connection options. Once you connect, you can start publishing. The default connection options are:

- A small "keep-alive" message is sent every 15 seconds to prevent the TCP/IP connection from being closed.
- The session is started without checking for the completion of previous publications.
- The interval between trying to send a message again is 15 seconds.
- No last will and testament message is created for the connection.
- The standard SocketFactory is used to create the connection.

Change the connection options by creating a `ConnectionOptions` object and passing it as an additional parameter to `client.connect`.

11. Publish.

```
MqttDeliveryToken token = topic.publish(message);
```

The example sends the "Hello World" publication on the topic "MQTT Examples" to the queue manager.

- When the `publish` method returns, the message is safely transferred to the MQTT client, but not yet transferred to the server. If the message has QoS 1 or 2, then the message is stored locally, in case the client fails before delivery is completed.
- `publish` returns a delivery token, which is used to check whether an acknowledgment has been received from the server yet.

12. Wait for acknowledgment from the server.

```
token.waitForCompletion(Example.timeout);
```

The `PubSync` example waits for an acknowledgment from the server, which confirms the message has been delivered.

- Without the timeout, the client would wait indefinitely. The task, [“Creating an asynchronous publisher for MQ Telemetry Transport using Java” on page 460](#) shows how to receive acknowledgments without waiting by using a callback object.

13. Disconnect the client from the server.

```
client.disconnect();
```

The client disconnects from the server and waits for any `MqttCallback` methods that are running to finish. It then waits for up to 30 seconds to finish any remaining work. You can specify a quiesce timeout as an additional parameter.

14. Save changes to `PubSync.java` and `Example.java`

Eclipse automatically compiles the Java. You are now ready to see the results by running the program.

Results

To see the publications using WebSphere MQ, create a topic, a queue, and a durable subscription, all called "MQTTExampleTopic" using the script in [Figure 87 on page 459](#). Run the client to publish on the MQTT Examples topic, and then run the sample program **amqsbcg** to browse the publications on the MQTTExamples queue.

1. Start a queue manager, and start its telemetry (MQXR) service running. Make sure the TCP/IP address and port configured for the telemetry channel match the values you use in the MQTT application.
2. Configure a durable subscription by creating the `mqttexamples.txt` command script, and running it using **runmqsc**:

```
DEFINE TOPIC('MQTTExampleTopic') TOPICSTR('MQTT Example') REPLACE
DEFINE QLOCAL('MQTTExampleQueue') REPLACE
DEFINE SUB('MQTTExampleSub') DEST('MQTTExampleQueue') TOPICOBJ('MQTTExampleTopic') REPLACE
```

Figure 87. *mqttExampleTopic.txt*

To run the script on Windows, type the command:

```
runmqsc queue manager name < mqttExampleTopic.txt
```

3. Run the client as a Java application from within Eclipse, or by running Java in a command window:

```
java -cp jar_dir\com.ibm.micro.client.mqttv3.jar
      com.ibm.mq.id.classname.class
```

Note: The command window must be open in the directory containing the path, com\ibm\mq\id.

4. Either browse the results using WebSphere MQ Explorer, or run the command:

```
amqsbcg MQTTExampleQueue queue manager name
```

Example code

[PubSync.java](#) is a full listing of the code described in [Procedure](#). Modify the Example class in [Figure 89](#) on [page 460](#) to override the default parameters used in [PubSync.java](#).

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class PubSync {
    public static void main(String[] args) {
        try {
            MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);
            MqttTopic topic = client.getTopic(Example.topicString);
            MqttMessage message = new MqttMessage(Example.publication.getBytes());
            message.setQos(Example.QoS);
            client.connect();
            System.out.println("Waiting for up to " + Example.sleepTimeout / 1000
                + " seconds for publication of \"" + message.toString()
                + "\" with QoS = " + message.getQos());
            System.out.println("On topic \"" + topic.getName()
                + "\" for client instance: \"" + client.getClientId()
                + "\" on address " + client.getServerURI() + "\"");
            MqttDeliveryToken token = topic.publish(message);
            token.waitForCompletion(Example.sleepTimeout);
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" has been received: " + token.isComplete());
            client.disconnect();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Figure 88. *PubSync.java*

```

package com.ibm.mq.id;
import java.util.Properties;
import java.util.UUID;
public final class Example {
    public static final String TCPAddress =
        System.getProperty("TCPAddress", "tcp://localhost:1883");
    public static final String SSLAddress =
        System.getProperty("SSLAddress", "ssl://localhost:8883");
    public static final String username =
        System.getProperty("username", System.getProperty("user.name"));
    public static final char [] password =
        System.getProperty("password", "Password").toCharArray();
    public static final String clientId =
        String.format("%-23.23s", username + "_" +
            System.getProperty("clientId",
                (UUID.randomUUID().toString()).trim().replace('-', '_'));
    public static final String topicString =
        System.getProperty("topicString", "MQTT Example");
    public static final String publication =
        System.getProperty("publication", "Hello World " +
            String.format("%tc", System.currentTimeMillis()));
    public static final int quiesceTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final int sleepTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final boolean cleanSession =
        Boolean.parseBoolean(System.getProperty("cleanSession", "false"));
    public static final int QoS =
        Integer.parseInt(System.getProperty("QoS", "1"));
    public static final boolean retained =
        Boolean.parseBoolean(System.getProperty("retained", "false"));
    public static final Properties getSSLSettings() {
        final Properties properties = new Properties();
        properties.setProperty("com.ibm.ssl.keyStore",
            "C:\\IBM\\MQ\\Data\\ClientKeyStore.jks");
        properties.setProperty("com.ibm.ssl.keyStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.keyStorePassword",
            "password");
        properties.setProperty("com.ibm.ssl.trustStore",
            "C:\\IBM\\MQ\\Data\\ClientTrustStore.jks");
        properties.setProperty("com.ibm.ssl.trustStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.trustStorePassword",
            "password");
        return properties;
    }
}

```

Figure 89. Example.java

Related concepts

[MQTT publish/subscribe applications](#)

Creating an asynchronous publisher for MQ Telemetry Transport using Java

In this task, you follow a tutorial to modify your first publisher application. The modifications enable the application to send publications without waiting for delivery acknowledgments. The delivery acknowledgments are received by a callback class that you create.

Before you begin

Install the WebSphere MQ Telemetry feature on a server that has IBM WebSphere MQ Version 7.1 or later installed.

The client application uses the `com.ibm.mq.micro.client.mqttv3` package in the IBM WebSphere MQ Telemetry Software Development Toolkit (SDK). The SDK is part of the IBM WebSphere MQ Telemetry installation. The client connects to the IBM WebSphere MQ Telemetry feature to exchange messages with IBM WebSphere MQ.

You must also install the telemetry updates for IBM WebSphere MQ Explorer Version 7.1 to administer IBM WebSphere MQ Telemetry. The updates are part of the IBM WebSphere MQ Telemetry installation.

An MQTT client, running on Java SE, requires version 6.0 of Java SE, or later. IBM Java SE v6.0 is part of the IBM WebSphere MQ Version 7.1 installation. It is located at *WebSphere MQ installation directory\java\jre*

About this task

The example is a publish application, PubAsync. PubAsync publishes Hello World on the topic MQTT Examples, without waiting for confirmation that the publication has been delivered to the queue manager. The delivery acknowledgments are received in a callback class, Callback.

By setting up a durable subscription to MQTT Examples you can check that the application works.

The procedure uses Eclipse to develop, build, and run the client. You can download Eclipse from the Eclipse project website at www.eclipse.org.

The steps in Procedure modify the PubSync.java application in “Creating your first MQ Telemetry Transport publisher application using Java” on page 455.

Alternatively, you can copy the code, “Example code” on page 463, into a new directory .\com\ibm\mq\id. Create three Java files, Example.java, Callback.java, and PubAsync.java. Compile the examples using the command,

```
javac -cp jar_dir\com.mq.micro.client.mqttv3.jar
      com.ibm.mq.id.PubAsync.java com.ibm.mq.id.Callback.java com.ibm.mq.id.Example.java
```

Run PubAsync using the command,

```
java -cp jar_dir\com.mq.micro.client.mqttv3.jar
      com.ibm.mq.id.PubAsync.class
```

Procedure

1. In the com.ibm.mq.id package, create a file, Callback.java. Copy the code in Figure 92 on page 463 into the file.
Callback.java implements the MqttCallback interface. In the example, an additional constructor initializes the callback with some instance data.
2. In the com.ibm.mq.id package, right click PubSync.java and copy it. Paste it in the same package, renaming it to PubAsync.
3. Just before the client.connect(); line of code, instantiate the Callback class, passing the client identifier.

```
Callback callback = new Callback(Example.clientId);
client.setCallback(callback);
```

- The Callback class implements MqttCallback. One callback instance is required, per client identifier. In this example, the constructor passes the client identifier to save as instance data. It is used in the callback to identify which instance of the callback has been started.
- You must implement three methods in the callback class:
public void messageArrived(MqttTopic topic, MqttMessage message)
Receives a publication that has been subscribed to.
public void connectionLost(Throwable cause)
Called when the connection is lost.
public void deliveryComplete(MqttDeliveryToken token)
Called when a delivery token is received for a QoS 1 or 2 message that has been published.
- The callback is activated by MqttClient.connect.

4. Disconnect the client

- a) Remove the statement containing the `token.waitForCompletion` expression.
The main thread continues without waiting for the publication to be delivered.
- b) Test whether the client is already disconnected.
The MQTT client disconnects following an error returned to the `lostConnection` method in `MqttCallback`, or the client application might disconnect. Test to see if there is an open connection.
- c) Use the constant, `Example.quiesceTimeout`, to set the maximum time to quiesce the client.

```
if (client.isConnected())
    client.disconnect(Example.quiesceTimeout);
```

The client finishes when a combination of the following three conditions is true:

- a. The callback has been called for all messages that have been published in this session, or if the session was restarted, in previous sessions.
- b. Messages are in-flight and the quiesce interval has expired. By default the quiesce interval is 30 seconds. You can change the quiesce timeout by passing the number of milliseconds to wait as a parameter of `client.disconnect`.
- c. `client.disconnect` was called after some messages were published and queued by the client but before the messages were sent. Queued messages are not yet in-flight. If the session is restartable, the messages are resent when the session restarts.

Results

To see the publications using WebSphere MQ, create a topic, a queue, and a durable subscription, all called "MQTTExampleTopic" using the script in [Figure 90 on page 462](#). Run the client to publish on the MQTT Examples topic, and then run the sample program **amqsbcg** to browse the publications on the MQTTExamples queue.

1. Start a queue manager, and start its telemetry (MQXR) service running. Make sure the TCP/IP address and port configured for the telemetry channel match the values you use in the MQTT application.
2. Configure a durable subscription by creating the `mqttexamples.txt` command script, and running it using **runmqsc**:

```
DEFINE TOPIC('MQTTExampleTopic') TOPICSTR('MQTT Example') REPLACE
DEFINE QLOCAL('MQTTExampleQueue') REPLACE
DEFINE SUB('MQTTExampleSub') DEST('MQTTExampleQueue') TOPICOBJ('MQTTExampleTopic') REPLACE
```

Figure 90. mqttExampleTopic.txt

To run the script on Windows, type the command:

```
runmqsc queue_manager_name < mqttExampleTopic.txt
```

3. Run the client as a Java application from within Eclipse, or by running Java in a command window:

```
java -cp jar_dir\com.mq.micro.client.mqttv3.jar
      com.ibm.mq.id.classname.class
```

Note: The command window must be open in the directory containing the path, `com\ibm\mq\id`.

4. Either browse the results using WebSphere MQ Explorer, or run the command:

```
amqsbcg MQTTExampleQueue queue_manager_name
```

Example code

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class PubAsync {
    public static void main(String[] args) {
        try {
            MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);
            MqttTopic topic = client.getTopic(Example.topicString);
            MqttMessage message = new MqttMessage(Example.publication.getBytes());
            Callback callback = new Callback(Example.clientId);
            client.setCallback(callback);
            client.connect();
            System.out.println("Publishing \"" + message.toString()
                + "\" on topic \"" + topic.getName() + "\" with QoS = "
                + message.getQos());
            System.out.println("For client instance \"" + client.getClientId()
                + "\" on address " + client.getServerURI() + "\"");
            MqttDeliveryToken token = topic.publish(message);
            System.out.println("With delivery token \"" + token.hashCode()
                + "\" delivered: " + token.isComplete());
            if (client.isConnected())
                client.disconnect(Example.quiesceTimeout);
            System.out.println("Disconnected: delivery token \"" + token.hashCode()
                + "\" received: " + token.isComplete());
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Figure 91. PubAsync.java

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class Callback implements MqttCallback {
    private String instanceData = "";
    public Callback(String instance) {
        instanceData = instance;
    }
    public void messageArrived(MqttTopic topic, MqttMessage message) {
        try {
            System.out.println("Message arrived: \"" + message.toString()
                + "\" on topic \"" + topic.toString() + "\" for instance \""
                + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
    public void connectionLost(Throwable cause) {
        System.out.println("Connection lost on instance \"" + instanceData
            + "\" with cause \"" + cause.getMessage() + "\" Reason code "
            + ((MqttException)cause).getReasonCode() + "\" Cause \""
            + ((MqttException)cause).getCause() + "\"");
        cause.printStackTrace();
    }
    public void deliveryComplete(MqttDeliveryToken token) {
        try {
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" received by instance \"" + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Figure 92. Callback.java

```

package com.ibm.mq.id;
import java.util.Properties;
import java.util.UUID;
public final class Example {
    public static final String      TCPAddress =
        System.getProperty("TCPAddress", "tcp://localhost:1883");
    public static final String      SSLAddress =
        System.getProperty("SSLAddress", "ssl://localhost:8883");
    public static final String      username =
        System.getProperty("username", System.getProperty("user.name"));
    public static final char []      password =
        System.getProperty("password", "Password").toCharArray();
    public static final String      clientId =
        String.format("%-23.23s", username + "_" +
            System.getProperty("clientId",
                (UUID.randomUUID().toString()).trim().replace('-', '_'));
    public static final String      topicString =
        System.getProperty("topicString", "MQTT Example");
    public static final String      publication =
        System.getProperty("publication", "Hello World " +
            String.format("%tc", System.currentTimeMillis()));
    public static final int         quiesceTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final int         sleepTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final boolean     cleanSession =
        Boolean.parseBoolean(System.getProperty("cleanSession", "false"));
    public static final int         QoS =
        Integer.parseInt(System.getProperty("QoS", "1"));
    public static final boolean     retained =
        Boolean.parseBoolean(System.getProperty("retained", "false"));
    public static final Properties getSSLSettings() {
        final Properties properties = new Properties();
        properties.setProperty("com.ibm.ssl.keyStore",
            "C:\\IBM\\MQ\\Data\\ClientKeyStore.jks");
        properties.setProperty("com.ibm.ssl.keyStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.keyStorePassword",
            "password");
        properties.setProperty("com.ibm.ssl.trustStore",
            "C:\\IBM\\MQ\\Data\\ClientTrustStore.jks");
        properties.setProperty("com.ibm.ssl.trustStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.trustStorePassword",
            "password");
        return properties;
    }
}

```

Figure 93. Example.java

Creating a recoverable asynchronous publisher for MQ Telemetry Transport using Java

In this task, you follow a tutorial to modify your asynchronous publisher application. The modifications enable the application to complete delivery of publications that were not acknowledged the last time the client ran.

Before you begin

Install the WebSphere MQ Telemetry feature on a server that has IBM WebSphere MQ Version 7.1 or later installed.

The client application uses the `com.ibm.mq.micro.client.mqttv3` package in the IBM WebSphere MQ Telemetry Software Development Toolkit (SDK). The SDK is part of the IBM WebSphere MQ Telemetry installation. The client connects to the IBM WebSphere MQ Telemetry feature to exchange messages with IBM WebSphere MQ.

You must also install the telemetry updates for IBM WebSphere MQ Explorer Version 7.1 to administer IBM WebSphere MQ Telemetry. The updates are part of the IBM WebSphere MQ Telemetry installation.

An MQTT client, running on Java SE, requires version 6.0 of Java SE, or later. IBM Java SE v6.0 is part of the IBM WebSphere MQ Version 7.1 installation. It is located at *WebSphere MQ installation directory\java\jre*

About this task

The example is a publish application, `PubAsyncRestartable`. `PubAsyncRestartable` publishes Hello World on the topic `MQTT_Examples`, without waiting for confirmation that the publication has been delivered to the queue manager. The delivery acknowledgments are received in a callback class, `CallBack`. Any delivery tokens for publications that were not completed in a previous instance can be examined. They are also processed by the callback class.

By setting up a durable subscription to `MQTT_Examples` you can check that the application works.

The procedure uses Eclipse to develop, build, and run the client. You can download Eclipse from the Eclipse project website at www.eclipse.org.

The steps in [Procedure](#) modify the `PubAsync.java` application in “[Creating an asynchronous publisher for MQ Telemetry Transport using Java](#)” on page 460.

Alternatively, you can copy the code, “[Example code](#)” on page 467, into a new directory `.\com\ibm\mq\id`. Create three Java files, `Example.java`, `CallBack.java`, and `PubAsyncRestartable.java`. Compile the examples using the command,

```
javac -cp jar_dir\com.mq.micro.client.mqttv3.jar
      com.ibm.mq.id.PubAsyncRestartable.java com.ibm.mq.id.CallBack.java
      com.ibm.mq.id.Example.java
```

Run `PubAsyncRestartable` using the command,

```
java -cp jar_dir\com.mq.micro.client.mqttv3.jar
      com.ibm.mq.id.PubAsyncRestartable.class
```

Procedure

1. In the `com.ibm.mq.id` package, right click `PubAsync.java` and copy it. Paste it in the same package, renaming it to `PubAsyncRestartable`.
2. Create a reusable client identifier.

```
Example.clientId = String.format(
    "%-23.23s",
    (System.getProperty("user.name") + "-" + (System.getProperty(
        "clientId", "PubAsyncRestartable-"))).trim()).replace('-', '_');
```

Figure 94. Reusable client identifier

The applications in “[Creating your first MQ Telemetry Transport publisher application using Java](#)” on page 455 and “[Creating an asynchronous publisher for MQ Telemetry Transport using Java](#)” on page 460 used a new client identifier for each client connection. For a restartable publisher or subscriber, you must use the same client identifier each time the client is connected, but different clients must use different identifiers; see “[Client identifier](#)” on page 506. The reusable client identifier is built from the user name and the name of the class. It is limited to 23 bytes long. It must only have characters that are valid in queue manager object names. The code removes any hyphens that might have been inserted.

3. The QoS of the message is set to 2 rather than the default, 1, to avoid duplicate messages.

```
message.setQos(Example.QoS);
```

You need to either change the value of `Example.QoS` to 2, or pass the QoS property as an argument using the `-DQoS=2` option on the Java command line.

4. Create an `MqttConnectOptions` object, and set its `cleanSession` attribute to `false`.
 - a) Create an `MqttConnectOptions` object.

```
MqttConnectOptions conOptions = new MqttConnectOptions();
```

conOptions is an option parameter on the MqttClient constructor.

- b) Set the clearSession attribute.

```
conOptions.setCleanSession(Example.cleanSession);
```

By default, the parameter Example.cleanSession is set to true, matching the default setting of MqttConnectionOptions.cleanSession.

When PubAsyncRestartable restarts, it can start with a "clean session", and clear any pending delivery tokens for messages of QoS 1 or 2.

Set Example.cleanSession to false to keep all the pending delivery tokens. The tokens are processed by the MqttCallback class when the client is connected again.

5. If the session is being restarted, then retrieve any pending delivery tokens, and print their contents.

```
if (!conOptions.isCleanSession()) {
    MqttDeliveryToken tokens[] = client.getPendingDeliveryTokens();
    System.out.println("Starting a previous session for instance \"
        + client.getClientId() + \" with \" + tokens.length
        + \" delivery tokens pending");
    for (int i = 0; i < tokens.length; i++) {
        System.out.println("Message \" + tokens[i].getMessage().toString()
            + \" with QoS=\" + tokens[i].getMessage().getQos()
            + \" recovered by instance \" + client.getClientId()
            + \" and assigned delivery token \" + tokens[i].hashCode()
            + \"");
    }
} else
    System.out.println("Starting a clean session for instance \"
        + client.getClientId());
```

6. Pass the conOptions parameter to the MqttClient constructor.

```
client.connect(conOptions);
```

7. On disconnecting, set a maximum disconnect interval.

```
client.disconnect(Example.timeout);
```

To be able to show pending delivery tokens being processed, a previous instance must finish without completing delivery. To run the example with the possibility of not acknowledging publications before PubAsyncRestartable finishes, set Example.timeout to 0.

Results

To see the publications using WebSphere MQ, create a topic, a queue, and a durable subscription, all called "MQTTExampleTopic" using the script in Figure 95 on page 466. Run the client to publish on the MQTT Examples topic, and then run the sample program **amqsbcg** to browse the publications on the MQTTExamples queue.

1. Start a queue manager, and start its telemetry (MQXR) service running. Make sure the TCP/IP address and port configured for the telemetry channel match the values you use in the MQTT application.
2. Configure a durable subscription by creating the mqttexamples.txt command script, and running it using **runmqsc**:

```
DEFINE TOPIC('MQTTExampleTopic') TOPICSTR('MQTT Example') REPLACE
DEFINE QLOCAL('MQTTExampleQueue') REPLACE
DEFINE SUB('MQTTExampleSub') DEST('MQTTExampleQueue') TOPICOBJ('MQTTExampleTopic') REPLACE
```

Figure 95. mqttExampleTopic.txt

To run the script on Windows, type the command:

```
runmqsc queue manager name < mqttExampleTopic.txt
```

3. Run the client as a Java application from within Eclipse, or by running Java in a command window:

```
java -cp jar_dir\com.ibm.mq.micro.client.mqttv3.jar  
com.ibm.mq.id.classname.class
```

Note: The command window must be open in the directory containing the path, com\ibm\mq\id.

4. Either browse the results using WebSphere MQ Explorer, or run the command:

```
amqsbcg MQTTEExampleQueue queue manager name
```

Example code

```
package com.ibm.mq.id;  
import com.ibm.micro.client.mqttv3.MqttClient;  
import com.ibm.micro.client.mqttv3.MqttConnectOptions;  
import com.ibm.micro.client.mqttv3.MqttDeliveryToken;  
import com.ibm.micro.client.mqttv3.MqttMessage;  
import com.ibm.micro.client.mqttv3.MqttTopic;  
public class PubAsyncRestartable {  
    public static void main(String[] args) {  
        Example.clientId = String.format(  
            "%-23.23s",  
            (System.getProperty("user.name") + "_" + (System.getProperty(  
                "clientId", "PubAsyncRestartable-"))).trim()).replace('-', '_');  
        try {  
            MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);  
            MqttTopic topic = client.getTopic(Example.topicString);  
            MqttMessage message = new MqttMessage(Example.publication.getBytes());  
            message.setQos(Example.QoS);  
            Callback callback = new Callback(Example.clientId);  
            client.setCallback(callback);  
            MqttConnectOptions conOptions = new MqttConnectOptions();  
            conOptions.setCleanSession(Example.cleanSession);  
            if (!conOptions.isCleanSession()) {  
                MqttDeliveryToken tokens[] = client.getPendingDeliveryTokens();  
                System.out.println("Starting a previous session for instance \"\"  
                    + client.getClientId() + "\" with " + tokens.length  
                    + " delivery tokens pending");  
                for (int i = 0; i < tokens.length; i++) {  
                    System.out.println("Message \"\" + tokens[i].getMessage().toString()  
                        + "\" with QoS=" + tokens[i].getMessage().getQos()  
                        + " recovered by instance \"\" + client.getClientId()  
                        + "\" and assigned delivery token \"\" + tokens[i].hashCode()  
                        + "\"");  
                }  
            } else  
                System.out.println("Starting a clean session for instance \"\"  
                    + client.getClientId() + "\"");  
            client.connect(conOptions);  
            System.out.println("Publishing \"\" + message.toString()  
                + "\" on topic \"\" + topic.getName() + "\" with QoS = "  
                + message.getQos());  
            System.out.println("For client instance \"\" + client.getClientId()  
                + "\" on address " + client.getServerURI() + "\"");  
            MqttDeliveryToken token = topic.publish(message);  
            System.out.println("With delivery token \"\" + token.hashCode()  
                + " delivered: " + token.isComplete());  
            if (client.isConnected())  
                client.disconnect(Example.quiesceTimeout);  
            System.out.println("Disconnected: delivery token \"\" + token.hashCode()  
                + "\" has been received: " + token.isComplete());  
        } catch (Exception e) {  
            e.printStackTrace();  
        }  
    }  
}
```

Figure 96. PubAsyncRestartable.java

```

package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class CallBack implements MqttCallback {
    private String instanceData = "";
    public CallBack(String instance) {
        instanceData = instance;
    }
    public void messageArrived(MqttTopic topic, MqttMessage message) {
        try {
            System.out.println("Message arrived: \"" + message.toString()
                + "\" on topic \"" + topic.toString() + "\" for instance \""
                + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
    public void connectionLost(Throwable cause) {
        System.out.println("Connection lost on instance \"" + instanceData
            + "\" with cause \"" + cause.getMessage() + "\" Reason code "
            + ((MqttException)cause).getReasonCode() + "\" Cause \""
            + ((MqttException)cause).getCause() + "\"");
        cause.printStackTrace();
    }
    public void deliveryComplete(MqttDeliveryToken token) {
        try {
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" received by instance \"" + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}

```

Figure 97. CallBack.java

```

package com.ibm.mq.id;
import java.util.Properties;
import java.util.UUID;
public final class Example {
    public static final String TCPAddress =
        System.getProperty("TCPAddress", "tcp://localhost:1883");
    public static final String SSLAddress =
        System.getProperty("SSLAddress", "ssl://localhost:8883");
    public static final String username =
        System.getProperty("username", System.getProperty("user.name"));
    public static final char [] password =
        System.getProperty("password", "Password").toCharArray();
    public static final String clientId =
        String.format("%-23.23s", username + "_" +
            System.getProperty("clientId",
                (UUID.randomUUID().toString()).trim().replace('-', '_'));
    public static final String topicString =
        System.getProperty("topicString", "MQTT Example");
    public static final String publication =
        System.getProperty("publication", "Hello World " +
            String.format("%tc", System.currentTimeMillis()));
    public static final int quiesceTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final int sleepTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final boolean cleanSession =
        Boolean.parseBoolean(System.getProperty("cleanSession", "false"));
    public static final int QoS =
        Integer.parseInt(System.getProperty("QoS", "1"));
    public static final boolean retained =
        Boolean.parseBoolean(System.getProperty("retained", "false"));
    public static final Properties getSSLSettings() {
        final Properties properties = new Properties();
        properties.setProperty("com.ibm.ssl.keyStore",
            "C:\\IBM\\MQ\\Data\\ClientKeyStore.jks");
        properties.setProperty("com.ibm.ssl.keyStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.keyStorePassword",
            "password");
        properties.setProperty("com.ibm.ssl.trustStore",
            "C:\\IBM\\MQ\\Data\\ClientTrustStore.jks");
        properties.setProperty("com.ibm.ssl.trustStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.trustStorePassword",
            "password");
        return properties;
    }
}

```

Figure 98. Example.java

Creating a subscriber for MQ Telemetry Transport using Java

In this task, you follow a tutorial to create a subscriber application. The subscriber creates a subscription to a topic and receives publications for the subscription.

Before you begin

Install the WebSphere MQ Telemetry feature on a server that has IBM WebSphere MQ Version 7.1 or later installed.

The client application uses the `com.ibm.mq.micro.client.mqttv3` package in the IBM WebSphere MQ Telemetry Software Development Toolkit (SDK). The SDK is part of the IBM WebSphere MQ Telemetry installation. The client connects to the IBM WebSphere MQ Telemetry feature to exchange messages with IBM WebSphere MQ.

You must also install the telemetry updates for IBM WebSphere MQ Explorer Version 7.1 to administer IBM WebSphere MQ Telemetry. The updates are part of the IBM WebSphere MQ Telemetry installation.

An MQTT client, running on Java SE, requires version 6.0 of Java SE, or later. IBM Java SE v6.0 is part of the IBM WebSphere MQ Version 7.1 installation. It is located at *WebSphere MQ installation directory\java\jre*

About this task

The example is a subscriber application, `Subscribe`. `Subscribe` creates a subscription topic, MQTT Examples, and waits for publications on the subscription for 30 seconds.

A subscriber can create subscription and wait for publications. It can also receive publications sent to a subscription created previously, for the same client identifier. The `MqttConnectionOptions.cleanSession` Boolean attribute controls whether publications sent previously are received or not; see “Subscriptions” on page 517.

You can use the publish example programs to create publications, or use WebSphere MQ explorer to create a test publication on the MQTT Examples topic.

The procedure uses Eclipse to develop, build, and run the client. You can download Eclipse from the Eclipse project website at www.eclipse.org.

The instructions in [Procedure](#) assume that you have already created the `com.ibm.mq.id` package in one of the earlier tasks, and copied in the `Example.java` and `Callback.java` classes.

Procedure

1. Create the class, `Subscribe` in the package `com.ibm.mq.id`.
2. Create a reusable client identifier.

```
Example.clientId = String.format(
    "%-23.23s",
    (System.getProperty("user.name") + " " + (System.getProperty(
        "clientId", "Subscribe."))).trim()).replace('-', '_');
```

Figure 99. Reusable client identifier

The applications in “Creating your first MQ Telemetry Transport publisher application using Java” on page 455 and “Creating an asynchronous publisher for MQ Telemetry Transport using Java” on page 460 used a new client identifier for each client connection. For a restartable publisher or subscriber, you must use the same client identifier each time the client is connected, but different clients must use different identifiers; see “Client identifier” on page 506. The reusable client identifier is built from the user name and the name of the class. It is limited to 23 bytes long. It must only have characters that are valid in queue manager object names. The code removes any hyphens that might have been inserted.

3. Create a try-catch block.

```
try { ...
} catch (Exception e) {
    e.printStackTrace();
}
```

The MQTT client throws `MqttException`, `MqttPersistenceException` or `MqttSecurityException`. `MqttPersistenceException` and `MqttSecurityException` are subclasses of `MqttException`.

Use the `MqttException.getReasonCode` method to find out the reason for the exception. If a `MqttPersistenceException` or `MqttSecurityException` is thrown, use the `getCause` method to return the underlying throwable exception.

4. Create a new `MqttClient` instance.

```
MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);
```

Provide the client with a server address, which is used later to connect to WebSphere MQ. Set the client identifier to name the client.

- Optionally, you can provide an implementation of the `MqttClientPersistence` interface to replace the default implementation. The default `MqttPersistence` implementation stores QoS 1 and 2 messages awaiting delivery as files; see [“Message persistence in MQTT clients”](#) on page 510.
- The default IBM WebSphere MQ TCP/IP port for MQTT is 1883. For SSL, it is 8883. In the example, the default address is set to `tcp://localhost:1883`.
- Typically, it is important to be able to identify a specific physical client using the client identifier. The client identifier must be unique across all clients connect to a server; see [“Client identifier”](#) on page 506. Using the same client identifier as a previous instance indicates the present instance is an instance of the same client. If you duplicate a client identifier in two running clients, an exception is thrown in both clients, and one client terminates.
- The length of the client identifier is limited to 23 bytes. An exception is thrown, if the length is exceeded. The client identifier must contain only characters permitted in a queue manager name; for example, no hyphens or spaces.
- Until you call the `MqttClient.connect` method, no message processing takes place.

Use the client object to publish and subscribe topics and recover information about publications that have not been delivered yet.

5. Just before the `client.connect()`; line of code, instantiate the `Callback` class, passing the client identifier.

```
Callback callback = new Callback(Example.clientId);
client.setCallback(callback);
```

- The `Callback` class implements `MqttCallback`. One callback instance is required, per client identifier. In this example, the constructor passes the client identifier to save as instance data. It is used in the callback to identify which instance of the callback has been started.
- You must implement three methods in the callback class:

public void messageArrived(MqttTopic topic, MqttMessage message)

Receives a publication that has been subscribed to.

public void connectionLost(Throwable cause)

Called when the connection is lost.

public void deliveryComplete(MqttDeliveryToken token)

Called when a delivery token is received for a QoS 1 or 2 message that has been published.

- The callback is activated by `MqttClient.connect`.
6. Create an `MqttConnectOptions` object, and set its `cleanSession` attribute.
 - a) Create an `MqttConnectOptions` object.

```
MqttConnectOptions conOptions = new MqttConnectOptions();
```

`conOptions` is an option parameter on the `MqttClient` constructor.

- b) Set the `cleanSession` attribute.

```
conOptions.setCleanSession(Example.cleanSession);
```

By default, the parameter `Example.cleanSession` is set to `true`, matching the default setting of `MqttConnectOptions.cleanSession`.

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, any old subscriptions for the client are removed when the client connects. Any new subscriptions the client makes during the session are removed when it disconnects.

If you set `MqttConnectOptions.cleanSession` to `false` before connecting, any subscriptions the client creates are added to all the subscriptions that existed for the client before it connected. All the subscriptions remain active when the client disconnects.

Another way of understanding the way the `cleanSession` attribute affects subscriptions is to think of it as a modal attribute. In its default mode, `cleanSession=true`, the client creates subscriptions and receives publications only within the scope of the session. In the alternative mode, `cleanSession=false`, subscriptions are durable. The client can connect and disconnect and its subscriptions remain active. When the client reconnects, it receives any undelivered publications. While it is connected, it can modify the set of subscriptions that are active on its behalf.

You must set the `cleanSession` mode before connecting; the mode lasts for the whole session. To change its setting, you must disconnect and reconnect the client. If you change modes from using `cleanSession=false` to `cleanSession=true`, all previous subscriptions for the client, and any publications that have not been received, are discarded.

7. Pass the `conOptions` parameter to the `MqttClient` constructor.

```
client.connect(conOptions);
```

8. Create a subscription.

```
client.subscribe(Example.topicString, Example.QoS);
```

The example uses a `MqttClient.subscribe` method which passes one topic filter with a QoS option. `MqttClient.subscribe` method has four signatures, and you can pass arrays of subscription filters as well as a single filter.

The example uses topic string used by the publish examples as a topic filter, so it receives any publications they create.

Each time you run the example, `subscribe.java`, it creates a subscription. If you do not change `Example.topicString`, it recreates the same subscription again. If a subscription is recreated it does not result in two identical subscriptions. A client does not receive duplicate copies of publications which match an identical subscription.

Subscriptions are described in [“Subscriptions” on page 517](#), and filters in [“Topic strings and topic filters in MQTT clients” on page 519](#).

9. Wait for some publications to arrive and then disconnect the client.

```
Thread.sleep(Example.sleepTimeout);  
client.disconnect();
```

Publications are received by the implementation of the `MqttCallback.messageArrived` method.

The subscribe application has not published any messages and so it does not wait for any delivery tokens. `client.disconnect` takes place without any delay.

Example code

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.MqttClient;
import com.ibm.micro.client.mqttv3.MqttConnectOptions;
public class Subscribe {
    public static void main(String[] args) {
        Example.clientId = String.format(
            "%-23.23s",
            (System.getProperty("user.name") + "_" + System.getProperty("clientId",
                "Subscribe.")).trim());
        try {
            MqttClient client = new MqttClient(Example.TCPAddress, Example.clientId);
            Callback callback = new Callback(Example.clientId);
            client.setCallback(callback);
            MqttConnectOptions conOptions = new MqttConnectOptions();
            conOptions.setCleanSession(Example.cleanSession);
            client.connect(conOptions);
            System.out.println("Subscribing to topic \"" + Example.topicString
                + "\" for client instance \"" + client.getClientId()
                + "\" using QoS " + Example.QoS + ". Clean session is "
                + Example.cleanSession);
            client.subscribe(Example.topicString, Example.QoS);
            System.out.println("Going to sleep for " + Example.sleepTimeout / 1000
                + " seconds");
            Thread.sleep(Example.sleepTimeout);
            client.disconnect();
            System.out.println("Finished");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Figure 100. *Subscribe.java*

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class Callback implements MqttCallback {
    private String instanceData = "";
    public Callback(String instance) {
        instanceData = instance;
    }
    public void messageArrived(MqttTopic topic, MqttMessage message) {
        try {
            System.out.println("Message arrived: \"" + message.toString()
                + "\" on topic \"" + topic.toString() + "\" for instance \""
                + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
    public void connectionLost(Throwable cause) {
        System.out.println("Connection lost on instance \"" + instanceData
            + "\" with cause \"" + cause.getMessage() + "\" Reason code "
            + ((MqttException)cause).getReasonCode() + "\" Cause \""
            + ((MqttException)cause).getCause() + "\"");
        cause.printStackTrace();
    }
    public void deliveryComplete(MqttDeliveryToken token) {
        try {
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" received by instance \"" + instanceData + "\"");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

Figure 101. *Callback.java*

```

package com.ibm.mq.id;
import java.util.Properties;
import java.util.UUID;
public final class Example {
    public static final String      TCPAddress =
        System.getProperty("TCPAddress", "tcp://localhost:1883");
    public static final String      SSLAddress =
        System.getProperty("SSLAddress", "ssl://localhost:8883");
    public static final String      username =
        System.getProperty("username", System.getProperty("user.name"));
    public static final char []      password =
        System.getProperty("password", "Password").toCharArray();
    public static final String      clientId =
        String.format("%-23.23s", username + "_" +
            System.getProperty("clientId",
                (UUID.randomUUID().toString()).trim().replace('-', '_'));
    public static final String      topicString =
        System.getProperty("topicString", "MQTT Example");
    public static final String      publication =
        System.getProperty("publication", "Hello World " +
            String.format("%tc", System.currentTimeMillis()));
    public static final int         quiesceTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final int         sleepTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final boolean     cleanSession =
        Boolean.parseBoolean(System.getProperty("cleanSession", "false"));
    public static final int         QoS =
        Integer.parseInt(System.getProperty("QoS", "1"));
    public static final boolean     retained =
        Boolean.parseBoolean(System.getProperty("retained", "false"));
    public static final Properties getSSLSettings() {
        final Properties properties = new Properties();
        properties.setProperty("com.ibm.ssl.keyStore",
            "C:\\IBM\\MQ\\Data\\ClientKeyStore.jks");
        properties.setProperty("com.ibm.ssl.keyStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.keyStorePassword",
            "password");
        properties.setProperty("com.ibm.ssl.trustStore",
            "C:\\IBM\\MQ\\Data\\ClientTrustStore.jks");
        properties.setProperty("com.ibm.ssl.trustStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.trustStorePassword",
            "password");
        return properties;
    }
}

```

Figure 102. Example.java

Related concepts

[MQTT publish/subscribe applications](#)

Authenticating an MQTT Java client using JAAS

Learn how to authenticate a client using JAAS. Modify the sample program `JAASLoginModule.java` and the example Java program `PubSync.java`. Configure a telemetry channel to require JAAS authentication, and run the modified publisher, checking its username and password using JAAS.

Before you begin

You are assumed to have installed the MQTT v3 client jar files, Javadoc, Eclipse, configured telemetry channels and coded and run `PubSync.java` before performing this task. You have an Eclipse workspace that includes a running version of `PubSync.java`.

The task is written for Windows. Change the directory paths for Linux.

About this task

The task is based on modifying the sample JAASLoginModule class in *WMQ Installation directory\mqxr\samples\JAASLoginModule.java* to create *MyLogin.java*. In the task you also modify the example code, *PubSync.java* in “Creating your first MQ Telemetry Transport publisher application using Java” on page 455 to set a username and password. As a test, *MyLogin.java* randomly accepts or rejects username and password.

The steps in the task are written as a programming exercise. You must adapt the procedure to perform real authentication in a production environment.

In a typical explanation of how to program JAAS authentication, it is assumed that the login module is authenticating the context that loaded JAAS. When the telemetry (MQXR) service calls JAAS, the context that loaded JAAS is the telemetry (MQXR) service. There is no point in authenticating the telemetry (MQXR) service context; it is always mqm. Instead, the telemetry (MQXR) service sets up the client username and password to be available to the login module class. The username and password are passed to the login module using two callbacks.

```
javax.security.auth.callback.Callback[] callbacks =
    new javax.security.auth.callback.Callback[2];
callbacks[0] =
    new javax.security.auth.callback.NameCallback("NameCallback");
callbacks[1] =
    new javax.security.auth.callback.PasswordCallback("PasswordCallback", false);
callbackHandler.handle(callbacks);
String username =
    ((javax.security.auth.callback.NameCallback) callbacks[0]).getName();
char[] password =
    ((javax.security.auth.callback.PasswordCallback) callbacks[1]).getPassword();
```

The username and password of the client are the only information about the client that is available to the login module.

Procedure

1. Create two packages, *samples* and *security.jaas* in the same Java project as *PubSync.java*.

The *samples* package is used only for reference. Make the code changes in the *security.jaas* package.

2. Import *JAASLoginModule.java* and *JAASPrincipal.java* into both packages.

If necessary, refactor the package statements in the Java source to eliminate compilation errors.

3. Refactor the class name, *JAASLoginModule*, in the *security.jaas* package to *MyLogin*
4. In *MyLogin.java*, replace some of the code in the *login* method to show the module working.

a) Replace the code:

```
// Accept everything.
if (true)
    loggedIn = true;
else
    throw new javax.security.auth.login.FailedLoginException("Login failed");
```

b) With the code:

```
// login half the users randomly
PrintWriter pw = new PrintWriter(new FileWriter(System.getProperty("user.dir")
    + "\\MyLogin.log", true));
pw.println("Called JAASLogin.login at "
    + System.getProperty("publication", "Hello World "
    + String.format("%tc", System.currentTimeMillis())));
if (Math.random() < 0.5)
    loggedIn = true;
pw.println("Username: \"\" + username + "\", Password: \"\"
    + String.valueOf(password) + "\" loggedIn: \" + loggedIn);
pw.close();
if (!loggedIn)
    throw new javax.security.auth.login.FailedLoginException("Login failed");
principal= new JAASPrincipal(username);
```

The complete source for `MyLogin.java` is in [Figure 105 on page 478](#). The source for `JAASPrincipal.java`, with the package name refactored to `security.jaas` is in [Figure 106 on page 479](#).

5. Set the class path in `service.env` to point to the directory containing the path to `security/jaas/MyLogin.class` and `security/jaas/JAASPrincipal.class`.

```
CLASSPATH=C:\WMQTelemetryApps\MQTTSecureExamples\bin
```

See [Telemetry channel JAAS configuration](#) for information about using `service.env` to pass a class path to a WebSphere MQ service.

6. Add a login module stanza to `jaas.config`.

```
MyLoginExample {  
    security.jaas.MyLogin required debug=true;  
};
```

See [Telemetry channel JAAS configuration](#) for information about using `jaas.config` to define a JAAS login module.

7. Add a telemetry channel using the **New telemetry channel** wizard in WebSphere MQ Explorer, configuring the channel to require JAAS authentication. Refer it to the `MyLoginExample` stanza.

For example, adapt the information you type into the wizard from this stanza in the `mqxr_win.properties` file. If you are working in Linux the file is called `mqxr_unix.properties`. Do not edit the telemetry properties file directly; use the wizard.

```
com.ibm.mq.MQXR.channel/JAASMCUser: \  
com.ibm.mq.MQXR.Port=1884;\  
com.ibm.mq.MQXR.JAASConfig=MyLoginExample;\br/>com.ibm.mq.MQXR.UserName=Admin;\br/>com.ibm.mq.MQXR.StartWithMQXRService=true
```

Note: If you modify any of the telemetry channel parameters, or modify the `security.jaas.MyLogin` class, you must stop and restart the telemetry (MQXR) service. Only when you restart the service do the changes take effect.

8. Make a copy of `PubSync.java` in the `com.ibm.mq.id` package, and name the copy `PubSyncJAAS.java`.

See [“Creating your first MQ Telemetry Transport publisher application using Java” on page 455](#) for the steps to create `PubSyncJAAS.java` in the `com.ibm.mq.id` package.

9. Set the `MqttConnectOptions.username` and `MqttConnectOptions.password` in `PubSyncJAAS.java` program, and pass `MqttConnectOptions` as a parameter of `MqttClient.connect`.

```
MqttConnectOptions conOptions = new MqttConnectOptions();  
conOptions.setUserName(Example.username);  
conOptions.setPassword(Example.password);  
client.connect(conOptions);
```

See the italicized code in `PubSyncJAAS.java` using the constants set in `Example.java`.

10. Set `Example.TCPAddress` to the socket address of the telemetry channel you have configured to use the JAAS configuration, `MyLoginExample`. For example, use 1884 as the port number.
11. Run `PubSyncJAAS` a number of times to see the client logging in, and being accepted or rejected.

An exception is thrown each time the login attempt is rejected.

Results

[Figure 103 on page 477](#) shows the results of running `PubSyncJAAS.java` twice. The log records are shown in [Figure 104 on page 477](#).

```
Waiting for up to 10 seconds for publication of "Hello World Fri Jun 04 08:31:05 BST 2010" with
QoS = 1
On topic "MQTT Example" for client instance: "Admin_61c57a18_4bf7_40d" on address tcp://
localhost:1884"
With username "Admin" and password "Password"
Client exception caught
Client is not connected (32104)
    at
com.ibm.micro.client.mqttv3.internal.ExceptionHelper.createMqttException(ExceptionHelper.java:33
)
    at com.ibm.micro.client.mqttv3.internal.ClientComms.internalSend(ClientComms.java:88)
    at com.ibm.micro.client.mqttv3.internal.ClientComms.sendNowait(ClientComms.java:105)
    at com.ibm.micro.client.mqttv3.MqttTopic.publish(MqttTopic.java:68)
    at com.ibm.mq.id.PubSync.main(PubSync.java:24)
```

```
Waiting for up to 10 seconds for publication of "Hello World Fri Jun 04 08:31:40 BST 2010" with
QoS = 1
On topic "MQTT Example" for client instance: "Admin_1d1599a0_50f5_4ea" on address tcp://
localhost:1884"
With username "Admin" and password "Password"
Delivery token "1731749688" has been received: true
```

Figure 103. Console output from PubSyncJAAS.java

The log file `MyLogin.log` is stored in *WMQ Data directory*; for example,
`C:\IBM\MQ\Data\MyLogin.log`:

```
Called JAASLogin.login at Hello World Fri Jun 04 08:31:05 BST 2010
Username: "Admin", Password: "Password" loggedIn: false
Called JAASLogin.login at Hello World Fri Jun 04 08:31:40 BST 2010
Username: "Admin", Password: "Password" loggedIn: true
```

Figure 104. MyLogin.log

Examples

The italicized code in [Figure 105 on page 478](#) is the modification to the sample `JAASLoginModule.java`.

```

package security.jaas;
import java.io.FileWriter;
import java.io.PrintWriter;

public class JAASLogin implements javax.security.auth.spi.LoginModule {
    private javax.security.auth.Subject subject;
    private javax.security.auth.callback.CallbackHandler callbackHandler;
    JAASPrincipal principal;
    boolean loggedIn = false;
    public void initialize(javax.security.auth.Subject subject,
        javax.security.auth.callback.CallbackHandler callbackHandler,
        java.util.Map<String, ?> sharedState, java.util.Map<String, ?> options) {
        this.subject = subject;
        this.callbackHandler = callbackHandler;
    }
    public boolean login() throws javax.security.auth.login.LoginException {
        try {
            javax.security.auth.callback.Callback[] callbacks = new
javax.security.auth.callback.Callback[2];
            callbacks[0] = new javax.security.auth.callback.NameCallback(
                "NameCallback");
            callbacks[1] = new javax.security.auth.callback.PasswordCallback(
                "PasswordCallback", false);

            callbackHandler.handle(callbacks);
            String username = ((javax.security.auth.callback.NameCallback) callbacks[0])
                .getName();
            char[] password = ((javax.security.auth.callback.PasswordCallback) callbacks[1])
                .getPassword();
            // login half the users randomly
            PrintWriter pw = new PrintWriter(new FileWriter(System
                .getProperty("user.dir")
                + "\\mylogin.log", true));
            pw.println("Called JAASLogin.login at "
                + System.getProperty("publication", "Hello World "
                    + String.format("%tc", System.currentTimeMillis())));
            if (Math.random() < 0.5)
                loggedIn = true;
            pw.println("Username: \"" + username + "\", Password: \""
                + String.valueOf(password) + "\" loggedIn: " + loggedIn);
            pw.close();
            if (!loggedIn)
                throw new javax.security.auth.login.FailedLoginException("Login failed");
            principal = new JAASPrincipal(username);
        } catch (java.io.IOException exception) {
            throw new javax.security.auth.login.LoginException(exception.toString());
        } catch (javax.security.auth.callback.UnsupportedCallbackException exception) {
            throw new javax.security.auth.login.LoginException(exception.toString());
        }
        return loggedIn;
    }
    public boolean abort() throws javax.security.auth.login.LoginException {
        logout();
        return true;
    }
    public boolean commit() throws javax.security.auth.login.LoginException {
        if (loggedIn) {
            if (!subject.getPrincipals().contains(principal))
                subject.getPrincipals().add(principal);
        }
        return true;
    }
    public boolean logout() throws javax.security.auth.login.LoginException {
        subject.getPrincipals().remove(principal);
        principal = null;
        loggedIn = false;
        return true;
    }
}

```

Figure 105. *MyLogin.java*

Figure 106 on page 479 is the sample code `JAASLoginPrincipal.java`, copied into the package `security.jaas`. The purpose of `JAASLoginPrincipal` is to implement the

java.security.Principal interface to keep a record of the users that have been successfully logged in by MyLogin.

```
package security.jaas;
public class JAASPrincipal implements java.security.Principal,
    java.io.Serializable {
    private static final long serialVersionUID = 1L;
    String name;
    public JAASPrincipal(String name) {
        this.name = name;
    }
    public String getName() {
        return name;
    }
    public String toString() {
        return (name);
    }
    public boolean equals(Object object) {
        if (object != null && object instanceof JAASPrincipal
            && name.equals(((JAASPrincipal) object).getName()))
            return true;
        else
            return false;
    }
    public int hashCode() {
        return name.hashCode();
    }
}
```

Figure 106. JAASLoginPrincipal.java

The code in [PubSync.java](#) that is modified to add a username and password is italicized in [Figure 107](#) on page 479.

```
package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class PubSyncSSL {
    public static void main(String[] args) {
        try {
            MqttClient client = new MqttClient(Example.SSLAddress, Example.clientId);
            MqttTopic topic = client.getTopic(Example.topicString);
            MqttMessage message = new MqttMessage(Example.publication.getBytes());
            message.setQos(Example.QoS);
            MqttConnectOptions conOptions = new MqttConnectOptions();
            conOptions.setUserName(Example.username);
            conOptions.setPassword(Example.password);
            client.connect(conOptions);
            System.out.println("Waiting for up to " + Example.sleepTimeout / 1000
                + " seconds for publication of \"" + message.toString()
                + "\" with QoS = " + message.getQos());
            System.out.println("On topic \"" + topic.getName() + "\" for client instance: \""
                + client.getClientId() + "\" on address " + client.getServerURI() + "\"");
            System.out.println("With username \"" + conOptions.getUserName()
                + "\" and password \"" + String.valueOf(conOptions.getPassword()) + "\"");
            MqttDeliveryToken token = topic.publish(message);
            token.waitForCompletion(Example.sleepTimeout);
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" has been received: " + token.isComplete());
            client.disconnect();
        } catch (Exception e) {
            System.out.println("Client exception caught");
            e.printStackTrace();
        }
    }
}
```

Figure 107. PubSyncJAAS.java

Modify the constants in [Example.java](#) to match your configuration. Ignore the SSL settings for this example.

```

package com.ibm.mq.id;
import java.util.Properties;
import java.util.UUID;
public final class Example {
    public static final String      TCPAddress =
        System.getProperty("TCPAddress", "tcp://localhost:1883");
    public static final String      SSLAddress =
        System.getProperty("SSLAddress", "ssl://localhost:8883");
    public static final String      username =
        System.getProperty("username", System.getProperty("user.name"));
    public static final char []     password =
        System.getProperty("password", "Password").toCharArray();
    public static final String      clientId =
        String.format("%-23.23s", username + "_" +
            System.getProperty("clientId",
                (UUID.randomUUID().toString()).trim().replace('-', '_'));
    public static final String      topicString =
        System.getProperty("topicString", "MQTT Example");
    public static final String      publication =
        System.getProperty("publication", "Hello World " +
            String.format("%tc", System.currentTimeMillis()));
    public static final int         quiesceTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final int         sleepTimeout =
        Integer.parseInt(System.getProperty("timeout", "10000"));
    public static final boolean     cleanSession =
        Boolean.parseBoolean(System.getProperty("cleanSession", "false"));
    public static final int         QoS =
        Integer.parseInt(System.getProperty("QoS", "1"));
    public static final boolean     retained =
        Boolean.parseBoolean(System.getProperty("retained", "false"));
    public static final Properties getSSLSettings() {
        final Properties properties = new Properties();
        properties.setProperty("com.ibm.ssl.keyStore",
            "C:\\IBM\\MQ\\Data\\ClientKeyStore.jks");
        properties.setProperty("com.ibm.ssl.keyStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.keyStorePassword",
            "password");
        properties.setProperty("com.ibm.ssl.trustStore",
            "C:\\IBM\\MQ\\Data\\ClientTrustStore.jks");
        properties.setProperty("com.ibm.ssl.trustStoreType",
            "JKS");
        properties.setProperty("com.ibm.ssl.trustStorePassword",
            "password");
        return properties;
    }
}

```

Figure 108. Example.java

Authenticating an SSL telemetry connection using self-signed certificates

Use self-signed certificates generated using **Keytool** to authenticate an SSL connection. You have the option of authenticating the telemetry channel, or the telemetry channel and the clients that attach to it. Messages flowing on the connection are encrypted.

Before you begin

Do the task, [“Creating your first MQ Telemetry Transport publisher application using Java”](#) on page 455 before you start, to get `PubSync.java` working with an unsecured TCP/IP connection. In this task, you modify `PubSync.java` to work with an SSL connection.

About this task

The steps in the task are written as a programming exercise. You must adapt the procedure to perform real authentication in a production environment.

The task is written for Windows. Change the directory paths for Linux.

Procedure

1. Do the task, “[Modifying PubSync.java to use SSL](#)” on page 481, to modify [PubSync.java](#) to use SSL.
2. Configure the telemetry channel, and create the keystores to use SSL.

Authenticate just the telemetry channel, or the channel and clients that connect to it:

- Do the task, “[Authenticating the telemetry channel](#)” on page 482, to connect with SSL, authenticating the telemetry channel.
 - Do the task, “[Authenticating the telemetry channel and clients](#)” on page 483, to connect with SSL, authenticating the telemetry channel and clients that connect to it.
3. Stop and restart the telemetry (MQXR) service to pick up changes to telemetry channel configurations.
 4. Run the client program to see if the configuration works.

Modifying PubSync.java to use SSL

Modify the first publisher program example to connect to a telemetry channel using SSL. Set the SSL properties used by the modified program.

Before you begin

You are assumed to have installed the MQTT v3 client jar files, Javadoc, Eclipse, configured telemetry channels and coded and run [PubSync.java](#) before performing this task. You have an Eclipse workspace that includes a running version of [PubSync.java](#).

About this task

The task uses the publisher client, [PubSync.java](#), you created in “[Creating your first MQ Telemetry Transport publisher application using Java](#)” on page 455 as a base. Only small modifications are necessary to use SSL; see [Figure 109 on page 482](#) and [Figure 110 on page 482](#).

Procedure

1. Make a copy of [PubSync.java](#) in the `com.ibm.mq.id` package, and name the copy [PubSyncSSL.java](#).
See “[Creating your first MQ Telemetry Transport publisher application using Java](#)” on page 455 for the steps to create [PubSync.java](#) in the `com.ibm.mq.id` package.
2. Set `Example.SSLAddress` to the socket address of the telemetry channel you have configured to use for the SSL configuration.
3. Change the socket address parameter of the client constructor to use `Example.SSLAddress`.

```
MqttClient client = new MqttClient(Example.SSLAddress, Example.clientId);
```

4. Set the `MqttConnectOptions.SSLProperties` in [PubSyncSSL.java](#), and pass `MqttConnectOptions` as a parameter of `MqttClient.connect`.

```
MqttConnectOptions conOptions = new MqttConnectOptions();  
conOptions.setSSLProperties(Example.getSSLSettings());  
client.connect(conOptions);
```

See the italicized code [PubSyncSSL.java](#) using the constants set in [Example.java](#).

Examples

The modifications to [PubSync.java](#) to add SSL are shown in [Figure 109 on page 482](#) in italics.

```

package com.ibm.mq.id;
import com.ibm.micro.client.mqttv3.*;
public class PubSyncSSL {
    public static void main(String[] args) {
        try {
            MqttClient client = new MqttClient(Example.SSLAddress, Example.clientId);
            MqttTopic topic = client.getTopic(Example.topicString);
            MqttMessage message = new MqttMessage(Example.publication.getBytes());
            message.setQos(Example.QoS);
            MqttConnectOptions conOptions = new MqttConnectOptions();
            conOptions.setSSLProperties(Example.getSSLSettings());
            client.connect(conOptions);
            System.out.println("Waiting for up to " + Example.sleepTimeout / 1000
                + " seconds for publication of \"" + message.toString()
                + "\" with QoS = " + message.getQos());
            System.out.println("On topic \"" + topic.getName() + "\" for client instance: \""
                + client.getClientId() + "\" on address " + client.getServerURI() + "\"");
            System.out.println("SSL Properties" + conOptions.getSSLProperties());
            MqttDeliveryToken token = topic.publish(message);
            token.waitForCompletion(Example.sleepTimeout);
            System.out.println("Delivery token \"" + token.hashCode()
                + "\" has been received: " + token.isComplete());
            client.disconnect();
        } catch (Exception e) {
            System.out.println("Client exception caught");
            e.printStackTrace();
        }
    }
}

```

Figure 109. PubSyncSSL . Java

The modifications to [Example.java](#) are shown in [Figure 110](#) on page 482.

```

public static final String      SSLAddress =
    System.getProperty("SSLAddress", "ssl://localhost:8883");

public static final Properties getSSLSettings() {
    final Properties properties = new Properties();
    properties.setProperty("com.ibm.ssl.keyStore", "C:\\Certificates\\SSClientKey.jks");
    properties.setProperty("com.ibm.ssl.keyStoreType", "JKS");
    properties.setProperty("com.ibm.ssl.keyStorePassword", "password");
    properties.setProperty("com.ibm.ssl.trustStore", "C:\\Certificates\\SSClientTrust.jks");
    properties.setProperty("com.ibm.ssl.trustStoreType", "JKS");
    properties.setProperty("com.ibm.ssl.trustStorePassword", "password");
    return properties;
}

```

Figure 110. Modifications to Example . java

Authenticating the telemetry channel

Clients authenticate the telemetry channel to encrypt the contents of the messages flowing on the channel, and to ensure a client connects to the correct telemetry channel. The server does not authenticate the client.

About this task

You can use a number of different keystore editors to create and manage self-signed certificates. The task uses the command line **keytool** command, which is part of the JRE. You can use the GUI tool **iKeyman**, which is shipped with WebSphere MQ to browse keystores and generate keys. Launch **iKeyman** using the command **strmqikm**.

Procedure

1. Create a telemetry channel, **SSLSSOptClients** that requires an SSL connection using the **New telemetry channel** wizard. The channel accepts anonymous clients.

Adapt your channel configuration from the following configuration stanza. Do not edit the telemetry properties file directly; use the wizard.

```
com.ibm.mq.MQXR.channel/SSLSSOptClients: \  
com.ibm.mq.MQXR.Port=8883;\  
com.ibm.mq.MQXR.Backlog=4096;\br/>com.ibm.mq.MQXR.KeyFileName=C:\\Certificates\\SSServerOptKey.jks;\br/>com.ibm.mq.MQXR.PassPhrase=password;\br/>com.ibm.mq.MQXR.ClientAuth=OPTIONAL;\br/>com.ibm.mq.MQXR.UserName=Admin;\br/>com.ibm.mq.MQXR.StartWithMQXRService=true
```

2. Generate the keys for client to authenticate the telemetry channel.

- a) Generate a self-signed key-pair for the telemetry channel in a new keystore, `SSServerOptKey.jks`:

```
Keytool -genkey -noprompt -alias SSServerPrivate  
-dname "CN=mqttserver.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"  
-keystore SSServerOptKey.jks -storepass password -keypass password
```

- b) Export its public certificate as an ASCII file, using the `-rfc` option:

```
Keytool -export -noprompt -alias SSServerPrivate -file SSServerPublic.cer  
-keystore SSServerOptKey.jks -storepass password -rfc
```

If you are running the task on windows, double-click `SSServerPublic.cer` to inspect its contents.

- c) Import the public certificate into a new client truststore, `SSClientTrust.jks`:

```
Keytool -import -noprompt -alias SSServerPublic -file SSServerPublic.cer  
-keystore SSClientTrust.jks -storepass password
```

- d) Create an empty client keystore, `SSClientKey.jks`.

Keytool does not have a command to create an empty keystore. You have two choices:

- i) Run **strmqikm**, and create a keystore, `SSClientKey.jks`, but do not add any keys.
- ii) Perform step 3a in [“Authenticating the telemetry channel and clients” on page 483](#), but do not use the keys yet.

Authenticating the telemetry channel and clients

Clients authenticate the telemetry channel and the telemetry channel authenticates clients attaching to it. Messages flowing on the channel are encrypted.

About this task

You can use a number of different keystore editors to create and manage self-signed certificates. The task uses the command line **keytool** command, which is part of the JRE. You can use the GUI tool **iKeyman**, which is shipped with WebSphere MQ to browse keystores and generate keys. Launch **iKeyman** using the command **strmqikm**.

The telemetry channel is configured with a different keystore to the task, [“Authenticating the telemetry channel” on page 482](#). You can use the same keystore, and omit step [“2” on page 484](#) to add keys to the keystore.

Procedure

1. Create a telemetry channel, `SSLSSReqClients` that requires an SSL connection using the **New telemetry channel** wizard. The channel accepts only authenticated clients.

Adapt your channel configuration from the following configuration stanza:

```
com.ibm.mq.MQXR.channel/SSLSSReqClients: \  
com.ibm.mq.MQXR.Port=8884;\
```

```
com.ibm.mq.MQXR.Backlog=4096;\
com.ibm.mq.MQXR.KeyFileName=C:\\Certificates\\SSServerReqKey.jks;\
com.ibm.mq.MQXR.PassPhrase=password;\
com.ibm.mq.MQXR.ClientAuth=REQUIRED;\
com.ibm.mq.MQXR.UserName=Admin;\
com.ibm.mq.MQXR.StartWithMQXRService=true
```

2. Generate the keys for client to authenticate the telemetry channel.

- a) Generate a self-signed key-pair for the telemetry channel in a new keystore, `SSServerReqKey.jks`:

```
Keytool -genkey -noprompt -alias SSServerPrivate
        -dname "CN=mqttserver.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore SSServerReqKey.jks -storepass password -keypass password
```

- b) Export its public certificate as an ASCII file, using the `-rfc` option:

```
Keytool -export -noprompt -alias SSServerPrivate -file SSServerPublic.cer
        -keystore SSServerReqKey.jks -storepass password -rfc
```

If you are running the task on windows, double-click `SSServerPublic.cer` to inspect its contents.

- c) Import the public certificate into a new client truststore, `SSClientTrust.jks`:

```
Keytool -import -noprompt -alias SSServerPublic -file SSServerPublic.cer
        -keystore SSClientTrust.jks -storepass password
```

3. Generate the keys for the telemetry channel to authenticate a client.

- a) Generate a self-signed key-pair for the client in a new keystore, `SSClientKey.jks`:

```
Keytool -genkey -noprompt -alias SSClientPrivate
        -dname "CN=mqttclient.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore SSClientKey.jks -storepass password -keypass password
```

- b) Export its public certificate as an ASCII file, using the `-rfc` option:

```
Keytool -export -noprompt -alias SSClientPrivate -file SSClientPublic.cer
        -keystore SSClientKey.jks -storepass password -rfc
```

If you are running the task on windows, double-click `SSClientPublic.cer` to inspect its contents.

- c) Import the public certificate into the server keystore, `SSServerReqKey.jks`:

```
Keytool -import -noprompt -alias SSClientPublic -file SSClientPublic.cer
        -keystore SSServerReqKey.jks -storepass password
```

Telemetry channels use the same store for both private keys and trusted certificates.

Authenticating an SSL telemetry connection using a certificate chain

Use signed certificates obtained from a certificate authority, or from implementing your own certification procedure, to authenticate an SSL connection. You have the option of authenticating the telemetry channel, or the telemetry channel and the clients that attach to it. Messages flowing on the connection are encrypted.

Before you begin

Do the task, “[Authenticating an SSL telemetry connection using self-signed certificates](#)” on page 480 before you start, to get [PubSyncSSL . Java](#) working with a secured TCP/IP connection using self-signed certificates.

About this task

In this task, modify the tasks “Authenticating the telemetry channel” on page 482, and “Authenticating the telemetry channel and clients” on page 483 in “Authenticating an SSL telemetry connection using self-signed certificates” on page 480, to work with keys certified by a certificate chain.

You can either obtain the certificates for this task from a certificate authority, or you can use a websites such as <http://www.openca.org/> to obtain certificates. Commercial certificate authorities generally provide trial certificates for a short period at no charge. This task was tested using commercially obtained certificates.

Another option is to build your own certification process, and run it on your own computers, using tools from websites such as <https://www.openssl.org/>.

The JRE cacerts truststores are not used in this task. You can use the JRE cacerts truststore at the client in the task, “Authenticating the telemetry channel” on page 485, instead of using the specified truststore. The certificate chain might be signed by a well known certificate authority that already has its root certificate in the cacerts store at the client. In this case, do not specify a truststore at the client. Make sure, if there are multiple JREs installed at the client that you manage the correct cacerts store.

Procedure

1. If you have not already done so, do the task, “Modifying PubSync.java to use SSL” on page 481, to modify PubSync.java to use SSL.
2. Configure the telemetry channel, and create the keystores to use SSL.
Authenticate just the telemetry channel, or the channel and clients that connect to it:
 - Do the task, “Authenticating the telemetry channel” on page 485, to connect with SSL, authenticating the telemetry channel.
 - Do the task, “Authenticating the telemetry channel and clients” on page 487, to connect with SSL, authenticating the telemetry channel and clients that connect to it.
3. Stop and restart the telemetry (MQXR) service to pick up changes to telemetry channel configurations.
4. Run the client program to see if the configuration works.

Authenticating the telemetry channel

Clients authenticate the telemetry channel to encrypt the contents of the messages flowing on the channel, and to ensure that a client connects to the correct telemetry channel. The server does not authenticate the client.

About this task

You can use a number of different keystore editors to create and manage certificates. The task uses the command line **keytool** command, which is part of the JRE. You can use the GUI tool **iKeyman**, which is shipped with WebSphere MQ to browse keystores and generate keys. Launch **iKeyman** using the command **strmqikm**.

Procedure

1. Create a telemetry channel, SSLCAOptClients that requires an SSL connection using the **New telemetry channel** wizard. The channel accepts anonymous clients.

Adapt your channel configuration from the following configuration stanza. Do not edit the telemetry properties file directly; use the wizard.

```
com.ibm.mq.MQXR.channel/SSLCAOptClients: \  
com.ibm.mq.MQXR.Port=8885;\  
com.ibm.mq.MQXR.Backlog=4096;\  
com.ibm.mq.MQXR.KeyFileName=C:\\Certificates\\CASServerOptKey.jks;\  
com.ibm.mq.MQXR.PassPhrase=password;\  
com.ibm.mq.MQXR.ClientAuth=OPTIONAL;\  

```

```
com.ibm.mq.MQXR.UserName=Admin;\ncom.ibm.mq.MQXR.StartWithMQXRService=true
```

2. Generate a CA-signed key for client to authenticate the telemetry channel.

- a) Generate a self-signed key-pair for the telemetry channel in a new keystore, `SSServerOptKey.jks`:

```
Keytool -genkey -noprompt -alias CASServerPrivate -keyalg RSA\n-dname "CN=mqttserverOpt.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"\n-keystore CASServerOptKey.jks -storepass password -keypass password
```

The key algorithm is set to RSA because some certificate authorities require it. The common name of the certificate must be unique, some certificate authorities do not issue keys with identical common names.

- b) Create a certificate signing request (CSR) as an ASCII file

```
Keytool -certreq -noprompt -alias CASServer -file CASServerOptKey.csr\n-dname "CN=mqttserverReq.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"\n-keystore CASServerOptKey.jks -storepass password -keypass password
```

- c) Run the certificate authority software, or logon to their website. Paste in the contents of `CASServerOptKey.csr` when asked for the CSR file.
- d) The certificate authority returns one or two certificates, and a signed response file as ASCII files. Paste the contents into two or three files:

Root certificate

Paste into `CARoot.cer`

Intermediate certificate

Paste into `CAInter.cer`

Server signed response file

Paste into `CASServerOpt.rsp`

The JRE certificates store are not used in this task. If you received one root certificate, and a signed response from the CA, use the root certificate and the signed response in the following steps. If you received a root and an intermediate certificate, use the intermediate certificate and the signed response.

- e) Receive the signed server response into the server keystore from which you issued the certificate request.

Receiving the response modifies the self-signed certificate so that it is signed by the CA. If you look at the certificate in the keystore before and after receiving the response, the signer changes. If it does not, an error is reported by the key management tool. Before using the certificate, inspect it, and verify that the signer is now the CA.

```
Keytool -import -noprompt -alias CASServer -file CASServerOpt.rsp\n-keystore CASServerOptKey.jks -storepass password
```

In some key management software, such as **iKeyman**, you receive, rather than import, response files.

- f) Import the CA certificate into the client truststore.

Import either the intermediate certificate if you received two certificates from the CA, or the root certificate, if you received only one certificate.

Either:

```
keytool -import -alias CAInter -file CAInter.cer\n-keystore CAClientTrust.jks -storepass password
```

Or:

```
keytool -import -alias CARoot -file CARoot.cer
        -keystore CAClientTrust.jks -storepass password
```

Authenticating the telemetry channel and clients

Clients authenticate the telemetry channel, and the telemetry channel authenticates clients attaching to it. Messages flowing on the channel are encrypted.

About this task

You can use a number of different keystore editors to create and manage certificates. The task uses the command line **keytool** command, which is part of the JRE. You can use the GUI tool **iKeyman**, which is shipped with WebSphere MQ to browse keystores and generate keys. Launch **iKeyman** using the command **strmqikm**.

The telemetry channel is configured with a different keystore to the one in the task, “Authenticating the telemetry channel” on page 485. You can use the same keystore, and omit step “2” on page 487 to add keys to the keystore.

Procedure

1. Create a telemetry channel, SSLCAReqClients that requires an SSL connection using the **New telemetry channel** wizard. The channel accepts only authenticated clients.

Adapt your channel configuration from the following configuration stanza. Do not edit the telemetry properties file directly; use the wizard.

```
com.ibm.mq.MQXR.channel/SSLCAReqClients: \
com.ibm.mq.MQXR.Port=8886;\
com.ibm.mq.MQXR.Backlog=4096;\
com.ibm.mq.MQXR.KeyFileName=C:\\Certificates\\CASServerReqKey.jks;\
com.ibm.mq.MQXR.PassPhrase=password;\
com.ibm.mq.MQXR.ClientAuth=REQUIRED;\
com.ibm.mq.MQXR.UserName=Admin;\
com.ibm.mq.MQXR.StartWithMQXRService=true
```

2. Generate a CA-signed key for client to authenticate the telemetry channel.

- a) Generate a self-signed key-pair for the telemetry channel in a new keystore, CASServerReqKey.jks:

```
Keytool -genkey -noprompt -alias CASServerPrivate -keyalg RSA
        -dname "CN=mqttsrverReq.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore CASServerReqKey.jks -storepass password -keypass password
```

The key algorithm is set to RSA because some certificate authorities require it. The common name of the certificate must be unique, some certificate authorities do not issue keys with identical common names.

- b) Create a certificate signing request (CSR) as an ASCII file

```
Keytool -certreq -noprompt -alias CASServer -file CASServerReqKey.csr
        -dname "CN=mqttsrverReq.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore CASServerReqKey.jks -storepass password -keypass password
```

- c) Run the certificate authority software, or logon to their website. Paste in the contents of CASServerReqKey.csr when asked for the CSR file.
- d) The certificate authority returns one or two certificates, and a signed response file as ASCII files. Paste the contents into two or three files:

Root certificate

Paste into CARoot.cer

Intermediate certificate

Paste into CAInter.cer

Server signed response file

Paste into CAserverReq.rsp

The JRE certificates store is not used in this task. If you received one root certificate, and a signed response from the CA, use the root certificate and the signed response in the following steps. If you received a root and an intermediate certificate, use the intermediate certificate and the signed response.

- e) Receive the signed server response into the server keystore from which you issued the certificate request.

Receiving the response modifies the self-signed certificate so that it is signed by the CA. If you look at the certificate in the keystore before and after receiving the response, the signer changes. If it does not, and error is reported by the key management tool. Before using the certificate, inspect it, and verify that the signer is now the CA.

```
Keytool -import -noprompt -alias CAserver -file CAserverReq.rsp
        -keystore CAserverReqKey.jks -storepass password
```

In some key management software, such as **iKeyman**, you receive, rather than import, response files.

- f) Import the CA certificate into the client truststore.

Import either the intermediate certificate if you received two certificates from the CA, or the root certificate, if you received only one certificate.

Either:

```
keytool -import -alias CAInter -file CAInter.cer
        -keystore CAClientTrust.jks -storepass password
```

Or:

```
keytool -import -alias CARoot -file CARoot.cer
        -keystore CAClientTrust.jks -storepass password
```

3. Generate a CA-signed key for the telemetry channel to authenticate clients.

- a) Generate a self-signed key-pair for the clients in a new keystore, CAClientKey.jks:

```
Keytool -genkey -noprompt -alias CAClientPrivate -keyalg RSA
        -dname "CN=mqttserverReq.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore CAClientKey.jks -storepass password -keypass password
```

The key algorithm is set to RSA because some certificate authorities require it. The common name of the certificate must be unique, some certificate authorities do not issue keys with identical common names.

- b) Create a certificate signing request (CSR) as an ASCII file

```
Keytool -certreq -noprompt -alias CAClient -file CAClientKey.csr
        -dname "CN=mqttserverReq.ibm.com, OU=ID, O=IBM, L=Hursley, S=Hants, C=GB"
        -keystore CAClientKey.jks -storepass password -keypass password
```

- c) Run the certificate authority software, or logon to their website. Paste in the contents of CAClientKey.csr when asked for the CSR file.
- d) The certificate authority returns one or two certificates, and a signed response file as ASCII files. Paste the contents into two or three files:

Root certificate

Paste into CARoot.cer

Intermediate certificate

Paste into CAInter.cer

Client signed response file

Paste into CAClient.rsp

The JRE certificates store are not used in this task. If you received one root certificate, and a signed response from the CA, use the root certificate and the signed response in the following steps. If you received a root and an intermediate certificate, use the intermediate certificate and the signed response.

- e) Receive the signed client response into the client keystore from which you issued the certificate request.

Receiving the response modifies the self-signed certificate so that it is signed by the CA. If you look at the certificate in the keystore before and after receiving the response, the signer changes. If it does not, and error is reported by the key management tool. Before using the certificate, inspect it, and verify that the signer is now the CA.

```
Keytool -import -noprompt -alias CAClient -file CAClient.rsp  
-keystore CAClientKey.jks -storepass password
```

In some key management software, such as **iKeyman**, you receive, rather than import, response files.

- f) Import the CA certificate into the server keystore.

Import either the intermediate certificate if you received two certificates from the CA, or the root certificate, if you received only one certificate.

Either:

```
keytool -import -alias CAInter -file CAInter.cer  
-keystore CAServerReqKey.jks -storepass password
```

Or:

```
keytool -import -alias CARoot -file CARoot.cer  
-keystore CAServerReqKey.jks -storepass password
```

Creating your first MQ Telemetry Transport publisher application using C

The steps to create an MQTT client publisher application are described in tutorial fashion. Each line of C code is explained. At the end of the task, you will have created an MQTT publisher.

Before you begin

The client application developed uses the client MQTT v3 C client libraries. The application connects to the WebSphere MQ Telemetry daemon for devices to publish messages. See [Creating your first publisher](#) for an example of a client communicating with WebSphere MQ Telemetry.

About this task

The example is a publish application, `pubsync.c`. The program `pubsync.c` publishes a message with the payload `Hello World!` to the topic `MQTT Example`, and waits for confirmation that the publication has been delivered to the daemon.

For simplicity, return codes from some functions used are not tested for correct completion. In production code, return codes can be checked to make sure that the program behaves as expected. Appropriate action must be taken if an unexpected error occurs.

By setting up a subscriber to `MQTT Example` you can check that the application works.

Use your selected C development environment to develop, build, and run the client. If you prefer, you can copy the code directly from the examples.

Procedure

1. Create a new, empty source file, `pubsync.c`

2. Create a file, `settings.h`. Copy the code in Figure 2 into the file.

All the parameters used in the program are defined in `settings.h`. You can override the values by changing the values in the file.

3. The steps that follow explain the code. Follow the steps or copy the code from [Figure 1](#) into `pubsync.c`.
4. Add the header file include statements for the required standard libraries and the `MQTTClient.h` and `settings.h` files.

```
#include "stdio.h"
#include "stdlib.h"
#include "string.h"
#include "MQTTClient.h"
#include "settings.h"
```

5. Start the definition of the `main()` function.

```
int main(int argc, char* argv[])
{
```

6. Define the local variables used in the program.

```
MQTTClient client;
MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
MQTTClient_message pubmsg;
MQTTClient_deliveryToken token;
int rc;
```

Note: Connection options are required by the `MQTTClient_connect` function. `MQTTClient_connectOptions_initializer` contains the default options.

7. Create a client.

```
MQTTClient_create(&client, ADDRESS, CLIENTID, MQTTCLIENT_PERSISTENCE_NONE, NULL);
```

- `&client` is a pointer to a handle for the newly created client. When this function returns with a 0 return code, it contains a handle to the new client. The example assumes success. Test the error code for correct completion in production code.
 - `ADDRESS` is the URI of the MQTT port that the daemon monitors for incoming client connection requests.
 - `CLIENTID` is the name used to identify the client to the daemon. Each active client must have a unique name. If you duplicate a client identifier in two running clients, an exception is thrown in both clients, and one client terminates. The name is used by the daemon to recognize a client that is reconnecting following a disconnection, see [The client identifier](#).
 - `MQTTCLIENT_PERSISTENCE_NONE` specifies that client state is held in memory and is lost if a system failure occurs. `MQTTCLIENT_PERSISTENCE_DEFAULT` specifies file system-based persistence, providing some protection against failures. For more specialized applications, you can use `MQTTCLIENT_PERSISTENCE_USER`, which provides an interface for you to implement your own persistence mechanism. For more details, see the API documentation for `MQTTClientPersistence.h`. Whether persistence is required is an application design question. For more details, see [Message persistence](#).
 - The default daemon TCP/IP port for MQTT is 1883. In the example, the default address is set to `tcp://localhost:1883`.
 - Until you call the `MQTTClient_connect` function, no message processing takes place.
8. Connect the client to the daemon.

```
if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS) {
    printf("Failed to connect, return code %d\n", rc);
    exit(-1);
}
```

- The `MQTTClient_connect` function is called, passing the client handle and a pointer to the connection options as arguments.
 - The return code from the `MQTTClient_connect` call is tested to make sure the connect request is successful.
 - If the `MQTTClient_connect` fails, the program ends with an error code of -1.
 - After the application connects, you can start publishing and subscribing.
 - A small "keep-alive" message is sent every 20 seconds to prevent the TCP/IP connection from being closed. This option is set by `conn_opts.keepAliveInterval`.
 - The session is started without checking for the completion of inflight messages remaining from a previous connection because `conn_opts.cleansession` is set to true. For more details, see [Clean sessions](#).
 - No last will and testament message is created for the connection. For more details, see [Last will and testament](#).
9. Populate the `MQTTClient_message` structure with the data to define the message payload and its attributes.

```
pubmsg.payload = PAYLOAD;
pubmsg.payloadlen = strlen(PAYLOAD);
pubmsg.qos = QOS;
pubmsg.retained = 0;
```

- `PAYLOAD` is our message content.
 - The example uses a string payload but MQTT payloads are byte arrays. The string length is required to specify the payload size.
 - The example publishes a QoS=1 message, so set the value accordingly
 - The retained attribute is set to false (0) as the message is not to be retained by the daemon. For more details, see [Retained publications](#).
10. Publish the message.

```
MQTTClient_publishMessage(client, TOPIC, &pubmsg, &token);
```

- The publish function specifies the client, the topic, and the payload to be sent to the daemon.
 - `TOPIC` is defined in `settings.h` as `MQTT_Example`.
 - The function is also passed a pointer to an `MQTTClient_deliveryToken`. This pointer is populated with a token representing the message when the function returns.
 - The message is now safely transferred to the MQTT client, but not yet transferred to the daemon. If the message has QoS=1 or 2, then the message is stored locally, in case the client fails before delivery is completed.
 - This function returns an error code that you can test for correct completion in production code.
11. Wait for acknowledgment from the server.

```
rc = MQTTClient_waitForCompletion(client, token, TIMEOUT);
```

- The `pubsync.c` example waits for an acknowledgment from the server, which confirms the message has been delivered.
 - The client and token arguments identify the specific message the program is waiting to be completed.
 - `TIMEOUT` limits how long the program waits for the message to complete delivery. The task, [Creating an asynchronous publisher for MQ Telemetry Transport using C](#) shows how to receive acknowledgments without waiting by using callback functions.
 - This function returns an error code that can be tested for correct completion in production code.
12. Disconnect the client from the daemon.

```
MQTTClient_disconnect(client, 10000);
```

- The client disconnects from the server and waits for any callback functions (not used in this example) for inflight messages to complete.
 - The second argument specifies a quiesce timeout in milliseconds. The example waits for up to 10 seconds to finish any other work it must do before disconnecting.
 - This function returns an error code that must be tested for correct completion in production code.
13. Free up memory used by the client and end the program.

```
MQTTClient_destroy(&client);
}
```

Results

To see the publications sent by this client, create a subscriber to the MQTT Example topic. For more details, see [Creating a subscriber for MQ Telemetry Transport using C](#)

Example

Figure 1 is a full listing of the code described in Procedure. The `settings.h` file in Figure 2 allows you to change the default parameters used in `pubsync.c`.

```
#include "stdio.h"
#include "stdlib.h"
#include "MQTTClient.h"
#include "settings.h"

int pubsync_main(int argc, char* argv[]) {
    MQTTClient client;
    MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
    MQTTClient_message pubmsg;
    MQTTClient_deliveryToken token;
    int rc;

    MQTTClient_create(&client, ADDRESS, CLIENTID, MQTTCLIENT_PERSISTENCE_NONE, NULL);
    if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS) {
        printf("Failed to connect, return code %d\n", rc);
        exit(-1);
    }
    pubmsg.payload = PAYLOAD;
    pubmsg.payloadlen = strlen(PAYLOAD);
    pubmsg.qos = QOS;
    pubmsg.retained = 0;
    MQTTClient_publishMessage(client, TOPIC, &pubmsg, &token);
    printf("Waiting for up to %d seconds for publication of %s\n",
           "on topic %s for client with ClientID: %s\n",
           TIMEOUT/1000, PAYLOAD, TOPIC, CLIENTID);
    rc = MQTTClient_waitForCompletion(client, token, TIMEOUT);
    printf("Message with delivery token %d delivered\n", token);
    MQTTClient_disconnect(client, 10000);
    MQTTClient_destroy(&client);
}
```

Figure 111. `pubsync.c`

```
#define ADDRESS "tcp://localhost:1883"
#define CLIENTID "ExampleClientPub"
#define TOPIC "MQTT Example"
#define PAYLOAD "Hello World!"
#define QOS 1
#define TIMEOUT 10000L
```

Figure 112. `settings.h`

Creating an asynchronous publisher for MQ Telemetry Transport using C

The steps to create an MQTT client asynchronous publisher application are described in tutorial fashion. Each line of C code is explained. At the end of the task, you will have created an MQTT asynchronous publisher.

In this task, you follow a tutorial to modify your first publisher application. The modifications enable the application to send publications without waiting for delivery acknowledgments. The delivery acknowledgments are received by a callback function that you create.

Before you begin

The client application developed uses the client MQTT v3 C client libraries. The application connects to the WebSphere MQ Telemetry daemon for devices to publish messages. See [Creating your first publisher](#) for an example of a client communicating with WebSphere MQ Telemetry.

About this task

The example is a publish application, `pubasync.c`. The program `pubasync.c` publishes a message with the payload `Hello World!` to the topic `MQTT Example`, without waiting for confirmation that the publication has been delivered to the daemon. The delivery acknowledgments are received in a callback function, `MQTTClient_deliveryComplete`.

For simplicity, return codes from some functions used are not tested for correct completion. In production code, return codes can be checked to make sure that the program behaves as expected. Appropriate action must be taken if an unexpected error occurs.

By setting up a subscriber to `MQTT Example` you can check that the application works.

Use your selected C development environment to develop, build, and run the client.

The steps in [Procedure](#) modify the `pubsync.c` application from “[Creating your first MQ Telemetry Transport publisher application using C](#)” on page 489. If you prefer, you can copy the code directly from the examples.

Procedure

1. Create a new, empty source file, `callback.h`.
2. Copy the code in [Figure 2](#) into the file.
 - `callback.h` declares the three callback methods needed for asynchronous client operation.
 - A variable, *deliveredtoken*, is also declared. This is accessed by the main program and the callback on different threads of execution. It is therefore declared as volatile. When using callbacks, take care to ensure that relevant variables are accessed in a threadsafe way.
3. Create a new, empty source file, `callback.c`.
4. Copy the code in [Figure 3](#) into the file.
 - `callback.c` implements the three callback methods used by the client for asynchronous operation, `delivered`, `msgarrvd`, and `connlost`.
5. Add an include statement for `callback.h` after the other includes in `pubasync.c`.

```
#include "callback.h"
```

6. Copy the contents of `pubsync.c` into a new file, `pubasync.c`.
7. Just before the `MQTTClient_connect` function call in `pubasync.c`, set the callback methods for the client.

```
MQTTClient_setCallbacks(client, NULL, connlost, msgarrvd, delivered);
```

- You must specify three callback functions. These functions are implemented in `callback.c`.

- MQTTClient_messageArrived is called when a message is sent to the client due to a matching subscription. This must return true when the received message has been successfully received by the client application. Returning false indicates to the client that your application had a problem receiving the message.
 - MQTTClient_connectionLost is called when the client loses its connection to the server.
 - MQTTClient_deliveryComplete is called when a QoS1 or QoS2 message has arrived and been acknowledged by the server. It is not called for QoS0 messages. In the example, this function saves the token from the delivered message in *deliveredtoken* to indicate that a message has arrived.
 - MQTTClient_setCallbacks must be called while the client is disconnected from the server.
 - The second argument allows you to pass contextual information to the callback functions. This is not used in the example, so is set to NULL.
8. Immediately before the call to MQTTClient_publishMessage, clear *deliveredtoken*. MQTTClient_deliveryComplete to sets *deliveredtoken* when a token is received.

```
deliveredtoken = 0;
```

9. Remove the MQTTClient_waitForCompletion call and the printf statement following it and replace with a loop waiting for a match of the original token and the token received in the callback.

```
while(deliveredtoken != token);
```

This is an example and does not cope with a number of situations that must be accommodated in production code design. These situations include:

- In case delivery is not completed, a timeout can be implemented
- Multiple messages can be inflight. The sample program only allows one delivery token to be checked at a time.

10. Disconnect the client from the daemon.

```
MQTTClient_disconnect(client, 10000);
```

- The client disconnects from the server and waits for any callback functions for inflight messages to complete.
 - The second argument specifies a quiesce timeout in milliseconds. The example waits for up to 10 seconds to finish any other work it must do before disconnecting.
 - This function returns an error code that must be tested for correct completion in production code.
11. Free up memory used by the client and end the program.

```
MQTTClient_destroy(&client);
}
```

Results

To see the publication sent by this client, create a subscriber to the MQTT Example topic. For more details, see [Creating a subscriber for MQ Telemetry Transport](#)

Example

pubasync.c, callbacks.c and callbacks.h are full listings of the code described in [Procedure](#).

```

#include "stdio.h"
#include "stdlib.h"
#include "MQTTClient.h"
#include "settings.h"
#include "callback.h"

int main(int argc, char* argv[]) {
    MQTTClient client;
    MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
    MQTTClient_message pubmsg;
    MQTTClient_deliveryToken token;
    int rc;

    MQTTClient_create(&client, ADDRESS, CLIENTID, MQTTCLIENT_PERSISTENCE_NONE, NULL);
    MQTTClient_setCallbacks(client, NULL, connlost, msgarrvd, delivered);
    if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS) {
        printf("Failed to connect, return code %d\n", rc);
        exit(-1);
    }
    pubmsg.payload = PAYLOAD;
    pubmsg.payloadlen = strlen(PAYLOAD);
    pubmsg.qos = QOS;
    pubmsg.retained = 0;
    deliveredtoken = 0;
    MQTTClient_publishMessage(client, TOPIC, &pubmsg, &token);
    printf("Waiting for publication of %s\n"
           "on topic %s for client with ClientID: %s\n", PAYLOAD, TOPIC, CLIENTID);
    while(deliveredtoken != token);
    MQTTClient_disconnect(client, 10000);
    MQTTClient_destroy(&client);
}

```

Figure 113. *pubasync.c*

```

MQTTClient_deliveryComplete delivered;
MQTTClient_messageArrived msgarrvd;
MQTTClient_connectionLost connlost;

extern volatile MQTTClient_deliveryToken deliveredtoken;

```

Figure 114. *callback.h*

```

#include "MQTTClient.h"

volatile MQTTClient_deliveryToken deliveredtoken;

void delivered(void *context, MQTTClient_deliveryToken dt)
{
    printf("Message with token value %d delivery confirmed\n", dt);
    deliveredtoken = dt;
}

int msgarrvd(void *context, char *topicName, int topicLen, MQTTClient_message *message)
{
    int i;
    char* payloadptr;

    printf("Message arrived\n");
    printf("    topic: %s\n", topicName);
    printf("    message: ");

    payloadptr = message->payload;
    for(i=0; i<message->payloadlen; i++) {
        putchar(*payloadptr++);
    }
    putchar('\n');
    MQTTClient_freeMessage(&message);
    free(topicName);
    return 1;
}

void connlost(void *context, char *cause)
{
    printf("\nConnection lost\n");
    printf("    cause: %s\n", cause);
}

```

Figure 115. *callback.c*

```

#define ADDRESS      "tcp://localhost:1883"
#define CLIENTID     "ExampleClientPub"
#define TOPIC        "MQTT Example"
#define PAYLOAD      "Hello World!"
#define QOS          1
#define TIMEOUT      10000L

```

Figure 116. *settings.h*

Creating a subscriber for MQ Telemetry Transport using C

The steps to create an MQTT client subscriber application are described in tutorial fashion. Each line of C code is explained. At the end of the task, you will have created an MQTT subscriber.

Before you begin

The client application developed uses the client MQTT v3 C client libraries. The application connects to the WebSphere MQ Telemetry daemon for devices to publish messages. See [Creating your first publisher](#) for an example of a client communicating with WebSphere MQ Telemetry.

About this task

The example is a subscriber application, `subscribe.c`. The program `subscribe.c` subscribes to the topic `MQTT Example` and waits for publications that match the subscription until the user ends the program.

A subscriber creates a subscription to a topic and waits for messages that match the subscription topic. Messages published while the client is disconnected, and that match a subscription created previously by the client, can be received when the client reconnects. The WebSphere MQ telemetry (MQXR) service or daemon for devices recognizes a client that has previously been connected by the client identifier. For more information, see [The client identifier](#). The `MQTTClient_connectOptions.cleansession`

boolean attribute controls whether publications sent previously are received or not. For more details, see [“Clean sessions”](#) on page 505.

For simplicity, return codes from some functions used are not tested for correct completion. In production code, return codes can be checked to ensure that the program behaves as expected. Appropriate action can be taken if an unexpected error occurs.

You can use the previously described publish example programs to send matching publications to WebSphere MQ Telemetry daemon for devices. Alternatively, use WebSphere MQ explorer to create test publications on the MQTT Example topic if you want to connect the client to a WebSphere MQ Telemetry channel.

The instructions in [Procedure](#) assume that you have already created `callback.c`, `callback.h`, and `settings.h` files in one of the earlier tasks.

Use your selected C development environment to develop, build, and run the client. If you prefer, you can copy the code directly from the examples.

Procedure

1. Create a copy of `settings.h` for this example, and change the `CLIENTID` define statement to the following:

```
#define CLIENTID "ExampleClientSub"
```

- If two clients with the same ID attempt to connect to a single server, one of them is forcibly disconnected. Typically, the new attempt to connect is successful and the older connection is disconnected.
 - Changing the `ClientID` allows you to use the previously developed publishing examples to send messages to this subscriber.
2. Create a new, empty source file, `subscribe.c`.
 3. The steps that follow explain the code. Follow the steps or copy the code from [Figure 117](#) on page 500 into the file `subscribe.c`.
 4. Add the header file include statements for the required standard libraries and the `MQTTClient.h` and `settings.h` files.

```
#include "stdio.h"
#include "stdlib.h"
#include "MQTTClient.h"
#include "settings.h"
```

5. Start the definition of the `main()` function.

```
int main(int argc, char* argv[]) {
```

6. Define the local variables used in the program.

```
MQTTClient client;
MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
MQTTClient_deliveryToken token;
int rc;
```

Connection options are required by the `MQTTClient_connect` function.
`MQTTClient_connectOptions_initializer` contains the default options.

7. Create a client.

```
MQTTClient_create(&client, ADDRESS, CLIENTID, MQTTCLIENT_PERSISTENCE_NONE, NULL);
```

- `&client` is a pointer to a handle for the newly created client. When this function returns with a 0 return code, the pointer contains a handle to the new client. The example assumes success. The error code can be tested for correct completion in production code.

- ADDRESS is the URI of the MQTT port that the daemon monitors for incoming client connection requests.
- CLIENTID is the name used to identify the client to the daemon. Each active client must have a unique name. If you duplicate a client identifier in two running clients, an exception is thrown in both clients, and one client terminates. The name is used by the daemon to recognize a client is reconnecting following a disconnection, see [The client identifier](#).
- MQTTCLIENT_PERSISTENCE_NONE specifies that client state is held in memory and is lost if a system failure occurs. MQTTCLIENT_PERSISTENCE#_DEFAULT specifies file system-based persistence, providing some protection against failures. For more specialized applications, you can use MQTTCLIENT_PERSISTENCE_USER, which provides an interface for you to implement your own persistence mechanism. Whether persistence is required is an application design question. For more details, see [Message persistence](#).
- The default daemon TCP/IP port for MQTT is 1883. In the example, the default address is set to tcp://localhost:1883.
- Until you call the MQTTClient_connect function, no message processing takes place.

8. Connect the client to the daemon

```
if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS) {
    printf("Failed to connect, return code %d\n", rc);
    exit(-1);
}
```

- The MQTTClient_connect function is called, passing the client handle and a pointer to the connection options as arguments.
- The return code from the MQTTClient_connect call is tested to make sure the connect request is successful.
- If the connect call fails, the program ends with an error code of -1.
- After the application connects, it can start publishing and subscribing.
- A small "keep-alive" message is sent every 20 seconds to prevent the TCP/IP connection from being closed. This option is set by conn_opts.keepAliveInterval.
- The session is started without checking for the completion of inflight messages remaining from a previous connection because conn_opts.cleansession is set to true. For more details, see [Clean sessions](#).
- No last will and testament message is created for the connection. For more detail, see [Last will and testament](#)

9. Subscribe to the topic.

```
MQTTClient_subscribe(client, TOPIC, QOS);
```

- Use the MQTTClient_subscribe function to subscribe the client application to the selected topic. The topic name can include wildcard characters. For more details, see [“Topic strings and topic filters in MQTT clients” on page 519](#).
- The QoS setting determines the maximum quality of service that is applied to messages sent to this subscriber. The server sends messages at the lower value of this setting and the QoS setting for the original message.
- This function returns an error code that can be tested for correct completion in production code.

10. Wait in a loop until the user enters a 'Q' character from the keyboard.

```
do {
    ch = getchar();
} while(ch!='Q' && ch != 'q');
```

The program now waits for messages to arrive. In this example, all message handling takes place in the callback function MQTTClient_messageArrived. For more details, see [“Receiving messages” on page 499](#).

11. Disconnect the client from the daemon.

```
MQTTClient_disconnect(client, 10000);
```

- The client disconnects from the server and waits for any callback functions (not used in this example) for inflight messages to complete.
- The second argument specifies a quiesce timeout in milliseconds. The example waits for up to 10 seconds to finish any other work it must perform before disconnecting.
- This function returns an error code that can be tested for correct completion in production code.

12. Free up memory used by the client and end the program.

```
MQTTClient_destroy(&client);  
}
```

Receiving messages

About this task

When messages arrive from the server, the `MQTTClient_messageArrived` function is started. The steps that follow explain the code.

Procedure

1. Start the definition of the callback function. This definition must match the `MQTTClient_messageArrived` function template.

```
int msgarrvd(void *context, char *topicName, int topicLen, MQTTClient_message *message) {
```

- `context` provides access to the context passed to the client library when the `MQTTClient_setCallbacks` function was called. This function is not used in the example.
- `topicName` is a pointer to the topic the received message is published to. If you have subscribed using wildcard characters, this parameter identifies the specific topic used for the message.
- `topicLen` is the length of the topic string. This option is provided for users who must embed NULL characters in topic strings.
- `message` is a pointer to the `MQTTClient_message` structure containing the message payload and attributes.

2. Define the local variables used.

```
int i;  
char* payloadptr;
```

These variables are used in the example to print out the payload by iterating over it.

3. Print out a message, displaying the topic and the payload of the message

```
printf("Message arrived\n");  
printf("    topic: %s\n", topicName);  
printf("    message: ");  
payloadptr = message->payload;  
for(i=0; i<message->payloadlen; i++){  
    putchar(*payloadptr++);  
}  
putchar('\n');
```

- The example assumes that the received payload is a sequence of printable characters.
- An MQTT payload is an array of bytes. The application is responsible for interpreting their meaning.

4. Free the memory used to store the message.

```
MQTTClient_freeMessage(&message);  
MQTTClient_free(topicName);
```

- In the example, all message handling takes place in the callback function.
- Ensure that the callback functions are short, and return control to its calling thread as soon as possible.
- The message pointer is passed for handling in the main part of the program.
- The main program must free the memory used by the message when processing is complete. `MQTTClient_freeMessage()` is a convenience function that returns the two memory blocks used to hold the `MQTTClient_message` structure and the message payload back to the system. The memory allocated to the `topicName` must be freed separately as shown.

5. Return a true value when the callback has successfully handled the message

```
    return 1;
}
```

- Returning a true value indicates that the client library can treat the message as successfully delivered.
- If the callback function cannot properly process the message, a false value is returned. For example, if the callback is putting messages onto a queue for the main program to process and the queue is full, returning false would be appropriate.
- For QoS1 and QoS2 messages, returning a false value indicates that the message was not delivered and further attempts to deliver it are made.

Example code

```
#include "stdio.h"
#include "stdlib.h"
#include "MQTTClient.h"
#include "settings.h"
#include "callback.h"

int main(int argc, char* argv[]) {
    MQTTClient client;
    MQTTClient_connectOptions conn_opts = MQTTClient_connectOptions_initializer;
    int rc;
    int ch;

    MQTTClient_create(&client, ADDRESS, CLIENTID, MQTTCLIENT_PERSISTENCE_NONE, NULL);

    MQTTClient_setCallbacks(client, NULL, connlost, msgarrvd, delivered);

    if ((rc = MQTTClient_connect(client, &conn_opts)) != MQTTCLIENT_SUCCESS) {
        printf("Failed to connect, return code %d\n", rc);
        exit(-1);
    }
    printf("Subscribing to topic %s\nfor client %s using QoS%d\n\n"
        "Press Q<Enter> to quit\n\n", TOPIC, CLIENTID, QOS);

    MQTTClient_subscribe(client, TOPIC, QOS);
    do {
        ch = getchar();
    } while(ch!='Q' && ch != 'q');
    MQTTClient_disconnect(client, 10000);
    MQTTClient_destroy(&client);
}
```

Figure 117. *subscriber.c*

```

#include "MQTTClient.h"

volatile MQTTClient_deliveryToken deliveredtoken;

void delivered(void *context, MQTTClient_deliveryToken dt) {
    printf("Message with token value %d delivery confirmed\n", dt);
    deliveredtoken = dt;
}

int msgarrvd(void *context, char *topicName, int topicLen, MQTTClient_message *message) {
    int i;
    char* payloadptr;

    printf("Message arrived\n");
    printf("    topic: %s\n", topicName);
    printf("    message: ");

    payloadptr = message->payload;
    for(i=0; i<message->payloadlen; i++) {
        putchar(*payloadptr++);
    }
    putchar('\n');
    MQTTClient_freeMessage(&message);
    MQTTClient_free(topicName);
    return 1;
}

void connlost(void *context, char *cause) {
    printf("\nConnection lost\n");
    printf("    cause: %s\n", cause);
}

```

Figure 118. *callback.h*

```

#define ADDRESS      "tcp://localhost:1883"
#define CLIENTID     "ExampleClientSub"
#define TOPIC        "MQTT Example"
#define PAYLOAD      "Hello World!"
#define QOS          1
#define TIMEOUT      10000L

```

Figure 119. *settings.h*

MQTT client programming concepts

The concepts described in this section help you to understand the Java, JavaScript and C client libraries for version 3.1 of the MQTT protocol. The concepts complement the API documentation accompanying the client libraries.

`com.ibm.micro.client.mqttv3` contains the classes that provide the public methods for the client libraries for the MQTT version 3.1 protocol. A version of the `com.ibm.micro.client.mqttv3` package, and the accompanying packages that implement the protocol for Java SE and ME, is provided with the installation of IBM WebSphere MQ Telemetry. To get the most recent version of the MQTT client libraries (Java, JavaScript and to view or download the API documentation, see ["MQTT client programming reference"](#).

To develop and run an MQTT client you need to copy or install these packages on the client device. You do not need to install a separate client runtime.

The licensing conditions for clients are associated with the server that you are connecting the clients to.

The MQTT client libraries are reference implementations of version 3.1 of the MQTT protocol. You can implement your own clients in different languages suitable for different device platforms. See [MQ Telemetry Transport format and protocol](#).

The API documentation makes no assumptions about which MQTT server the client is connected to. The behavior of the client might differ slightly when connected to different servers. The descriptions that follow describe the behavior of the client when connected to the IBM WebSphere MQ telemetry service.

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Callbacks

The `MqttCallback` interface has three callback methods; see an example implementation in [Callback.java](#).

connectionLost(`java.lang.Throwable` cause)

`connectionLost` is called when a communications error leads to the connection dropping. It is also called if the server drops the connection as a result of an error on the server after the connection has been established. Server errors are logged to the queue manager error log. The server drops the connection to the client, and the client calls `MqttCallback.connectionLost`. The only remote errors thrown as exceptions on the same thread as the client application are exceptions from `MqttClient.connect`. Errors detected by the server after the connection is established are reported back to the `MqttCallback.connectionLost` callback method as throwables.

Typical server errors that result in `connectionLost` are authorization errors. For example, the telemetry server tries to publish on a topic on behalf of a client that is not authorized to publish on the topic. Anything that results in a MQCC_FAIL condition code being returned to the telemetry server can result in the connection being dropped.

deliveryComplete(`MqttDeliveryToken` token)

`deliveryComplete` is called by the MQTT client to pass a delivery token back to the client application; see [“Delivery tokens”](#) on page 507. Using the delivery token, the callback can access the published message with the method `token.getMessage`.

When the application callback returns control to the MQTT client after being called by the `deliveryComplete` method, delivery is completed. Until delivery is completed, messages with QoS 1 or 2 are retained by the persistence class.

The call to `deliveryComplete` is a point of synchronization between the application and the persistence class. The `deliveryComplete` method is never called twice for the same message.

When the application callback returns from `deliveryComplete` to the MQTT client, the client calls `MqttClientPersistence.remove` for messages with QoS 1 or 2.

`MqttClientPersistence.remove` deletes the locally stored copy of the published message.

From a transaction processing perspective, the call to `deliveryComplete` is a single phase transaction that commits the delivery. If processing fails during the callback, on restart of the client `MqttClientPersistence.remove` is called again to delete the local copy of the published message. The callback is not called again. If you are using the callback to store a log of delivered messages, you cannot synchronize the log with the MQTT client. If you want to store a log reliably, then update the log in the `MqttClientPersistence` class.

The delivery token and message are referenced by the main application thread and the MQTT client. The MQTT client dereferences the `MqttMessage` object when delivery is completed, and the delivery token object when the client disconnects. The `MqttMessage` object can be garbage collected after delivery is completed if the client application dereferences it. The delivery token can be garbage collected after the session is disconnected.

You can get `MqttDeliveryToken` and `MqttMessage` attributes after a message has been published. If you attempt to set any `MqttMessage` attributes after the message has been published the result is undefined.

The MQTT client continues to process delivery acknowledgments if the client reconnects to the previous session with the same `ClientIdentifier`; see [“Clean sessions”](#) on page 505. The MQTT client application must set `MqttClient.CleanSession` to `false` for the previous session, and set it to `false` in the new session. The MQTT client creates new delivery tokens and message objects in the new session for pending deliveries. It recovers the objects using the

`MqttClientPersistence` class. If the application client still has references to the old delivery tokens and messages, dereference them. The application callback is called in the new session for any deliveries initiated in the previous session and completed in this session.

The application callback is called after the application client connects, when a pending delivery is completed. Before the application client connects, it can retrieve pending deliveries using the `MqttClient.getPendingDeliveryTokens` method.

Notice that the client application originally created the message object that is published, and its payload byte array. The MQTT client references these objects. The message object returned by the delivery token in the method `token.getMessage` is not necessarily the same message object created by the client. If a new MQTT client instance recreates the delivery token, the `MqttClientPersistence` class recreates the `MqttMessage` object. For consistency `token.getMessage` returns null if `token.isCompleted` is true, regardless of whether the message object was created by the application client or the `MqttClientPersistence` class.

`messageArrived(MqttTopic topic, MqttMessage message)`

`messageArrived` is called when a publication arrives for the client that matched a subscription topic. `topic` is the publication topic, not the subscription filter. The two can be different if the filter contains wildcards.

If the topic matches multiple subscriptions created by the client, the client receives multiple copies of the publication. If a client publishes to a topic that it also subscribes to, it receives a copy of its own publication.

If a message is sent with a QoS of 1 or 2, the message is stored by the `MqttClientPersistence` class before the MQTT client calls `messageArrived`. `messageArrived` behaves like `deliveryComplete`: it is only called once for a publication, and the local copy of the publication is removed by `MqttClientPersistence.remove` when `messageArrived` returns to the MQTT client. The MQTT client drops its references to the topic and message when `messageArrived` returns to the MQTT client. The topic and message objects are garbage collected, if the application client has not held onto a reference to the objects.

Callbacks, threading, and client application synchronization

The MQTT client calls a callback method on a separate thread to the main application thread. The client application does not create a thread for the callback, it is created by the MQTT client.

The MQTT client synchronizes callback methods. Only one instance of the callback method runs at a time. The synchronization makes it easy to update an object that tallies which publications have been delivered. One instance of the `MqttCallback.deliveryComplete` runs at a time, and so it is safe to update the tally without further synchronization. It is also the case that only one publication arrives at a time. Your code in the `messageArrived` method can update an object without synchronizing it. If you are referring to the tally, or the object that is being updated, in another thread, synchronize the tally or object.

The delivery token provides a synchronization mechanism between the main application thread and delivery of a publication. The method `token.waitForCompletion` waits until delivery of a specific publication is completed, or until an optional timeout expires. You might use `token.waitForCompletion` in a couple of simple ways to process one publication at a time:

1. To pause the application client until delivery of the publication is completed; see [Figure 88 on page 459](#).
2. To synchronize with the `MqttCallback.deliveryComplete` method. Only when `MqttCallback.deliveryComplete` returns to the MQTT Client does `token.waitForCompletion` resume. Using this mechanism you can synchronize running code in `MqttCallback.deliveryComplete` before code runs in the main application thread.

What if you wanted to publish without waiting for each publication to be delivered, but want confirmation when all the publications have been delivered? If you publish on a single thread, the last publication to be sent is also the last to be delivered.

Synchronization of requests sent to the server

Table 70 on page 504 describes the methods in the MQTT Java client that send a request to the server. Unless the application client sets an indefinite timeout, the client never waits indefinitely for the server. If the client hangs, it is either an application programming problem, or a defect in the MQTT client.

Table 70. Synchronization behavior of methods that result in requests to the server		
Method	Synchronization	Timeout interval
MqttClient.Connect	Waits for a connection to be established with the server.	Defaults to 30 seconds, or as set by a parameter, then throws an exception.
MqttClient.Disconnect	Waits for the MQTT client to finish any work it must do, and for the TCP/IP session to disconnect.	
MqttClient.Subscribe	Waits for completion of the Subscribe or UnSubscribe method.	
MqttClient.UnSubscribe		
MqttClient.Publish	Returns immediately to the application thread after passing the request to the MQTT client.	None.
MqttDeliveryToken.waitForCompletion	Waits for the delivery token to be returned.	Indefinite, or as set as a parameter.

Related concepts

[Clean sessions](#)

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

[Client identifier](#)

[Delivery tokens](#)

[Last will and testament publication](#)

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

[Message persistence in MQTT clients](#)

[Publications](#)

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

[Qualities of service provided by an MQTT client](#)

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

[Retained publications and MQTT clients](#)

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

[Subscriptions](#)

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

When you connect an MQTT client application using the `MqttClient.connect` method, the client identifies the connection using the client identifier and the address of the server. The server checks whether session information has been saved from a previous connection to the server. If a previous session still exists, and `cleanSession=true`, then the previous session information at the client and server is cleared. If `cleanSession=false` the previous session is resumed. If no previous session exists, a new session is started.

Note: The WebSphere MQ Administrator can forcibly close an open session and delete all the session information. If the client reopens the session with `cleanSession=false`, a new session is started.

Publications

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, all pending publication deliveries for the client are removed when the client connects.

The clean session setting has no effect on publications sent with `QoS=0`. For `QoS=1` and `QoS=2`, using `cleanSession=true` might result in losing a publication.

Subscriptions

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, any old subscriptions for the client are removed when the client connects. Any new subscriptions the client makes during the session are removed when it disconnects.

If you set `MqttConnectOptions.cleanSession` to `false` before connecting, any subscriptions the client creates are added to all the subscriptions that existed for the client before it connected. All the subscriptions remain active when the client disconnects.

Another way of understanding the way the `cleanSession` attribute affects subscriptions is to think of it as a modal attribute. In its default mode, `cleanSession=true`, the client creates subscriptions and receives publications only within the scope of the session. In the alternative mode, `cleanSession=false`, subscriptions are durable. The client can connect and disconnect and its subscriptions remain active. When the client reconnects, it receives any undelivered publications. While it is connected, it can modify the set of subscriptions that are active on its behalf.

You must set the `cleanSession` mode before connecting; the mode lasts for the whole session. To change its setting, you must disconnect and reconnect the client. If you change modes from using `cleanSession=false` to `cleanSession=true`, all previous subscriptions for the client, and any publications that have not been received, are discarded.

Related concepts

[Callbacks and synchronization in MQTT client applications](#)

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Client identifier

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Client identifier

The client identifier is a 23 byte string that identifies an MQTT client. Each identifier must be unique to only one connected client at a time. The identifier must contain only characters valid in a queue manager name. Within these constraints, you are able to use any identification string. It is important to have a procedure for allocating client identifiers, and a means of configuring a client with its chosen identifier.

The client identifier is used in the administration of an MQTT system. With potentially hundreds of thousands of clients to administer, you need to be able to identify a particular client rapidly. Suppose for example, a device has malfunctioned, and you are notified, perhaps by a customer ringing a help desk. How does the customer identify the device, and how do you correlate that identification with the server that is typically connected to the client? Do you have to consult a database that maps each device to a client identifier and to a server? Does the name of the device identify which server it is attached to? When you browse through MQTT client connections, each connection is labeled with the client identifier. Do you need to look up a table to map a client identifier to a physical device?

Does the client identifier identify a particular device, a user, or an application running at the client? If a customer replaces a faulty device with a new one, does the new device have the same identifier as the old device? Do you allocate a new identifier? If you change a physical device, but keep the same identifier, outstanding publications and active subscriptions are automatically transferred to the new device.

How do you ensure that client identifiers are unique? As well as a system for generating unique identifiers, you must have a reliable process for setting the identifier on the client. Perhaps the client device is a "black-box", with no user interface. Do you manufacture the device with a client identifier - such as using

its MAC address? Or do you have a software installation and configuration process that configures the device before it is activated?

You might create a client identifier from the 48 bit device MAC address, to keep the identifier short and unique. If transmission size is not a critical issue, you might use the remaining 17 bytes to make the address easier to administer.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Delivery tokens

When a client publishes on a topic a new delivery token is created. Use the delivery token to monitor the delivery of a publication, or to block the client application until delivery is complete.

The token is an `MqttDeliveryToken` object. It is created by calling the `MqttTopic.publish()` method and is retained by the MQTT client until the client session is disconnected and the delivery is completed.

The normal use of the token is to check whether delivery is complete. Block the client application until delivery is complete by using the returned token to call `token.waitForCompletion`. Alternatively,

provide a `MqttCallback` handler. When the MQTT client has received all the acknowledgments it expects as part of delivering the publication, it calls `MqttCallback.deliveryComplete` passing the delivery token as a parameter.

Until delivery is complete, you can inspect the publication using the returned delivery token by calling `token.getMessage`.

Completed deliveries

The completion of deliveries is asynchronous and depends on the quality of service associated with the publication.

At most once

`QoS=0`

Delivery is complete immediately on return from `MqttTopic.publish`.
`MqttCallback.deliveryComplete` is called immediately.

At least once

`QoS=1`

Delivery is complete when an acknowledgment to the publication has been received from the queue manager. `MqttCallback.deliveryComplete` is called when the acknowledgment is received. The message might be delivered more than once before `MqttCallback.deliveryComplete` is called, if communications are slow or unreliable.

Exactly once

`QoS=2`

Delivery is complete when the client receives a completion message that the publication has been published to subscribers. `MqttCallback.deliveryComplete` is called as soon as the publication message is received. It does not wait for the completion message.

In rare circumstances, your client application might not return to the MQTT client from `MqttCallback.deliveryComplete` normally. You know that delivery has completed, because the `MqttCallback.deliveryComplete` was called. If the client restarts the same session, `MqttCallback.deliveryComplete` does not get called again.

Incomplete deliveries

If the delivery is not complete after the client session is disconnected you can connect the client again and complete the delivery. You can only complete the delivery of a message if the message was published in a session with the `MqttConnectionOptions` attribute set to `false`.

Create the client using the same client identifier and server address, and then connect, setting the `cleanSession` `MqttConnectionOptions` attribute to `false` again. If you set `cleanSession` to `true`, pending delivery tokens are thrown away.

You can check if there are any pending deliveries by calling `MqttClient.getPendingDeliveryTokens`. You can call `MqttClient.getPendingDeliveryTokens` before connecting the client.

Related concepts

[Callbacks and synchronization in MQTT client applications](#)

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

[Clean sessions](#)

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to

run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Create a topic for the last will and testament. You might create a topic such as `MQTTManagement/Connections/server URI/client identifier/Last`.

Set up a "last will and testament" using the `MqttConnectionOptions.setWill(MqttTopic lastWillTopic, byte [] lastWillPayload, int lastWillQos, boolean lastWillRetained)` method.

Consider creating a time stamp in the `lastWillPayload` message. Include other client information that assists in identifying the client and the circumstances of the connection. Pass the `MqttConnectionOptions` object to the `MqttClient` constructor.

Set `lastWillQos` to 1 or 2, to make the message persistent in WebSphere MQ, and to guarantee delivery. To retain the last lost connection information, set the `lastWillRetained` to `true`.

The "last will and testament" publication is sent to subscribers if the connection ends unexpectedly. It is sent if the connection ends without the client calling the `MqttClient.disconnect` method.

To monitor connections, complement the "last will and testament" publication with other publications to record connections and programmed disconnections.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Delivery tokens

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Message persistence in MQTT clients

Publication messages are made persistent if they are sent with a quality of service of "at least once", or "exactly once". You can implement your own persistence mechanism on the client, or use the default persistence mechanism that is provided with the client. Persistence works in both directions, for publications sent to or from the client.

In MQTT, message persistence has two aspects; how the message is transferred, and whether it is queued in IBM MessageSight and IBM WebSphere MQ as a persistent message.

1. The MQTT client couples message persistence with quality of service. Depending on what quality of service you choose for a message, the message is made persistent. Message persistence is necessary to implement the required quality of service.

If you specify "at most once", `QoS=0`, the client discards the message as soon as it is published. If there is any failure in the upstream processing of the message, the message is not sent again. Even if the client remains active the message is not sent again. The behavior of `QoS=0` messages is the same as IBM WebSphere MQ fast nonpersistent messages.

If a message is published by a client with `QoS` of 1 or 2, it is made persistent. The message is stored locally, and only discarded from the client when it is no longer needed to guarantee "at least once", `QoS=1`, or "exactly once", `QoS=2`, delivery.

2. If a message is marked as QoS 1 or 2, it is queued in IBM MessageSight and IBM WebSphere MQ as a persistent message. If it is marked as QoS=0, then it is queued in IBM MessageSight and IBM WebSphere MQ as a nonpersistent message. In IBM WebSphere MQ nonpersistent messages are transferred between queue managers "exactly once", unless the message channel has the NPMSPEED attribute set to FAST.

A persistent publication is stored on the client until it is received by a client application. For QoS=2, the publication is discarded from the client when the application callback returns control. For QoS=1 the application might receive the publication again, if a failure occurs. For QoS=0, the callback receives the publication no more than once. It might not receive the publication if there is a failure, or if the client is disconnected at the time of publication.

When you subscribe to a topic, you can reduce the QoS with which the subscriber receives messages to match its persistence capabilities. Publications that are created at a higher QoS are sent with the highest QoS that the subscriber requested.

Storing messages

The implementation of data storage on small devices varies a great deal. The model of temporarily saving persistent messages in storage that is managed by the MQTT client might be too slow, or demand too much storage. In mobile devices, the mobile operating system might provide a storage service that is ideal for MQTT messages.

To provide flexibility in meeting the constraints of small devices, the MQTT client has two persistence interfaces. The interfaces define the operations that are involved in storing persistent messages. The interfaces are described in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see [MQTT client programming reference](#). You can implement the interfaces to suit a device. The MQTT client that runs on Java SE has a default implementation of the interfaces that store persistent messages in the file system. It uses the `java.io` package. The client also has a default implementation for Java ME, `MqttDefaultMIDPPersistence`.

Persistence classes

MqttClientPersistence

Pass an instance of your implementation of `MqttClientPersistence` to the MQTT client as a parameter of the `MqttClient` constructor. If you omit the `MqttClientPersistence` parameter from the `MqttClient` constructor, the MQTT client stores persistent messages using the class `MqttDefaultFilePersistence` or `MqttDefaultMIDPPersistence`.

MqttPersistable

`MqttClientPersistence` gets and puts `MqttPersistable` objects using a storage key. You must provide an implementation of `MqttPersistable` as well as the implementation of `MqttClientPersistence` if you are not using the `MqttDefaultFilePersistence` or `MqttDefaultMIDPPersistence`.

MqttDefaultFilePersistence

The MQTT client provides the `MqttDefaultFilePersistence` class. If you instantiate `MqttDefaultFilePersistence` in your client application, you can provide the directory to store persistent messages as a parameter of the `MqttDefaultFilePersistence` constructor.

Alternatively, the MQTT client can instantiate `MqttDefaultFilePersistence` and place files in a default directory. The name of the directory is *client identifier-tcp hostname portnumber*. `"\"`, `"\\\"`, `"/\"`, `":"` and `" "` are removed from the directory name string.

The path to the directory is the value of the system property `rcp.data`. If `rcp.data` is not set, the path is the value of the system property `usr.data`.

`rcp.data` is a property associated with installation of an OSGi or Eclipse Rich Client Platform (RCP).

`usr.data` is the directory in which the Java command that started the application was launched.

MqttDefaultMIDPPersistence

MqttDefaultMIDPPersistence has a default constructor and no parameters. It uses the `javax.microedition.rms.RecordStore` package to store messages.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

An `MqttMessage` has a byte array as its payload. Aim to keep messages as small as possible. The maximum length of message permitted by the MQTT protocol is 250 MB.

Typically, an MQTT client program uses `java.lang.String` or `java.lang.StringBuffer` to manipulate message contents. For convenience, the `MqttMessage` class has a `toString` method to convert its payload to a string. To create the byte array payload from a `java.lang.String` or `java.lang.StringBuffer`, use the `getBytes` method.

The `getBytes` method converts a string to the default character set for the platform. The default character set is generally UTF-8. MQTT publications that contain only text are usually encoded in UTF-8. Use the method `getBytes("UTF8")` to override the default character set.

In IBM WebSphere MQ, an MQTT publication is received as a `jms-bytes` message. The message includes an `MQRFH2` folder containing an `<mqtt>`, and an `<mmps>` folder. The `<mqtt>` folder contains the `clientId` and `qos`, but this content might change in the future.

An `MqttMessage` has three additional attributes: quality of service, whether it is retained, and whether it is a duplicate. The duplicate flag is only set if the quality of service is "at least once" or "exactly once". If the message was sent previously, and not acknowledged quickly enough by the MQTT client, the message is sent again, with the duplicate attribute set to `true`.

Publishing

To create a publication in an MQTT client application, create an `MqttMessage`. Set its payload, quality of service and whether it is retained, and call the `MqttTopic.publish(MqttMessage message)` method; `MqttDeliveryToken` is returned and the completion of the publication is asynchronous.

Alternatively, the MQTT client can create a temporary message object for you from the parameters on the `MqttTopic.publish(byte [] payload, int qos, boolean retained)` method when it creates a publication.

If the publication has an "at least once" or an "exactly once" quality of service, `QoS=1` or `QoS=2`, the MQTT client calls the `MqttClientPersistence` interface. It calls `MqttClientPersistence` to store the message before returning a delivery token to the application.

The application can choose to block until the message is delivered to the server, using the `MqttDeliveryToken.waitForCompletion` method. Alternatively, the application can continue without blocking. If you want to check if publications are delivered, without blocking, register an instance of a callback class that implements `MqttCallback` with the MQTT client. The MQTT client calls the `MqttCallback.deliveryComplete` method as soon as the publication has been delivered. Depending on the quality of service, the delivery might be almost immediate for `QoS=0`, or it might take some time for `QoS=2`.

Use the `MqttDeliveryToken.isComplete` method to poll if delivery is complete. While the value of `MqttDeliveryToken.isComplete` is `false`, you can call `MqttDeliveryToken.getMessage` to get the message contents. If the result of calling `MqttDeliveryToken.isComplete` is `true`, the message has been discarded and calling `MqttDeliveryToken.getMessage` would throw a null pointer exception. There is no built-in synchronization between `MqttDeliveryToken.getMessage` and `MqttDeliveryToken.isComplete`.

If the client disconnects before receiving all the pending delivery tokens, a new instance of the client can query pending delivery tokens before connecting. Until the client connects, no new deliveries are completed, and it is safe to call `MqttDeliveryToken.getMessage`. Use the `MqttDeliveryToken.getMessage` method to find out which publications have not been delivered. Pending delivery tokens are discarded if you connect with `MqttConnectOptions.cleanSession` set to its default value, `true`.

Subscribing

A queue manager or IBM MessageSight is responsible for creating publications to send to an MQTT subscriber. The queue manager checks if the topic filter in a subscription created by an MQTT client matches the topic string in a publication. The match can either be an exact match, or the match can include wildcards. Before the publication is forwarded to the subscriber by the queue manager, the queue manager checks the topic attributes associated with the publication. It follows the search procedure described in [Subscribing using a topic string that contains wildcard characters](#) to identify if an administrative topic object grants the user authority to subscribe.

When the MQTT client receives a publication with "at least once" quality of service, it calls the `MqttCallback.messageArrived` method to process the publication. If the quality of service of the

publication is "exactly once", QoS=2, the MQTT client calls the `MqttClientPersistence` interface to store the message when it is received. It then calls `MqttCallback.messageArrived`.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server.

Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

The quality of service of a publication is an attribute of `MqttMessage`. It is set by the method `MqttMessage.setQos`.

The method `MqttClient.subscribe` can reduce the quality of service applied to publications sent to a client on a topic. The quality of service of a publication forwarded to a subscriber might be different to the quality of service of the publication. The lower of the two values is used to forward a publication.

At most once

QoS=0

The message is delivered at most once, or it is not delivered at all. Its delivery across the network is not acknowledged.

The message is not stored. The message might be lost if the client is disconnected, or if the server fails.

QoS=0 is the fastest mode of transfer. It is sometimes called "fire and forget".

The MQTT protocol does not require servers to forward publications at QoS=0 to a client. If the client is disconnected at the time the server receives the publication, the publication might be discarded, depending on the server. The telemetry (MQXR) service does not discard messages sent with QoS=0. They are stored as nonpersistent messages, and are only discarded if the queue manager stops.

At least once

QoS=1

QoS=1 is the default mode of transfer.

The message is always delivered at least once. If the sender does not receive an acknowledgment, the message is sent again with the DUP flag set until an acknowledgment is received. As a result receiver can be sent the same message multiple times, and might process it multiple times.

The message must be stored locally at the sender and the receiver until it is processed.

The message is deleted from the receiver after it has processed the message. If the receiver is a broker, the message is published to its subscribers. If the receiver is a client, the message is delivered to the subscriber application. After the message is deleted, the receiver sends an acknowledgment to the sender.

The message is deleted from the sender after it has received an acknowledgment from the receiver.

Exactly once

QoS=2

The message is always delivered exactly once.

The message must be stored locally at the sender and the receiver until it is processed.

QoS=2 is the safest, but slowest mode of transfer. It takes at least two pairs of transmissions between the sender and receiver before the message is deleted from the sender. The message can be processed at the receiver after the first transmission.

In the first pair of transmissions, the sender transmits the message and gets acknowledgment from the receiver that it has stored the message. If the sender does not receive an acknowledgment, the message is sent again with the DUP flag set until an acknowledgment is received.

In the second pair of transmissions, the sender tells the receiver that it can complete processing the message, "PUBREL ". If the sender does not receive an acknowledgment of the "PUBREL " message, the "PUBREL " message is sent again until an acknowledgment is received. The sender deletes the message it saved when it receives the acknowledgment to the "PUBREL " message

The receiver can process the message in the first or second phases, provided that it does not reprocess the message. If the receiver is a broker, it publishes the message to subscribers. If the receiver is a client, it delivers the message to the subscriber application. The receiver sends a completion message back to the sender that it has finished processing the message.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server.

Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Use the `MqttMessage.setRetained` method to specify whether a publication on a topic is retained, or not.

To delete a retained publication in IBM WebSphere MQ, run the **CLEAR TOPICSTR** MQSC command.

If you create a publication with a null payload, the empty publication is forwarded to subscribers. Other MQTT brokers might not forward an empty publication to subscribers.

If you publish a non-retained publication to a topic that has a retained publication, the retained publication is not affected. Current subscribers receive the new publication. New subscribers receive the retained publication first, then receive any new publications.

When you create or update a retained publication, send the publication with a QoS of 1 or 2. If you send it with a QoS of 0, IBM WebSphere MQ creates a nonpersistent retained publication. The publication is not retained if the queue manager stops.

Use retained publications to record the latest value of a measurement. New subscribers to the retained topic immediately receive the most recent value of the measurement. If no new measurements are taken since the subscriber last subscribed to the publication topic, and if the subscriber subscribes again, the subscriber receives the most recent retained publication on the topic again.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to

run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

[Client identifier](#)

[Delivery tokens](#)

[Last will and testament publication](#)

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

[Message persistence in MQTT clients](#)

[Publications](#)

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

[Qualities of service provided by an MQTT client](#)

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

[Subscriptions](#)

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

[Topic strings and topic filters in MQTT clients](#)

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Subscriptions

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

Create subscriptions using the `MqttClient.subscribe` methods, passing one or more topic filters and quality of service parameters. The quality of service parameter sets the maximum quality of service that the subscriber is prepared to use to receive a message. Messages sent to this client cannot be delivered with a higher quality of service. The quality of service is set to the lower of the original value when the message was published and the level specified for the subscription. The default quality of service for receiving messages is `QoS=1`, at least once.

The subscription request itself is sent with `QoS=1`.

Publications are received by a subscriber when the MQTT client calls the `MqttCallback.messageArrived` method. The `messageArrived` method also passes the topic string with which the message was published to the subscriber.

You can remove a subscription, or a set of subscriptions, using the `MqttClient.unsubscribe` methods.

A WebSphere MQ command can remove a subscription. List subscriptions using WebSphere MQ Explorer, or by using **runmqsc** or PCF commands. All MQTT client subscriptions are named. They are given a name of the form: *ClientIdentifier:Topic name*

If you use the default `MqttConnectOptions`, or set `MqttConnectOptions.cleanSession` to `true` before connecting the client, any old subscriptions for the client are removed when the client connects. Any new subscriptions the client makes during the session are removed when it disconnects.

If you set `MqttConnectOptions.cleanSession` to `false` before connecting, any subscriptions the client creates are added to all the subscriptions that existed for the client before it connected. All the subscriptions remain active when the client disconnects.

Another way of understanding the way the `cleanSession` attribute affects subscriptions is to think of it as a modal attribute. In its default mode, `cleanSession=true`, the client creates subscriptions and receives publications only within the scope of the session. In the alternative mode, `cleanSession=false`, subscriptions are durable. The client can connect and disconnect and its subscriptions remain active. When the client reconnects, it receives any undelivered publications. While it is connected, it can modify the set of subscriptions that are active on its behalf.

You must set the `cleanSession` mode before connecting; the mode lasts for the whole session. To change its setting, you must disconnect and reconnect the client. If you change modes from using `cleanSession=false` to `cleanSession=true`, all previous subscriptions for the client, and any publications that have not been received, are discarded.

Publications that match active subscriptions are sent to the client as soon as they are published. If the client is disconnected, they are sent to the client if it reconnects to the same server with the same client identifier and `MqttConnectOptions.cleanSession` set to `false`.

Subscriptions for a particular client are identified by the client identifier. You can reconnect the client from a different client device to the same server, and continue with the same subscriptions and receive undelivered publications.

Related concepts

Callbacks and synchronization in MQTT client applications

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

Clean sessions

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

Client identifier

Delivery tokens

Last will and testament publication

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

Message persistence in MQTT clients

Publications

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

Qualities of service provided by an MQTT client

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

Retained publications and MQTT clients

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Topic strings and topic filters in MQTT clients

Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.

Topics strings are used to send publications to subscribers. Create a topic string using the method, `MqttClient.getTopic(java.lang.String topicString)`.

Topic filters are used to subscribe to topics and receive publications. Topic filters can contain wildcards. With wildcards, you can subscribe to multiple topics. Create a topic filter by using a subscription method; for example, `MqttClient.subscribe(java.lang.String topicFilter)`.

Topic strings

The syntax of a IBM WebSphere MQ topic string is described in [Topic Strings](#). The syntax of MQTT topic strings is described in the `MqttClient` class in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see [MQTT client programming reference](#).

The syntax of each type of topic string is almost identical. There are four minor differences:

1. Topic strings sent to IBM WebSphere MQ by MQTT clients must follow the convention for queue manager names. In particular, topic strings cannot contain hyphens.
2. The maximum lengths differ. IBM WebSphere MQ topic strings are limited to 10,240 characters. An MQTT client can create topic strings of up to 65535 bytes.
3. A topic string created by an MQTT client cannot contain a null character.
4. In WebSphere Message Broker, a null topic level, '...//...' was invalid. Null topic levels are supported by IBM WebSphere MQ.

Unlike IBM WebSphere MQ publish/subscribe, the mqttv3 protocol does not have a concept of an administrative topic object. You cannot construct a topic string from a topic object and a topic string. However, a topic string is mapped to an administrative topic in IBM WebSphere MQ. The access control associated with the administrative topic determines whether a publication is published to the topic, or discarded. The attributes that are applied to a publication when it is forwarded to subscribers, are influenced by the attributes of the administrative topic.

Topic filters

The syntax of a IBM WebSphere MQ topic filter is described in [Topic-based wildcard scheme](#). The syntax of the topic filters you can construct with an MQTT client are described in the `MqttClient` class in the API documentation for the MQTT client for Java. For links to client API documentation for the MQTT client libraries, see [MQTT client programming reference](#).

The syntax of each type of topic filter is almost identical. The only difference is in the way different MQTT brokers interpret a topic filter. In WebSphere Message Broker V6, a multilevel wildcard could only be used at the end of a topic filter. In IBM WebSphere MQ, a multilevel wildcard can be used at any level in the topic tree; for example `USA/#/Dutchess County`.

Related concepts

[Callbacks and synchronization in MQTT client applications](#)

The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements `MqttCallback`.

[Clean sessions](#)

The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to

run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting `MqttConnectOptions.cleanSession` before connecting.

[Client identifier](#)

[Delivery tokens](#)

[Last will and testament publication](#)

If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.

[Message persistence in MQTT clients](#)

[Publications](#)

Publications are instances of `MqttMessage` that are associated with a topic string. MQTT client can create publications to send to IBM WebSphere MQ, and subscribe to topics on IBM WebSphere MQ to receive publications.

[Qualities of service provided by an MQTT client](#)

An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.

[Retained publications and MQTT clients](#)

If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.

[Subscriptions](#)

Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.

C client programming concepts

Differences between the C and Java client for version 3.1 of the MQ Telemetry Transport are described in this topic. The topic complements the client concepts and the C reference information.

The topic is organized in the same way as "[MQTT client programming concepts](#)" on page 501. Each heading corresponds to a topic in *WebSphere(r) MQ Telemetry Transport client programming concepts*. The sections describe differences between the C client and the Java client. Small differences in the signatures between the Java methods and the C functions are not described.

The C client is most often used to implement a lightweight adapter between a telemetry device and the WebSphere MQ Telemetry daemon for devices. The daemon is commonly used as a network concentrator between very lightweight telemetry devices and the telemetry (MQXR) service.

The WebSphere MQ Telemetry daemon for devices is also a C client and differences in its behaviour from the telemetry (MQXR) service are described. The daemon does not provide an implementation of JAAS or SSL for clients connecting to it.

`mqttclient.dll` and `mqttclient.lib` are the 32-bit Windows libraries that contain client functions for the C implementation of the MQ Telemetry Transport version 3.1 protocol. The 32-bit Linux libraries are `libmqttclient.so` and `libmqttclient.a`. Two header files contain the function and other declarations needed by client applications: `MQTTClient.h` and `MQTTClientPersistence.h`. These files are provided with the installation of WebSphere MQ Telemetry.

To develop and run an MQ Telemetry Transport client you need to copy these files onto the client device. Unlike WebSphere MQ clients, you do not need to install a separate client runtime.

Consult the licensing conditions associated with the WebSphere MQ Telemetry feature that govern connecting MQ Telemetry Transport clients to WebSphere MQ and the WebSphere MQ Telemetry daemon for devices.

The C client is a reference implementation of version 3.1 of the MQ Telemetry Transport. You can implement your own clients in different languages suitable for different device platforms. Refer to [MQ Telemetry Transport format and protocol](#) for details.

The MQTT client identifier

“Client identifier” on page 506	The client identifier is a 23 byte string that identifies an MQTT client. Each identifier must be unique to only one connected client at a time. The identifier must contain only characters valid in a queue manager name. Within these constraints, you are able to use any identification string. It is important to have a procedure for allocating client identifiers, and a means of configuring a client with its chosen identifier.
-------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- No differences.

Publications

“Publications” on page 512	Publications are instances of <code>MqtttMessage</code> that are associated with a topic string. MQTT
--------------------------------------------	-------------------------------------------------------------------------------------------------------

- The callback function is not called for publications with the "fire and forget", QoS=0, quality of service.

Delivery tokens

“Delivery tokens” on page 507	When a client publishes on a topic a new delivery token is created. Use the delivery token to monitor the delivery of a publication, or to block the client application until delivery is complete.
-----------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- A delivery token is an `int`. It has a typedef of `MQTTClient_deliveryToken`
- The callback function is not called for publications with the "fire and forget", QoS=0, quality of service.

Retained publications

“Retained publications and MQTT clients” on page 516	If you create a subscription to a topic that has a retained publication, the most recent retained publication on the topic is immediately forwarded to you.
----------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------

- Retained messages are only saved in the daemon if persistence is configured; see [Saving retained messages and subscriptions](#).

For WebSphere MQ, the quality of service affects whether a retained message is saved permanently. If a client is attached to the telemetry service, retained messages with "fire and forget", QoS=0 quality of service are discarded, if the queue manager shuts down.

Subscriptions

“Subscriptions” on page 517	Create subscriptions to register an interest in publication topics using a topic filter. A client can create multiple subscriptions, or a subscription containing a topic filter that uses wildcards, to register an interest in multiple topics. Publications on topics matching the filters are sent to the client. Subscriptions can remain active while a client is disconnected. The publications are sent to the client when it reconnects.
---------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- Durable subscriptions are only saved in the daemon if persistence is configured; see [Saving retained messages and subscriptions](#).
- Publications can be received synchronously. Call the MQTTClient_receive function.

Callbacks and synchronization

“Callbacks and synchronization in MQTT client applications” on page 502	<p>The MQTT client programming model uses threads extensively. The threads decouple an MQTT client application, as much as they can, from delays in transmitting messages to and from the server. Publications, delivery tokens, and connection lost events are delivered to the methods in a callback class that implements MqttCallback.</p>
-----------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- The operation of synchronization in the C client is modal. Calling MQTTClient_setCallback puts the client into asynchronous mode.
- In synchronous mode, the application client must voluntarily yield control so the MQTT client can process acknowledgements and issue MQTT pings to keep the network alive. Yield control by calling MQTTClient_receive or MQTTClient_yield.

Topic strings and filters

“Topic strings and topic filters in MQTT clients” on page 519	<p>Topic strings and topic filters are used to publish and to subscribe. The syntax of topic strings and filters in MQTT clients is largely the same as topic strings in IBM WebSphere MQ.</p>
-------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- The WebSphere MQ Telemetry daemon for devices handles the multi-level wildcard # differently from WebSphere MQ v7. /# must be the last two characters in the filter string for # to behave as a wildcard. In WebSphere MQ v7, . . /# / . . is a valid use of the multilevel wildcard. WebSphere MQ Telemetry daemon for devices treats the multi-level wildcard the same as WebSphere MQ Broker v6.

Quality of service

“Qualities of service provided by an MQTT client” on page 514	<p>An MQTT client provides three qualities of service for delivering publications to WebSphere MQ and to the MQTT client: "at most once", "at least once" and "exactly once". When an MQTT client sends a request to WebSphere MQ to create a subscription, the request is sent with the "at least once" quality of service.</p>
-------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- No differences.

Message persistence

“Message persistence in MQTT clients” on page 510	<p>Publication messages are made persistent if they are sent with a quality of service of "at least once", or "exactly once". You can implement your own persistence mechanism on the client, or use the default persistence mechanism that is provided with the client. Persistence works in both directions, for publications sent to or from the client.</p>
-------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- Because of language binding differences, set the message persistence mechanism in the C client as follows. Call the MQTT C client with one of the three options set as the fourth parameter to MQTTClient_create:

MQTTCLIENT_PERSISTENCE_DEFAULT

A file based persistence, the details of which are specific to the client platform.

MQTTCLIENT_PERSISTENCE_NONE

Data is only held in memory and lost when the client stops. The WebSphere MQ Telemetry daemon for devices supports only this option.

MQTTCLIENT_PERSISTENCE_USER

You can develop functions to implement your own persistence mechanism. Pass a structure, `MQTTClient_persistence` containing pointers to your functions on the `MQTTClient_create` call. Read the MQTT C client reference information for details.

Clean sessions

“Clean sessions” on page 505	The MQTT client, and the telemetry (MQXR) service maintain session state information. The state information is used to ensure "at least once" and "exactly once" delivery, and "exactly once" receipt of publications. Session state also includes subscriptions created by an MQTT client. You can choose to run an MQTT client with or without maintaining state information between sessions. Change the clean session mode by setting <code>MqttConnectOptions.cleanSession</code> before connecting.
----------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- No differences.

Last will and testament

“Last will and testament publication” on page 509	If an MQTT client connection unexpectedly ends, you can configure WebSphere MQ Telemetry to send a "last will and testament" publication. Predefine the content of the publication, and the topic to send it to. The "last will and testament" is a connection property. Create it before connecting the client.
-------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- No differences.

Handling program errors

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

Whenever possible, the queue manager returns any errors as soon as an MQI call is made. These are *locally determined errors*.

When sending messages to a remote queue, errors might not be apparent when the MQI call is made. In this case, the queue manager that identifies the errors reports them by sending another message to the originating program. These are *remotely determined errors*.

Locally determined errors

Information about locally determined errors which include: failure on an MQI call, system interruptions, and messages containing incorrect data.

The three most common causes of errors that the queue manager can report immediately are:

- Failure of an MQI call; for example, because a queue is full
- An interruption to the running of some part of the system on which your application depends; for example, the queue manager
- Messages containing data that cannot be processed successfully

If you are using the asynchronous put facility, errors are not reported immediately. Use the MQSTAT call to retrieve status information about previous asynchronous put operations.

Failure of an MQI call

The queue manager can report immediately any errors in the coding of an MQI call. It does this using a set of predefined return codes. These are divided into completion codes and reason codes.

To show whether a call is successful, the queue manager returns a *completion code* when the call completes. There are three completion codes, indicating success, partial completion, and failure of the call. The queue manager also returns a *reason code* that indicates the reason for the partial completion or the failure of the call.

The completion and reason codes for each call are listed with the description of that call in [Return codes](#). For more detailed information, including ideas for corrective action, see:

- [Reason codes](#) for all other WebSphere MQ platforms

Design your programs to handle all the return codes that can arise from each call.

System interruptions

Your application might be unaware of any interruption if the queue manager to which it is connected has to recover from a system failure. However, you must design your application to ensure that your data is not lost if such an interruption occurs.

The methods that you can use to make sure that your data remains consistent depends on the platform on which your queue manager is running:

UNIX, Linux and Windows systems

In these environments, you can make your MQPUT and MQGET calls in the usual way, but you must declare sync points by using the MQCMIT and MQBACK calls (see “Committing and backing out units of work” on page 309). In the CICS environment, MQCMIT and MQBACK commands are disabled, because you can make your MQPUT and MQGET calls within units of work that are managed by CICS.

Use persistent messages for carrying all data that you cannot afford to lose. Persistent messages are reinstated on queues if the queue manager has to recover from a failure. With WebSphere MQ on UNIX, Linux, and Windows systems, an MQGET or MQPUT call within your application will fail at the point of filling all the log files, with the message MQRC_RESOURCE_PROBLEM. For more information about log files on AIX, HP-UX, Linux, Solaris, and Windows systems, see [Administering](#).

If the queue manager is stopped by an operator while an application is running, the quiesce option is usually used. The queue manager enters a quiescing state in which applications can continue to do work, but they must terminate as soon as convenient. Small, quick applications can probably ignore the quiescing state and continue until they terminate as normal. Longer running applications, or ones that wait for messages to arrive, should use the *fail if quiescing* option when they use the MQOPEN, MQPUT, MQPUT1, and MQGET calls. These options mean that the calls fail when the queue manager quiesces, but the application might still have time to terminate cleanly by issuing calls that ignore the quiescing state. Such applications could also commit, or back out, changes that they have made, and then terminate.

If the queue manager is forced to stop (that is, stop without quiescing), applications will receive the MQRC_CONNECTION_BROKEN reason code when they make MQI calls. Exit the application or, alternatively, on UNIX, Linux, and Windows systems, issue an MQDISC call.

Messages containing incorrect data

When you use units of work in your application, if a program cannot successfully process a message that it retrieves from a queue, the MQGET call is backed out.

The queue manager maintains a count (in the *BackoutCount* field of the message descriptor) of the number of times that happens. It maintains this count in the descriptor of each message that is affected. This count can provide valuable information about the efficiency of an application. Messages with backout counts that are increasing over time are being repeatedly rejected; design your application so that it analyzes the reasons for this and handles such messages accordingly.

On WebSphere MQ for Windows, UNIX and Linux systems, the backout count always survives the queue manager restarts.

Using report messages for problem determination

The remote queue manager cannot report errors such as failing to put a message on a queue when you make your MQI call, but it can send you a report message to say how it has processed your message.

Within your application you can create (MQPUT) report messages as well as select the option to receive them (in which case they are sent by either another application or by a queue manager).

Creating report messages

Report messages enable an application to tell another application that it cannot deal with the message that was sent.

However, the *Report* field must initially be analyzed to determine whether the application that sent the message is interested in being informed of any problems. Having determined that a report message is required, you have to decide:

- Whether you want to include the entire original message, just the first 100 bytes of data, or none of the original message.
- What to do with the original message. You can discard it or let it go to the dead-letter queue.
- Whether the contents of the *MsgId* and *CorrelId* fields are needed as well.

Use the *Feedback* field to indicate the reason for the report message being generated. Put your report messages on an application's reply-to queue. See [Feedback](#) for further information.

Requesting and receiving (MQGET) report messages

When you send a message to another application, you are not informed of any problems unless you complete the *Report* field to indicate the feedback that you require. See [Structure of the report field](#) for the options available.

Queue managers always put report messages on an application's reply-to queue and it is recommended that your own applications do the same. When you use the report message facility, specify the name of your reply-to queue in the message descriptor of your message; otherwise, the MQPUT call fails.

Your application must contain procedures that monitor your reply-to queue and process any messages that arrive on it. Remember that a report message can contain all the original message, the first 100 bytes of the original message, or none of the original message.

The queue manager sets the *Feedback* field of the report message to indicate the reason for the error; for example, the target queue does not exist. Your programs should do the same.

For more information about report messages, see [“Report messages” on page 11](#).

Remotely determined errors

When you send messages to a remote queue, even when the local queue manager has processed your MQI call without finding an error, other factors can influence how your message is handled by a remote queue manager.

For example, the queue that you are targeting might be full, or might not even exist. If your message has to be handled by other intermediate queue managers on the route to the target queue, any of these could find an error.

Problems delivering a message

When an MQPUT call fails, you can try to put the message on the queue again, return it to the sender, or put it on the dead-letter queue.

Each option has its merits, but you might not want to try putting a message again if the reason that the MQPUT failed was because the destination queue was full. In this instance, putting it on the dead-letter queue allows you to deliver it to the correct destination queue later on.

Retry message delivery

Before the message is put on a dead-letter queue, a remote queue manager attempts to put the message on the queue again if the attributes *MsgRetryCount* and *MsgRetryInterval* have been set for the channel, or if there is a retry exit program for it to use (the name of which is held in the channel attribute *MsgRetryExitId* field).

If the *MsgRetryExitId* field is blank, the values in the attributes *MsgRetryCount* and *MsgRetryInterval* are used.

If the *MsgRetryExitId* field is not blank, the exit program of this name runs. For more information about using your own exit programs, see [“Channel-exit programs for messaging channels”](#) on page 378.

Return message to sender

You return a message to the sender by requesting a report message to be generated to include all of the original message.

See [“Report messages”](#) on page 11 for details on report message options.

Using the dead-letter (undelivered message) queue

When a queue manager cannot deliver a message, it attempts to put the message on its dead-letter queue. This queue should be defined when the queue manager is installed.

Your programs can use the dead-letter queue in the same way that the queue manager uses it. You can find the name of the dead-letter queue by opening the queue manager object (using the MQOPEN call) and inquiring about the *DeadLetterQName* attribute (using the MQINQ call).

When the queue manager puts a message on this queue, it adds a header to the message, the format of which is described by the dead-letter header (MQDLH) structure; see [MQDLH - Dead-letter header](#). This header includes the name of the target queue and the reason that the message was put on the dead-letter queue. It must be removed and the problem must be resolved before the message is put on the intended queue. Also, the queue manager changes the *Format* field of the message descriptor (MQMD) to indicate that the message contains an MQDLH structure.

MQDLH structure

You are recommended to add an MQDLH structure to all messages that you put on the dead-letter queue; however, if you intend to use the dead-letter handler provided by certain WebSphere MQ products, you **must** add an MQDLH structure to your messages.

The addition of the header to a message might make the message too long for the dead-letter queue, so always make sure that your messages are shorter than the maximum size allowed for the dead-letter queue, by at least the value of the MQ_MSG_HEADER_LENGTH constant. The maximum size of messages allowed on a queue is determined by the value of the *MaxMsgLength* attribute of the queue. For the dead-letter queue, make sure that this attribute is set to the maximum allowed by the queue manager. If your application cannot deliver a message, and the message is too long to be put on the dead-letter queue, follow the advice given in the description of the MQDLH structure.

Ensure that the dead-letter queue is monitored, and that any messages arriving on it get processed. The dead-letter queue handler runs as a batch utility and can be used to perform various actions on selected messages on the dead-letter queue. For further details, see [“Dead-letter queue processing”](#) on page 527.

If data conversion is necessary, the queue manager converts the header information when you use the MQGMO_CONVERT option on the MQGET call. If the process putting the message is an MCA, the header is followed by all the text of the original message.

Messages put on the dead-letter queue might be truncated if they are too long for this queue. A possible indication of this situation is the messages on the dead-letter queue being the same length as the value of the *MaxMsgLength* attribute of the queue.

Dead-letter queue processing

This information contains general-use programming interface information when using dead-letter queue processing.

Dead-letter queue processing depends on local system requirements, but consider the following things when you draw up the specification:

- The message can be identified as having a dead-letter queue header because the value of the format field in the MQMD, is MQFMT_DEAD_LETTER_HEADER.
- On WebSphere MQ for z/OS using CICS, if an MCA puts this message to the dead-letter queue, the *PutApplType* field is MQAT_CICS, and the *PutApplName* field is the *ApplId* of the CICS system followed by the transaction name of the MCA.
- The reason for the message to be routed to the dead-letter queue is contained in the *Reason* field of the dead-letter queue header.
- The dead-letter queue header contains details of the destination queue name and queue manager name.
- The dead-letter queue header contains fields that have to be reinstated in the message descriptor before the message is put to the destination queue. These are:
 1. *Encoding*
 2. *CodedCharSetId*
 3. *Format*
- The message descriptor is the same as PUT by the original application, except for the three fields shown (Encoding, CodedCharSetId, and Format).

Your dead-letter queue application must do one or more of the following things:

- Examine the *Reason* field. A message might have been put by an MCA for the following reasons:
 - The message was longer than the maximum message size for the channel
The reason is MQRC_MSG_TOO_BIG_FOR_CHANNEL
 - The message could not be put to its destination queue
The reason is any MQRC_* reason code that can be returned by an MQPUT operation
 - A user exit has requested this action
The reason code is that supplied by the user exit, or the default MQRC_SUPPRESSED_BY_EXIT
- Try to forward the message to its intended destination, where this is possible.
- Retain the message for a certain length of time before discarding when the reason for the diversion is determined, but not immediately correctable.
- Give instructions to administrators correct problems where these have been determined.
- Discard messages that are corrupted or otherwise not processible.

There are two ways to deal with the messages that you have recovered from the dead-letter queue:

1. If the message is for a local queue:
 - Carry out any code translations required to extract the application data
 - Carry out code conversions on that data if this is a local function
 - Put the resulting message on the local queue with all the detail of the message descriptor restored
2. If the message is for a remote queue, put the message on the queue.

For information about how undelivered messages are handled in a distributed queuing environment, see [What happens when a message cannot be delivered?](#).

Multicast programming

Use this information to learn about the WebSphere MQ Multicast programming tasks such as connecting to a queue manager and exception reporting.

WebSphere MQ Multicast was designed to be as transparent to the user as possible and yet still be compatible with existing applications. Defining a COMMINFO object and setting the TOPIC object's **MCAST** and **COMMINFO** parameters, means that existing WebSphere MQ applications do not require substantial rewriting to use multicast. However, there might be some limitations (see “Multicast and the Message Queue Interface” on page 528 for more information) and some security issues to consider (see [Multicast security](#) for more information).

Multicast and the Message Queue Interface

Use this information to understand the major MQI concepts and how they relate to WebSphere MQ Multicast.

Multicast subscriptions are nondurable; because there are no physical queues involved, there is nowhere to store the offline messages that are created by durable subscriptions.

After an application has subscribed to a multicast topic, it is given back an object handle which it can consume or MQGET from, as if it were a handle to a queue. This means that only managed multicast subscriptions (subscriptions created with MQSO_MANAGED) are supported, that is; it is not possible to make a subscription and 'point' the messages at a queue. This means that messages must be consumed from the object handle returned on the subscription call. On the client, the messages are stored in a message buffer until they are consumed by the client; see [MessageBuffer stanza of the client configuration file](#) for more information. If the client does not keep up with the publishing rate, the messages are discarded as required, with the oldest messages discarded first.

It is normally an administration decision whether an application uses Multicast or not, specified by setting the MCAST attribute of a TOPIC object. If a publishing application must ensure that multicast is not used, it can use the MQ00_NO_MULTICAST option. Similarly, a subscribing application can ensure that multicast is not used by subscribing with the MQSO_NO_MULTICAST option.

WebSphere MQ Multicast supports the use of message selectors. A selector is used by an application to register its interest in only those messages with properties that satisfy the SQL92 query that the selection string represents. For more information about message selectors, see “[Selectors](#)” on page 22.

The following table lists all the major MQI concepts and how they relate to Multicast:

Table 71. MQI concepts and how they relate to multicast		
MQI Concept	Action when tried using multicast	Reason code
Putting a zero length message	Rejected	2005 (07D5) (RC2005): MQRC_BUFFER_LENGTH_ERROR
Grouping	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
Segmentation	Rejected	2443 (098B) (RC2443): MQRC_SEGMENTATION_NOT_ALLOWED
Distribution lists	Rejected	2154 (086A) (RC2154): MQRC_RECS_PRESENT_ERROR

Table 71. MQI concepts and how they relate to multicast (continued)

MQI Concept	Action when tried using multicast	Reason code
MQINQ	Rejected for topics handles: MQINQ and MQSET of topics is not supported.	2038 (07F6) (RC2038): MQRC_NOT_OPEN_FOR_INQUIRE
	Accepted for managed handle. Only Current Depth can be inquired.	<ul style="list-style-type: none"> • If the value is Current Depth, then there is no applicable reason code. • If the value is anything other than Current Depth, the reason code is 2067 (0813) (RC2067): MQRC_SELECTOR_ERROR.
MQSET	Rejected for all handles.	2040 (07F8) (RC2040): MQRC_NOT_OPEN_FOR_SET
Transactions (XA or not)	Rejected	2072 (0818) (RC2072): MQRC_SYNCPOINT_NOT_AVAILABLE
Message browse	Rejected	2036 (07F4) (RC2036): MQRC_NOT_OPEN_FOR_BROWSE
Lock messages	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
Browse with mark	Rejected	2036 (07F4) (RC2036): MQRC_NOT_OPEN_FOR_BROWSE
Pass context	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
MQPUT1	Rejected. It is invalid to try and MQPUT1 to a Multicast only topic.	2560 (0A00) (RC2560): MQRC_MULTICAST_ONLY
Durable subscription	Rejected if the topic is marked as "Multicast only", otherwise a non-Multicast subscription is made.	2436 (0984) (RC2436): MQRC_DURABILITY_NOT_ALLOWED
TopicString > 255	Rejected. If the topic string is greater than 255 characters, it is rejected in the client.	2425 (0979) (RC2425): MQRC_TOPIC_STRING_ERROR

Table 71. MQI concepts and how they relate to multicast (continued)		
MQI Concept	Action when tried using multicast	Reason code
Non-managed subscription made	Rejected if the topic is marked as "Multicast only", otherwise a non-Multicast subscription is made.	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR
MQPMO_NOT_OWN_SUBS	Rejected	2046 (07FE) (RC2046): MQRC_OPTIONS_ERROR

The following items expand on some of the MQI concepts from the previous table, and provides information on some of the MQI concepts that are not in the table:

Message persistence

For nondurable multicast subscribers, persistent messages from the publisher are delivered in an unrecoverable fashion.

Message truncation

Message truncation is supported, which means that it is possible for an application to:

1. Issue an MQGET.
2. Get MQRC_TRUNCATED_MSG_FAILED.
3. Allocate a larger buffer.
4. Reissue the MQGET to retrieve the message.

Subscription expiry

Subscription expiry is not supported. Any attempt to set an expiry is ignored.

High availability for multicast

Use this information to understand WebSphere MQ Multicast continuous peer-to-peer operation; although WebSphere MQ connects to a WebSphere MQ queue manager, messages do not flow through that queue manager.

Although a connection to a queue manager must be made in order to MQOPEN or MQSUB the multicast topic object, the messages themselves do not flow through the queue manager. Therefore, after the MQOPEN or MQSUB is completed on the multicast topic object, it is possible to continue transmitting multicast messages even if the connection to the queue manager has been lost. There are two modes of operation:

A normal connection is made to the queue manager

Multicast communication is possible while the connection to the queue manager exists. If the connection fails, the normal MQI rules are applied, for example; an MQPUT to the multicast object handle returns 2009 (07D9) (RC2009): MQRC_CONNECTION_BROKEN.

A reconnecting client connection is made to the queue manager

Multicast communication is possible even during the reconnection cycle. This means that even when the connection to the queue manager has been broken, the putting and consuming of multicast messages is not affected. The client attempts to reconnect to a queue manager, and if that reconnection fails, the connection handle becomes broken and all MQI calls, including multicast ones, fail. For more information, see: [Automated client reconnection](#)

If any application explicitly issues an MQDISC, then all multicast subscriptions and object handles are closed.

Multicast continuous peer-to-peer operation

One of the advantages of peer-to-peer communication between the clients is that the messages do not need to flow through the queue manager; therefore if the connection to the queue manager breaks, message transfer continues. The following restrictions apply to the continuous message requirements of this mode:

- The connection must be made using one of the MQCNO_RECONNECT_* options for continuous operation. This process means that although the communications session might be broken, the actual connection handle is not broken, and is in the reconnecting state instead. If reconnection fails, the connection handle is now broken which prevents all further MQI calls.
- Only MQPUT, MQGET, MQINQ, and Async Consume are supported in this mode. Any MQOPEN, MQCLOSE, or MQDISC verbs require reconnection to the queue manager to complete.
- Status flows to the queue manager stop; any state in the queue manager might therefore be stale or missing. This means that the clients might be sending and receiving messages and there is no status known on the queue manager. For more information, see: [Multicast application monitoring](#)

Data conversion in the MQI for multicast messaging

Use this information to understand how data conversion works for WebSphere MQ Multicast messaging.

WebSphere MQ Multicast is a shared, connectionless protocol, and so it is not possible for each client to make specific requests for data conversion. Every client subscribed to the same multicast stream receives the same binary data; therefore, if WebSphere MQ data conversion is required, the conversion is performed locally at each client.

Data is converted on the client for WebSphere MQ Multicast traffic. If the **MQGMO_CONVERT** option is specified, data conversion is done as requested. User defined formats need the data conversion exit installed on the client; see [“Writing data-conversion exits” on page 397](#) for information about which libraries are now in the client and server packages.

For information about administering data conversion, see [Enabling data conversion for Multicast messaging](#).

For more information about data conversion, see [Data conversion](#).

For more information about data conversion exits and ClientExitPath, see [ClientExitPath stanza of the client configuration file](#).

Multicast exception reporting

Use this information to learn about WebSphere MQ Multicast event handlers and reporting WebSphere MQ Multicast exceptions.

WebSphere MQ Multicast assists with problem determination by calling the event handler to report multicast events which are reported using the standard WebSphere MQ event handler mechanism.

An individual Multicast event can result in more than one WebSphere MQ event being called because there might be multiple MQHCONN connection handles using the same multicast transmitter or receiver. However, each multicast exception causes only one event handler to be called per WebSphere MQ connection.

The WebSphere MQ MQCBDO_EVENT_CALL constant enables applications to register a callback to receive only WebSphere MQ events, and the MQCBDO_MC_EVENT_CALL enable applications to register a callback to receive only multicast events. If both constants are used, both types of event are received.

Requesting Multicast events

WebSphere MQ Multicast events use the MQCBDO_MC_EVENT_CALL constant in the cbd.Options field. The following example demonstrates how to request multicast events:

```
cbd.CallbackType = MQCBT_EVENT_HANDLER;  
cbd.Options = MQCBDO_MC_EVENT_CALL;
```

```
cbd.CallbackFunction = EventHandler;
MQCB(Hcon,MQOP_REGISTER,&cbd,MQHO_UNUSABLE_HOBJ,NULL,NULL,&CompCode,&Reason);
```

When the MQCBDO_MC_EVENT_CALL option is specified for the `cbd.Options` field, the event handler is sent only WebSphere MQ Multicast events instead of connection level events. To request that both types of events are sent to the event handler, the application must specify the MQCBDO_EVENT_CALL constant in the `cbd.Options` field as well as the MQCBDO_MC_EVENT_CALL constant as shown in the following example:

```
cbd.CallbackType      = MQCBT_EVENT_HANDLER;
cbd.Options           = MQCBDO_EVENT_CALL | MQCBDO_MC_EVENT_CALL
cbd.CallbackFunction = EventHandler;
MQCB(Hcon,MQOP_REGISTER,&cbd,MQHO_UNUSABLE_HOBJ,NULL,NULL,&CompCode,&Reason);
```

If neither of these constants is used, only connection level events are sent to the event handler.

For more information about values for the `Options` field see [Options \(MQLONG\)](#).

Multicast event format

WebSphere MQ Multicast exceptions include some supporting information which is returned in the **Buffer** parameter of the callback function. The **Buffer** pointer points to an array of pointers and the MQCBC.DataLength field specifies the size, in bytes, of the array. The first element of the array always points to a short text description of the event. More parameters might be supplied depending on the type of event. The following table lists the exceptions:

Table 72. Multicast event code descriptions		
Event code	Description	Additional data
MQMCEV_PACKET_LOSS	Unrecoverable packet loss	Number of lost packets
MQMCEV_HEARTBEAT_TIMEOUT	Long absence of heartbeat control packet	N/A
MQMCEV_VERSION_CONFLICT	Reception of newer protocol version packets	N/A
MQMCEV_RELIABILITY	Different reliability modes of the transmitter and the receiver	N/A
MQMCEV_CLOSED_TRANS	Topic transmission is closed by 1 source	N/A
MQMCEV_STREAM_ERROR	Error detected on stream	N/A
MQMCEV_NEW_SOURCE	A new source starts to transmit on the topic	Source structure
MQMCEV_RECEIVE_QUEUE_TRIMMED	Packets removed from PacketQ due to time or space expiration	Number of trimmed packets
MQMCEV_PACKET_LOSS_NACK_EXPIRE	Unrecoverable packet loss due to NACK expiration	Number of lost packets
MQMCEV_ACK_RETRIES_EXCEEDED	Packets removed from history after max_ack_retries was exceeded	Number of packets removed
MQMCEV_STREAM_SUSPEND_NACK	NACKs have been suspended on a stream accepted by this topic	Suspend stream ID
		Time in milliseconds that the stream is suspended for

Table 72. Multicast event code descriptions (continued)		
Event code	Description	Additional data
MQMCEV_STREAM_RESUME_NACK	NACKs have been resumed after they have been suspended on a stream	Stream ID
MQMCEV_STREAM_EXPELLED	A stream accepted by this topic has been rejected due to an expel request	Stream ID
MQMCEV_FIRST_MESSAGE	First message from a source	Message number
MQMCEV_LATE_JOIN_FAILURE	Failed to start late join session	N/A
MQMCEV_MESSAGE_LOSS	Unrecoverable message loss	Number of lost messages
MQMCEV_SEND_PACKET_FAILURE	Multicast transmitter failed to send a multicast packet	N/A
MQMCEV_REPAIR_DELAY	Multicast receiver did not receive a repair packet for an outstanding NAK	N/A
MQMCEV_MEMORY_ALERT_ON	Receiver reception buffers are filling up	Buffer pool utilization percentage
MQMCEV_MEMORY_ALERT_OFF	Receiver reception buffers are down to normal	Buffer pool utilization percentage
MQMCEV_NACK_ALERT_ON	Receiver repair packet request rate reached high water mark	Current repair request rate in packets per second
MQMCEV_NACK_ALERT_OFF	Receiver repair packet request rate is down to normal	Current repair request rate in packets per second
MQMCEV_REPAIR_ALERT_ON	Transmitter repair packet send rate reached high water mark	N/A
MQMCEV_REPAIR_ALERT_OFF	Transmitter repair packet send rate is down to normal	N/A
MQMCEV_SHM_DEST_UNUSABLE	The Shared Memory region used by a transmitter topic destination has been detected to be unusable	N/A
MQMCEV_SHM_PORT_UNUSABLE	The Shared Memory port used by a receiver instance has been detected to be unusable	N/A
MQMCEV_CCT_GETTIME_FAILED	The get time from Coordinated Cluster Time failed	N/A
MQMCEV_DEST_INTERFACE_FAILURE	The network interface used by a transmitter topic destination has failed and a backup network interface is unavailable	
MQMCEV_DEST_INTERFACE_FAILOVER	The network interface used by a transmitter topic destination has failed and a successful failover to another Interface has been completed	

Table 72. Multicast event code descriptions (continued)		
Event code	Description	Additional data
MQMCEV_PORT_INTERFACE-FAILURE	The network interface used by a receiver rmmPort has failed and a backup network interface is unavailable (or has also failed)	RMM configuration
MQMCEV_PORT_INTERFACE_FAILOVER	The network interface used by a receiver rmmPort has failed and a successful failover to another Interface has been completed	RMM configuration

Using .NET

WebSphere MQ classes for .NET allow a program written in the .NET programming framework to connect to WebSphere MQ as a WebSphere MQ MQI client or to connect directly to a WebSphere MQ server.

If you have applications which use Microsoft's .NET Framework and want to take advantage of the facilities of WebSphere MQ, you must use WebSphere MQ classes for .NET.

The object-oriented WebSphere MQ .NET interface is different from the MQI interface in that it uses methods of objects rather than using the MQI verbs.

The procedural WebSphere MQ application programming interface is built around verbs such as those in the following list:

```
MQCONN, MQDISC, MQOPEN, MQCLOSE,
MQINQ, MQSET, MQGET, MQPUT, MQSUB
```

These verbs all take, as a parameter, a handle to the WebSphere MQ object on which they are to operate. Because .NET is object-oriented, the .NET programming interface turns this round. Your program consists of a set of WebSphere MQ objects, which you act upon by calling methods on those objects. You can write programs in any language supported by .NET.

When you use the procedural interface, you disconnect from a queue manager by using the call `MQDISC(Hconn, CompCode, Reason)`, where *Hconn* is a handle to the queue manager.

In the .NET interface, the queue manager is represented by an object of class `MQQueueManager`. You disconnect from the queue manager by calling the `Disconnect()` method on that class.

```
// declare an object of type queue manager
MQQueueManager queueManager=new MQQueueManager();
...
// do something...
...
// disconnect from the queue manager
queueManager.Disconnect();
```

WebSphere MQ classes for .NET is a set of classes that enable .NET applications to interact with WebSphere MQ. They represent the various components of WebSphere MQ which your application uses, such as queue managers, queues, channels and messages. For details of these classes, see [The WebSphere MQ .NET classes and interfaces](#).

Before you can compile any applications that you write, you must have a .NET Framework installed. For instructions on installing the WebSphere MQ classes for .NET and the .NET Framework, see [“Installing WebSphere MQ classes for .NET”](#) on page 536.

Related concepts

[Technical overview](#)

[“Connection options”](#) on page 535

There are three modes of connecting WebSphere MQ classes for .NET to a queue manager. Consider which type of connection best suits your requirements.

[“Using WebSphere MQ classes for .NET” on page 546](#)

This collection of topics describes how to configure your system to run the sample programs to verify your WebSphere MQ classes for .NET installation, and how to run your own programs.

[“Solving WebSphere MQ .NET problems” on page 548](#)

If a program does not complete successfully, run one of the sample applications, and follow the advice given in the diagnostic messages.

[“Writing and deploying WebSphere MQ .NET programs” on page 549](#)

To use WebSphere MQ classes for .NET to access WebSphere MQ queues, you write programs in any language supported by .NET containing calls that put messages onto, and get messages from, WebSphere MQ queues.

[“IBM WebSphere MQ custom channel for Microsoft Windows Communication Foundation \(WCF\)” on page 568](#)

The Microsoft Windows Communication Foundation (WCF) custom channel for IBM WebSphere MQ sends and receives messages between WCF clients and services.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Developing applications” on page 7](#)

IBM WebSphere MQ provides several ways in which you can develop applications to send and receive messages that you need to support your business processes. You can also develop applications to manage your queue managers and related resources.

Getting started with WebSphere MQ classes for .NET

WebSphere MQ classes for .NET allow a program written in the .NET programming framework to connect to WebSphere MQ as a WebSphere MQ MQI client or to connect directly to a WebSphere MQ server.

Connection options

There are three modes of connecting WebSphere MQ classes for .NET to a queue manager. Consider which type of connection best suits your requirements.

Client bindings connection

To use WebSphere MQ classes for .NET as a WebSphere MQ MQI client, you can install it, with the WebSphere MQ MQI client, either on the WebSphere MQ server machine, or on a separate machine. A client bindings connection can use XA or non-XA transactions

Server bindings connection

When used in server bindings mode, WebSphere MQ classes for .NET use the queue manager API, rather than communicating through a network. This provides better performance for WebSphere MQ applications than using network connections.

To use the bindings connection, you must install WebSphere MQ classes for .NET on the WebSphere MQ server.

Managed client connection

A connection made in this mode connects as a WebSphere MQ client to a WebSphere MQ server running either on the local or a remote machine.

The WebSphere MQ classes for .NET connecting in this mode remain in .NET managed code and make no calls to native services. For more information about managed code, refer to Microsoft documentation.

There are a number of limitations to using the managed client. For more information about these, see [“Managed client connections” on page 549](#).

Installing WebSphere MQ classes for .NET

WebSphere MQ classes for .NET, including samples, is installed with WebSphere MQ. There is a prerequisite of Microsoft .NET Framework.

The latest version of WebSphere MQ classes for .NET is installed by default as part of the standard WebSphere MQ installation in the *Java and .NET Messaging and Web Services* feature. For installation instructions, see [Installing IBM WebSphere MQ server on Windows](#) or [Installing an IBM WebSphere MQ client on Windows systems](#).

In a multiple installation environment, if you have previously installed the WebSphere MQ classes for .NET as a support pack, you cannot install WebSphere MQ unless you first uninstall the support pack. The WebSphere MQ classes for .NET feature that is installed with WebSphere MQ contains the same functionality as the support pack.

Sample applications, including source files, are also supplied; see [“Sample applications” on page 546](#).

To run WebSphere MQ classes for .NET on 32-bit or 64-bit platforms you must have installed Microsoft .NET Framework V2.0 or later.

Note: If the Microsoft .NET Framework v2.0 or higher is not installed before installing WebSphere MQ V7.0.1, then the WebSphere MQ product installation will continue without error, but the WebSphere MQ classes for .NET will not be available. If the .NET Framework is installed after installing WebSphere MQ 7.0.1, then the WebSphere .NET assemblies must be registered by running the `WMQInstallDir\bin\amqiRegisterdotNet.cmd` script, where `WMQInstallDir` is the directory where WebSphere MQ 7.0.1 is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of `amqi*.log` files recording the actions taken are created in the `%TEMP%` directory.

For information about using the WebSphere MQ custom channel for the Microsoft WCF with .NET 3, see: [“IBM WebSphere MQ custom channel for Microsoft Windows Communication Foundation \(WCF\)” on page 568](#)

Distributed transactions in .NET

Distributed transactions or global transactions allows client applications to include several different sources of data on two or more networked systems in one transaction.

In distributed transactions, a transaction manager coordinates and manages the transaction among two or more resource managers.

Transactions can be either single phase or two-phase commit process. The single-phase commit is a process where only one resource manager participates in the transaction and two-phase commit process is where there are more than one resource manager participating in the transaction. In the two-phase commit process, the transaction manager sends a prepare call to check whether all the resource managers are prepared to commit. When it receives the acknowledgment from all the resource managers, the commit call is issued. Else, a rollback on the whole transaction happens. See [Transaction management and support](#) for more details. The resource managers should inform the transaction managers of their participation in the transaction. When the resource manager informs the transaction manager of its participation, the resource manager gets callbacks from the transaction manager when the transaction is going to commit or roll back.

WebSphere MQ .NET classes already supports distributed transactions in unmanaged and server bindings mode connections. In these modes, WebSphere MQ .NET classes delegates all its calls to C extended transaction client, which manages the transaction processing on behalf of .NET.

WebSphere MQ.NET classes now support distributed transactions in managed mode where WebSphere MQ .NET Classes uses `System.Transactions` namespace for the distributed transactions support. The `System.Transactions` infrastructure makes transactional programming simple and efficient by supporting the transactions initiated in all the resource managers including WebSphere MQ. The WebSphere MQ .NET application can put and get messages using .NET implicit transaction programming or explicit transaction

programming model. In implicit transactions, the transaction boundaries are created by the application program that decides when to commit, rollback (for explicit transactions) or complete the transaction. In explicit transactions, you have to explicitly specify whether you want to commit, roll back, and complete the transaction.

WebSphere MQ.NET uses Microsoft distributed transaction coordinator (MS DTC) as the transaction manager, which coordinates and manages the transaction between multiple resource managers. WebSphere MQ is used as the resource manager.

WebSphere MQ.NET follows the X/Open Distributed Transaction Processing (DTP) model. The X/Open Distributed Transaction Processing model is a distributed transaction processing model proposed by the Open Group, a vendor consortium. This model is a standard among most of the commercial vendors in the transaction processing and database domains. Most of the commercial transaction management products support the X/DTP model.

Modes of transaction

- [“Distributed transactions in managed mode” on page 538](#)
- [Distributed transactions for unmanaged mode](#)

Coordinating transactions in various scenarios

- A connection might participate in several transactions, but only one transaction is active at any point of time.
- During a transaction, the MQQueueManager.Disconnect call is honored. In this case the transaction is asked to roll back.
- During a transaction, the MQQueue.Close or MQTopic.Close call is honored. In this case transaction is asked to roll back.
- The transaction boundaries are created by the application program that decides when to commit, rollback (for explicit transactions) or complete (for implicit transactions) the transaction.
- If the client application breaks during a transaction with an unexpected error before issuing a Put or Get call on a queue or topic call, the transaction is rolled back and a MQException is thrown.
- If MQCC_FAILED reason code is returned during a Put or Get call on a queue or Topic call, an MQException is thrown with reason code and the transaction is rolled. If a prepare call has been already issued by the transaction manager, then WebSphere MQ .NET returns the prepare request by forcibly rolling back the transaction. Then the transaction manager DTC causes a rollback on current work with all the resource managers in current ambient transactions.
- During a transaction involving multiple resource managers if some environmental reason causes the Put or Get call to hang indefinitely, the transaction manager waits until a stipulated time. After the time is out, it causes the rollback of all current work with all the resource managers in current ambient transactions. If this indefinite wait happens during the prepare phase, the transaction manager might timeout or issue an in-doubt call on the resource in which case the transaction is rolled back.
- Applications using transactions must Put or Get messages under SYNC_POINT. If a message Put or Get call is issued under a transactional context that is not under SYNC_POINT, the call fails with MQRC_UNIT_OF_WORK_NOT_STARTED reason code.

Behavioral differences between Managed and Unmanaged Client transaction support using Microsoft .NET System.Transactions namespace

Nested Transactions have a TransactionScope inside another TransactionScope

- WebSphere MQ .NET fully managed client does support nested TransactionScope
- WebSphere MQ .NET unmanaged client does not support nested TransactionScope

Dependent Transactions from System.Transactions

- WebSphere MQ .NET fully managed client does support the dependent transactions facility provided by System.Transactions.
- WebSphere MQ .NET unmanaged client does not support the dependent transactions facility provided by System.Transactions.

Product Samples

New product samples SimpleXAPut, and SimpleXAGet are available under WebSphere MQ\tools\dotnet\samples\cs\base . The samples are C# applications, which demonstrate using MQPUT and MQGET under Distributed Transactions using SystemTransactions namespace. For more information about these samples, see [“Creating simple put and get messages within a TransactionScope”](#) on page 541

Distributed transactions in managed mode

WebSphere MQ .NET classes use System.Transactions namespace for the distributed transactions support in managed mode. In the managed mode, MS DTC coordinates and manages distributed transactions across all the servers enlisted in a transaction.

WebSphere MQ .NET classes provide an explicit programming model based on the System.Transactions.Transaction class and an implicit programming model using the System.Transactions.TransactionScope, class where the transactions are automatically managed by the infrastructure.

Implicit Transaction

The following piece of code describes how a WebSphere MQ .NET application puts a message using .NET implicit transaction programming.

```
Using (TransactionScope scope = new TransactionScope ())
{
    Q.Put (putMsg,pmo);
    scope.Complete ();
}

Q.close();
qMgr.Disconnect();}
```

Explanation of the code flow of implicit transaction

The code creates *TransactionScope* and puts the message under the scope. It then calls *Complete* to inform the transaction coordinator of the completion of the transaction. The transaction coordinator now issues *prepare* and *commit* to complete the transaction. If an issue is detected, then a *rollback* is called.

Explicit Transaction

The following code describes how a WebSphere MQ .NET application puts messages using .NET explicit transaction programming model.

```
MQQueueManager qMgr = new MQQueueManager ("MQQM");
MQQueue Q = QMGR.AccessQueue("Q", MQC.MQOO_OUTPUT+MQC.MQOO_INPUT_SHARED);
MQPutMessageOptions pmo = new MQPutMessageOptions();
pmo.Options = MQC.MQPMO_SYNCPOINT;
MQMessage putMsg1 = new MQMessage();
Using(CommittableTransaction tx = new CommittableTransaction()){
    Transaction.Current = tx;
    try
    {
        Q.Put(MSG,pmo);
        tx.commit();
    }
    catch(Exception)
    {tx.rollback();}
}

Q.close();
qMgr.Disconnect();
}
```

Explanation of the code flow of explicit transaction

The piece of code creates transaction using *CommitableTransaction* class. It puts a message under that scope and then explicitly calls *commit* to complete the transaction. If there are any issues *rollback* is called.

Distributed transactions in unmanaged mode

WebSphere MQ.NET classes support unmanaged connections (client) using extended transaction client and COM+/MTS as the transaction coordinator, using either implicit or explicit transaction programming model. In the unmanaged mode, WebSphere MQ .NET classes delegate all its calls to C extended transaction client that manages the transaction processing on behalf of .NET.

The transaction processing is controlled by an external transaction manager, coordinating the global unit of work under the control of the API of the transaction manager. The MQBEGIN, MQCMIT, and MQBACK verbs are unavailable. WebSphere MQ .NET classes expose this support by way of its unmanaged transport mode (C client). See [Configuring XA-compliant transaction managers](#)

MTS is evolved as a transaction processing (TP) system to provide the same features on Windows NT as available in CICS, Tuxedo, and on other platforms. When the MTS is installed, a separate service is added to Windows NT called the Microsoft Distributed Transaction Coordinator (MSDTC). The MSDTC coordinates the transactions that span separate data stores or resources. To work, it requires each data store to implement its own proprietary resource manager.

WebSphere MQ becomes compatible with MSDTC by implementing an interface (proprietary resource manager interface) where it manages to map DTC XA calls to WebSphere MQ(X/Open) calls. WebSphere MQ plays the role of a resource manager.

When a component such as COM+ requests access to a WebSphere MQ, the COM usually checks with the appropriate MTS context object if a transaction is required. If a transaction is required, the COM informs the DTC and automatically starts an integral WebSphere MQ transaction for this operation. Then the COM works with the data through the MQMTS software, putting and getting messages as required. The object instance obtained from the COM calls the SetComplete or SetAbort method after all the actions on the data are over. When the application issues SetComplete, the call signals the DTC that the application has completed the transaction and the DTC can go ahead with the two-phase commit process. The DTC then issues calls to MQMTS which in turn issues calls to WebSphere MQ to commit or roll back the transaction.

Writing a WebSphere MQ .NET application using unmanaged client

To run within the context of COM+, a .NET class must inherit from *System.EnterpriseServices.ServicedComponent*. The rules and recommendations to create assemblies that use serviced components are the following:

Note: The following steps are relevant only if you are using *System.EnterpriseServices* mode.

- The class and method being started in COM+ must both be public (no internal classes, and no protected or static methods).
- The class and method attributes: The *TransactionOption* attribute dictates the transaction level of the class, that is whether the transactions are disabled, supported, or required. The *AutoComplete* attribute on the *ExecuteUOW()* method instructs COM+ to commit the transaction if no unhandled exception is thrown.
- Strong-naming an assembly: The assembly must be strong-named and registered in the Global Assembly Cache (GAC). The assembly is registered in COM+ explicitly or by lazy registration after it is registered in the GAC.
- Registering an assembly in COM+: Prepare the assembly to be exposed to COM clients. Then create a type library by using the Assembly Registration tool, *regasm.exe*.

```
regasm UnmanagedToManagedXa.dll
```

- Register the assembly into GAC `gacutil /i UnmanagedToManagedXa.dll`.
- Register the assembly in COM+ by using the .NET services installer tool, *regsvcs.exe*. See the type library created by *regasm.exe*:

```
Regsvcs /appname:UnmanagedToManagedXa /tlb:UnmanagedToManagedXa.tlb
UnmanagedToManagedXa.dll
```

- The assembly is deployed into the GAC, and later it is registered in COM+ by lazy registration. The .NET framework takes care of the registration after the code is run for the first time.

The example code flow using System.EnterpriseServices model and System.Transactions with COM+ are described in the following sections:

Example code flow using System.EnterpriseServices model

```
using System;
using IBM.WMQ;
using IBM.WMQ.Nmqi;
using System.Transactions;
using System.EnterpriseServices;

namespace UnmanagedToManagedXa
{
    [ComVisible(true)]
    [System.EnterpriseServices.Transaction(System.EnterpriseServices.TransactionOption.Required)]
    public class MyXa : System.EnterpriseServices.ServicedComponent
    {
        public MQQueueManager QMGR = null;
        public MQQueueManager QMGR1 = null;
        public MQQueue QUEUE = null;
        public MQQueue QUEUE1 = null;
        public MQPutMessageOptions pmo = null;
        public MQMessage MSG = null;

        public MyXa()
        {
        }

        [System.EnterpriseServices.AutoComplete()]
        public void ExecuteUOW()
        {
            QMGR = new MQQueueManager("usemq");

            QUEUE = QMGR.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE",
                                     MQC.MQOO_INPUT_SHARED +
                                     MQC.MQOO_OUTPUT +
                                     MQC.MQOO_BROWSE);

            pmo = new MQPutMessageOptions();
            pmo.Options = MQC.MQPMO_SYNCPOINT;
            MSG = new MQMessage();
            QUEUE.Put(MSG, pmo);
            QMGR.Disconnect();
        }
    }

    public void RunNow()
    {
        MyXa xa = new MyXa();
        xa.ExecuteUOW();
    }
}
```

Example code flow using System.Transactions for interactions with COM+

```
[STAThread]
public void ExecuteUOW()
{
    Hashtable t1 = new Hashtable();
    t1.Add(MQC.CHANNEL_PROPERTY, "SYSTEM.DEF.SVRCONN");
    t1.Add(MQC.HOST_NAME_PROPERTY, "localhost");
    t1.Add(MQC.PORT_PROPERTY, 1414);
    t1.Add(MQC.TRANSPORT_PROPERTY, MQC.TRANSPORT_MQSERIES_CLIENT);
    TransactionOptions opts = new TransactionOptions();

    using(TransactionScope scope = new TransactionScope(TransactionScopeOption.RequiresNew,
                                                         opts,
                                                         EnterpriseServicesInteropOption.Full))
    {
    }
}
```

```

QMGR = new MQQueueManager("usemq", t1);
QUEUE = QMGR.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE",
                          MQC.MQOO_INPUT_SHARED +
                          MQC.MQOO_OUTPUT +
                          MQC.MQOO_BROWSE);

pmo = new MQPutMessageOptions();
pmo.Options = MQC.MQPMO_SYNCPOINT;
MSG = new MQMessage();
QUEUE.Put(MSG, pmo);
scope.Complete();
}
QMGR.Disconnect();
}

```

Creating simple put and get messages within a TransactionScope

Product sample C# applications are available within WebSphere MQ. These simple applications demonstrate putting and getting messages within a TransactionScope. At the end of the task, you will be able to put and get messages from a queue or topic.

Before you begin

MSDTC service must be running and enabled for XA Transactions.

About this task

The example is a simple application, SimpleXAPut and SimpleXAGet. The programs SimpleXAPut and SimpleXAGet are C# applications available within WebSphere MQ. SimpleXAPut demonstrates using MQPUT, under Distributed Transactions using SystemTransactions namespace. SimpleXAGet demonstrates using MQGET, under Distributed Transactions using SystemTransactions namespace.

SimpleXAPut is located in WebSphere MQ\tools\dotnet\samples\cs\base

Procedure

The applications can be run with the command-line parameters from tools\dotnet\samples\cs\base\bin

```
SimpleXAPut.exe -d destinationURI [-h host -p port -l channel -tx transaction -tm mode -n
numberOfMsgs]
```

```
SimpleXAGet.exe -d destinationURI [-h host -p port -l channel -tx transaction -tm mode -n
numberOfMsgs]
```

where the parameters are:

-destinationURI

This can be queue or topic. For a queue, specify as queue://queueName and for a topic specify as topic://topicName.

-host

This can be a hostname such as localhost or an IP address.

-port

The port on which the queue manager is running.

-channel

The connection channel being used. The default is SYSTEM.DEF.SVRCONN

-transaction

The transaction result, for example commit or rollback.

-mode

The transport mode, for example managed or unmanaged.

-numberOfMsgs

The number of messages. The default is 1.

Example

```
SimpleXAPut -d topic://T01 -h localhost -p 2345 -tx rollback -tm unmanaged
```

```
SimpleXAGet -d queue://Q01 -h localhost -p 2345 -tx rollback -tm unmanaged
```

Recovering Transactions

This section describes the process of recovering transactions in WebSphere MQ .NET XA using managed mode.

Overview

In distributed transaction processing, the transactions can be successfully completed. But, there can be scenarios where a transaction may fail for many reasons. These reasons could include a system failure, hardware failure, network error, incorrect or invalid data, application errors or natural or man-made disasters. It is not possible to prevent transaction failures. The distributed transaction system must be capable of handling these failures. It must be able to detect and correct errors when they occur. This process is known as Transaction Recovery.

An important aspect of the Distributed Transaction Processing is to recover the incomplete or in doubt transactions. It is essential to run the recovery as the Unit of Work part of a particular transaction is held locked until it is recovered. Microsoft .NET from its System.Transactions class library provides the option for recovering incomplete/in-doubt transactions. This recovery support expects Resource Manager to maintain the transaction logs and run the recovery when in need.

Recovery Model

In Microsoft .NET's transaction recovery model, the Transaction Manager (System.Transactions, or Microsoft Distributed Transaction coordinator (MS DTC), or both), initiates, coordinates, and controls the transaction recovery. The OLE Tx Protocol (Microsoft's XA protocol) based Resource Managers provide the options to configure the DTC to drive, coordinate, and control the recovery for them. To do this, Resource Managers must register XA_Switch with MS DTC by using native interface.

XA_Switch provides the entry points of XA functions like xa_start, xa_end, and xa_recover in the Resource Manager to the Distributed Transaction Coordinator.

Recovery using Microsoft Distributed Transaction coordinator (DTC):

Microsoft Distributed Transaction coordinator provides two kinds of recovery processes.

Cold Recovery

Cold recovery is performed if the transaction manager process fails while a connection to a XA resource manager is open. When the transaction manager restarts, it reads the transaction manager logs and re-establishes the connection to the XA resource manager and then initiates recovery.

Hot Recovery

Hot recovery is performed if the transaction manager remains up while the connection between the transaction manager and the XA resource manager fails because the XA resource manager or the network fails. After the failure, the transaction manager periodically attempts to reconnect to the XA resource manager. When the connection is re-established, the transaction manager initiates XA recovery.

System.Transactions namespace provides managed implementation of Distributed transactions that are based on MS DTC as the transaction manager. It provides similar features as that of MS DTC's native interface but in fully managed environment. The only difference is about the transaction recovery. System.Transactions expects Resource Managers to drive the recovery by themselves and then coordinate with the Transaction Managers (MS DTC). Resource Manager must ask for recovery of a particular incomplete transaction and then Transaction Manager accepts it and coordinates based on the actual outcome of that particular transaction.

Transaction recovery process for WebSphere MQ .NET

This section describes how distributed transactions can be recovered with WebSphere MQ .NET classes.

Overview

To recover an incomplete transaction the recovery information is required. The transaction recovery information must be logged to storage by the resource managers. WebSphere MQ .NET classes follow a similar path. The transaction recovery information is logged to a system queue called SYSTEM.DOTNET.XARECOVERY.QUEUE.

Transaction recovery in WebSphere MQ .NET is a two stage process.

1. Logging of transaction recovery information.

- For every transaction, during the prepare phase a persistent message containing the recovery information is added to SYSTEM.DOTNET.XARECOVERY.QUEUE.
- The message is deleted if the commit call succeeds.

2. Recovering transactions using a monitor application WmqDotnetXAMonitor.

- WmqDotnetXAMonitor is a .NET managed application that processes messages in SYSTEM.DOTNET.XARECOVERY.QUEUE and recovers incomplete transactions

If the MCA is unable to put the message to the destination queue, it generates an exception report containing the original message, and puts it on a transmission queue to be sent to the reply-to queue specified in the original message. (If the reply-to queue is on the same queue manager as the MCA, the message is put directly to that queue, not to a transmission queue.)

SYSTEM.DOTNET.XARECOVERY.QUEUE

This is a system queue that holds transaction recovery information of incomplete transactions. This queue gets created when a queue manager is created.

Note: You should not delete SYSTEM.DOTNET.XARECOVERY.QUEUE queue.

WMQDotnetXAMonitor Application

WebSphere MQ .NET XA Monitor application monitors a given queue manager and recovers incomplete transactions if any. Following are deemed as incomplete transactions and are recovered:

Incomplete Transactions

- If the transaction is prepared but COMMIT did not complete within the timeout period.
- If the transaction is prepared but WebSphere MQ Queue manager has gone down.
- If the transaction is prepared but then Transaction Manager has gone down.

Monitor application must be run from the same system where your WebSphere MQ .NET client application is running. If there are applications running on multiple systems connecting to the same queue manager, the monitor application must be run from all the systems. Though each client machine has a monitor application running to recover the application, each monitor should be able to identify the message that corresponds to transaction that the current monitor's local MS DTC was coordinating so that it can reenlist and complete it.

Transaction recovery use cases for WebSphere MQ .NET

Following are the different usage scenarios:

- **WebSphere MQ Application using single DTC and single Queue Manager instance:** In this scenario, when you connect to the queue manager and run Unit of Work (UoW) under transaction, and if the transaction fails and becomes incomplete, the monitor application recovers the transaction and completes it.

In this scenario, there will be a single instance of monitor application running, as a single Queue manager is associated with the transactions.

- **Multiple WebSphere MQ Applications using single DTC and single Queue Manager instance:** In this scenario, there are more than one WMQ applications under single DTC and all are connecting to the same queue manager and running UoW under transactions.

If the transactions fail and become incomplete, the monitor application recovers them and completes the transactions pertaining to all the applications.

In this scenario, a single monitor application runs, as one queue manager is used in transactions.

- **Multiple WebSphere MQ Applications, multiple DTCs, different Queue Manager instances:** In this scenario, there are more than one WMQ applications under different DTCs (that is, each application is running on a different machine) and connecting to different queue managers.

If failure occurs and transaction becomes incomplete, monitor application checks the TransactionManagerWhereabouts in the message to determine the DTC address. If the TransactionManagerWhereabouts value matches with the DTC address under which the monitor is running, it completes the recovery, else, it continues to search until the message corresponding to its DTC is found.

In this scenario, there will be only one instance of monitor application running per client (user or computer) as each client has its own queue manager used in transactions.

- **Multiple WebSphere MQ Applications, multiple DTCs, multiple same Queue Managers instances:** In this scenario, there are more than one WMQ applications under different DTCs (that is each application is running on a different machine) and all are connecting to same queue manager.

If failure occurs and transaction becomes incomplete, monitor application verifies the TransactionManagerWhereabouts in the message to check if the DTC address and value match with the DTC under which the monitor is running. If both the values match, it completes the recovery else continues to search until it finds the message corresponding to its DTC.

In this scenario, there will be only a single instance of monitor application running per client (user or computer), as each client has its own queue manager association used in transactions.

- **Multiple WebSphere MQ Applications, single DTC, different Queue Manager instances:** In this scenario, there are more than one WMQ applications under a single DTC (that is, on a computer, there are more than one WMQ applications running) and connecting to different queue managers.

If transaction fails and becomes incomplete, monitor application recovers the transaction.

In this scenario, there will be as many instances of monitor application running as queue managers connected to, as each application has its own queue manager used in transactions and each of it must be recovered.

Note: If the monitor application is not running in the background, you can start it.

Using the WMQDotnetXAMonitor application

XA monitor application must be run manually. It can be started at any time. You can start it when you see the messages on the SYSTEM.DOTNET.XARECOVERY.QUEUE or you can keep it running in the background before you do any transactional work with the applications that are written using WebSphere MQ .NET classes.

Command to start the monitor application

```
WmqDotnetXAMonitor.exe -m <QueueManagerName> -n <ConnectionName> -c <Channel> -i
```


Where

- **n** - Connection name in host (port) format. Connection Name can contain more than one connection name. Multiple connection names must be given in comma separated list, for example "localhost (1414), localhost (1415), localhost (1416)". Monitor application run the recovery for each of the connection names specified in the comma separated list.
- **c** - Channel name.
- **m** - Queue manager name. Optional
- **i** - Heuristic branch completion. Optional

The monitor application performs the following actions:

1. Checks the queue depth of SYSTEM.DOTNET.XARECOVERY.QUEUE at an interval of 100 seconds.
2. If the queue depth is greater than zero the XA monitor browses the queue for messages and checks if the message satisfies the incomplete transaction criteria.
3. If any of the messages satisfy the incomplete transaction criteria, monitor pulls it out, and retrieves the transaction recovery information.
4. It then determines if the recovery information relates to the local MS DTC. If yes then it proceeds to recover the transaction. Otherwise it goes back to browse the next message.
5. It then makes calls to the queue manager to recover the incomplete transaction.

WmqDotNETXAMonitor application configuration file settings

To monitor the application, inputs can also be provided using the application configuration file. A sample application configuration file is shipped with WebSphere MQ .NET. This file can be modified according your requirements.

Application Configuration file takes the highest precedence while considering the input values. If input values are provided at both command line and Application Configuration file, then values from application configuration are considered.

Sample application configuration file.

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
<configSections>
<sectionGroup name="IBM.WMQ">
<section name="dnetxa" type="System.Configuration.NameValueFileSectionHandler" />
</sectionGroup>
</configSections>
<IBM.WMQ>
<dnetxa>
<add key="ConnectionName" value="" />
<add key="ChannelName" value="" />
<add key="QueueManagerName" value="" />
<add key="UserId" value="" />
<add key="SecurityExit" value="" />
<add key="SecurityExitUserData" value = "">
</dnetxa>
</dnetxa>
</configuration>
```

WmqDotNetXAMonitor Application log

Monitor Application creates a log file in the application directory for logging the Monitor's progress and transaction recovery status. Logging starts with the connection name and the channel details to show current queue manager for which the recovery is running.

Once the recovery starts, MessageId of the transaction recovery message, TransactionId of the incomplete transaction and actual outcome of the transaction as per Transaction Manager Coordination will be logged.

Sample log file:

```
Time|ProcessId|ThreadId|WMQ .NET XA Recovery Monitor, Running now for
ConnectionName:xxxx, Time|ProcessId|ThreadId|Channel=xxxx
Time|ProcessId|ThreadId|Current QueueDepth = n
```

```

Time|ProcessId|ThreadId|Current MessageId = xxxx
Time|ProcessId|ThreadId|Current Incomplete Transaction being recovered = xxxxx
Time|ProcessId|ThreadId|Actual Outcome of the transaction(as per DTC)= Commit/Rollback
Time|ProcessId|ThreadId|Recovery Completed for TransactionId= xxxxx
Time|ProcessId|ThreadId|Current QueueDepth = n
Time|ProcessId|ThreadId|Current MessageId = xxxx
Time|ProcessId|ThreadId|Current Incomplete Transaction being recovered = xxxxx
Time|ProcessId|ThreadId|Actual Outcome of the transaction(as per DTC)= Commit/Rollback
Time|ProcessId|ThreadId| Recovery Completed for TransactionId= xxxxx

```

Using WebSphere MQ classes for .NET

This collection of topics describes how to configure your system to run the sample programs to verify your WebSphere MQ classes for .NET installation, and how to run your own programs.

Configuring your queue manager to accept TCP/IP client connections

To configure a queue manager to accept incoming connection requests from the clients:

1. Define a server connection channel:
 - a. Start the queue manager.
 - b. Define a sample channel called NET.CHANNEL³:

```

DEF CHL('NET.CHANNEL') CHLTYPE(SVRCONN) TRPTYPE(TCP) MCAUSER(' ') +
DESCR('Sample channel for WebSphere MQ classes for .NET')

```

2. Start a listener:

```
runmqclsr -t tcp [-m qmname] [-p portnum]
```

Note: The square brackets indicate optional parameters; *qmname* is not required for the default queue manager, and the port number *portnum* is not required if you are using the default (1414).

Sample applications

To run your own .NET applications, use the instructions for the verification programs, substituting your application name in place of the sample applications.

Five sample applications are supplied:

- A put message application
- A get message application
- A 'hello world' application
- A publish/subscribe application
- An application using message properties

All these sample applications are supplied in the C# language, and some are also supplied in C++ and Visual Basic. You can write applications in any language supported by .NET.

"Put message" program SPUT (nmqsput.cs, mmqsput.cpp, vmqsput.vb)

This program shows how to put a message to a named queue. The program has three parameters:

- The name of a queue (required), for example, SYSTEM.DEFAULT.LOCAL.QUEUE
- The name of a queue manager (optional)
- The definition of a channel (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue manager name is given, the queue manager defaults to the default local queue manager.

If a channel is defined, it has the same format as the MQSERVER environment variable.

³ In this sample, we are not considering security implications. For a production system, consider using SSL or a security exit. See [Security](#) for more information.

"Get message" program SGET (nmqsget.cs, mmqsget.cpp, vmqsget.vb)

This program shows how to get a message from a named queue. The program has three parameters:

- The name of a queue (required), for example, SYSTEM.DEFAULT.LOCAL.QUEUE
- The name of a queue manager (optional)
- The definition of a channel (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue manager name is given, the queue manager defaults to the default local queue manager. If a channel is defined, it has the same format as the MQSERVER environment variable.

"Hello World" program (nmqwrld.cs, mmqwrld.cpp, vmqwrld.vb)

This program shows how to put and get a message. The program has three parameters:

- The name of a queue (optional), for example, SYSTEM.DEFAULT.LOCAL.QUEUE or SYSTEM.DEFAULT.MODEL.QUEUE
- The name of a queue manager (optional)
- A channel definition (optional), for example, SYSTEM.DEF.SVRCONN/TCP/hostname(1414)

If no queue name is given, the name defaults to SYSTEM.DEFAULT.LOCAL.QUEUE. If no queue manager name is given, the queue manager defaults to the default local queue manager.

"Publish/subscribe" program (MQPubSubSample.cs)

This program shows how to use WebSphere MQ publish/subscribe. It is supplied in C# only. The program has two parameters:

- The name of a queue manager (optional)
- A channel definition (optional)

"Message properties" program (MQMessagePropertiesSample.cs)

This program shows how to use message properties. It is supplied in C# only. The program has two parameters:

- The name of a queue manager (optional)
- A channel definition (optional)

You can verify your installation by compiling and running these applications.

The sample applications are installed to the following locations, according to the language in which they are written. *MQ_INSTALLATION_PATH* represents the high-level directory in which WebSphere MQ is installed.

C#

MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqswrld.cs

MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqspu.cs

MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\nmqsgt.cs

MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\MQPubSubSample.cs

MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\MQMessagePropertiesSample.cs

Managed C++

MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqswrld.cpp

MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqspu.cpp

MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\mmqsgt.cpp

Visual Basic

MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqswrld.vb

MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqspu.vb

MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\vmqsgt.vb

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqswrld.vb`

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqsput.vb`

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\xmqsget.vb`

To build the sample applications a batch file has been supplied for each language.

C#

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\cs\bldcssamp.bat`

The `bldcssamp.bat` file contains a line for each sample, which is all that is necessary to build this sample program:

```
csc /t:exe /r:System.dll /r:amqmdnet.dll /lib:MQ_INSTALLATION_PATH\bin
/out:nmqwrld.exe nmqwrld.cs
```

Managed C++

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\mcp\bldmcp Samp.bat`

The `bldmcp Samp.bat` file contains a line for each sample, which is all that is necessary to build this sample program:

```
cl /clr:oldsyntax MQ_INSTALLATION_PATH\bin mmqwrld.cpp
```

If you want to compile these applications on Microsoft Visual Studio 2003/.NET SDKv1.1, replace the compile command:

```
cl /clr:oldsyntax MQ_INSTALLATION_PATH\bin mmqwrld.cpp
```

with

```
cl /clr MQ_INSTALLATION_PATH\bin mmqwrld.cpp
```

Visual Basic

`MQ_INSTALLATION_PATH\Tools\dotnet\samples\vb\bldvbsamp.bat`

The `bldvbsamp.bat` file contains a line for each sample, which is all that is necessary to build this sample program:

```
vbc /r:System.dll /r:MQ_INSTALLATION_PATH\bin\amqmdnet.dll /out:vmqwrld.exe vmqwrld.vb
```

Solving WebSphere MQ .NET problems

If a program does not complete successfully, run one of the sample applications, and follow the advice given in the diagnostic messages.

These sample applications are described in [“Using WebSphere MQ classes for .NET” on page 546](#).

If the problems continue and you need to contact the IBM service team, you might be asked to turn on the trace facility.

Tracing the sample application

For instructions on using the trace facility, refer to [“Tracing WebSphere MQ .NET programs” on page 568](#).

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 WebSphere 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 WebSphere MQ is installed.

Information about this installation is written to the specified log file (`logfile.txt` in this example).

Writing and deploying WebSphere MQ .NET programs

To use WebSphere MQ classes for .NET to access WebSphere MQ queues, you write programs in any language supported by .NET containing calls that put messages onto, and get messages from, WebSphere MQ queues.

The WebSphere MQ documentation contains information only on the C#, C++ and Visual Basic languages.

This collection of topics provides information to assist with writing applications to interact with WebSphere MQ systems. For details of individual classes, see [The WebSphere MQ .NET classes and interfaces](#).

Connection differences

The way you program for WebSphere MQ .NET has some dependencies on the connection modes you want to use.

Managed client connections

When WebSphere MQ classes for .NET are used as a managed client, there are a number of differences from a standard WebSphere MQ MQI client.

The following features are not available to a managed client:

- Channel compression
- SSL support
- Channel exit chaining

If you try to use these features with a managed client, it will return an `MQException`. If the error is detected at the client end of a connection, it will use reason code `MQRC_ENVIRONMENT_ERROR`. If it is detected at the server end, the reason code returned by the server will be used.

Channel exits written for an unmanaged client do not work. You must write new exits specifically for the managed client. Check that there are no invalid channel exits specified in your client channel definition table (CCDT).

The name of a managed channel exit can be up to 999 characters long. However, if you use the CCDT to specify the channel exit name, it is limited to 128 characters.

Communication is supported only over TCP/IP.

When you stop a queue manager using the **endmqm** command, a server-connection channel to a .NET managed client can take longer to close than server-connection channels to other clients.

If you have set `NMQ_MQ_LIB` to managed in order to use managed WebSphere MQ problem diagnostics, none of the parameters `-i`, `-p`, `-s`, `-b`, or `-c` of the **strmqtrc** command is supported.

A managed .NET application using XA transactions will not work with a z/OS queue manager. A managed .Net client attempting to connect to a z/OS queue manager fails with an error, MQRC_UOW_ENLISTMENT_ERROR (mqrc=2354), on MQOPEN call. However, a managed .NET application using XA transactions will work with distributed queue manager.

Defining which connection type to use

The connection type is determined by the setting of the connection name, channel name, the customization value NMQ_MQ_LIB and the property MQC.TRANSPORT_PROPERTY.

You can specify the connection name as follows:

- Explicitly on an MQQueueManager constructor:

```
public MQQueueManager(String queueManagerName, MQLONG Options, string Channel,
string ConnName)
```

```
public MQQueueManager(String queueManagerName, string Channel, string ConnName)
```

- By setting the properties MQC.HOST_NAME_PROPERTY and, optionally, MQC.PORT_PROPERTY in a hashtable entry on an MQQueueManager constructor:

```
public MQQueueManager(String queueManagerName, Hashtable properties)
```

- As explicit MQEnvironment values

```
MQEnvironment.Hostname
```

```
MQEnvironment.Port(optional).
```

- By setting the properties MQC.HOST_NAME_PROPERTY and, optionally, MQC.PORT_PROPERTY in the MQEnvironment.properties hashtable.

You can specify the channel name as follows:

- Explicitly on an MQQueueManager constructor:

```
public MQQueueManager(String queueManagerName, MQLONG Options, string Channel,
string ConnName)
```

```
public MQQueueManager(String queueManagerName, string Channel, string ConnName)
```

- By setting the property MQC.CHANNEL_PROPERTY in a hashtable entry on an MQQueueManager constructor:

```
public MQQueueManager(String queueManagerName, Hashtable properties)
```

- As an explicit MQEnvironment value

```
MQEnvironment.Channel
```

- By setting the property MQC.CHANNEL_PROPERTY in the MQEnvironment.properties hashtable.

You can specify the transport property as follows:

- By setting the property MQC.TRANSPORT_PROPERTY in a hashtable entry on an MQQueueManager constructor:

```
public MQQueueManager(String queueManagerName, Hashtable properties)
```

- By setting the property MQC.TRANSPORT_PROPERTY in the MQEnvironment.properties hashtable.

Select the connection type you require by using one of the following values:

- MQC.TRANSPORT_MQSERIES_BINDINGS - connect as server
- MQC.TRANSPORT_MQSERIES_CLIENT - connect as non-XA client
- MQC.TRANSPORT_MQSERIES_XACLIENT - connect as XA client
- MQC.TRANSPORT_MQSERIES_MANAGED - connect as non-XA managed client

You can set the customization value NMQ_MQ_LIB to explicitly choose the connection type as shown in the following table

NMQ_MQ_LIB value	Connection type
mqic.dll	Connect as a non-XA client
mqicxa.dll	Connect as an XA client
mqm.dll	Connect as a server or as a non-XA client
managed	Connect as a non-XA managed client
Note: Values of mqic32.dll and mqic32xa.dll are accepted as synonyms of mqic.dll and mqicxa.dll for compatibility with earlier releases. However, mqm.dll and mqm.pdb are only part of the client package from version 7.1 onwards.	

If you choose a connection type which is unavailable in your environment, for example you specify mqic32xa.dll and do not have XA support, WebSphere MQ .NET throws an exception.

Setting NMQ_MQ_LIB to "managed" causes the client to use managed WebSphere MQ problem diagnostic tests, .NET data conversion, and other managed low-level WebSphere MQ functions.

All other values for NMQ_MQ_LIB cause the .NET process to use unmanaged WebSphere MQ problem diagnostic tests and data conversion, and other unmanaged low-level WebSphere MQ functions (assuming a WebSphere MQ MQI client or server is installed on the system).

WebSphere MQ .NET chooses the connection type as follows:

1. If MQC.TRANSPORT_PROPERTY is specified, it connects according to the value of MQC.TRANSPORT_PROPERTY.

Note, however, that setting MQC.TRANSPORT_PROPERTY to MQC.TRANSPORT_MQSERIES_MANAGED does not guarantee that the client process runs managed. Even with this setting, the client is not managed in the following cases:

- If another thread in the process has connected with MQC.TRANSPORT_PROPERTY set to something other than MQC.TRANSPORT_MQSERIES_MANAGED.
 - If NMQ_MQ_LIB is not set to "managed", problem diagnostic tests, data conversion, and other low-level functions are not fully managed (assuming a WebSphere MQ MQI client or server is installed on the system).
2. If a connection name has been specified without a channel name, or a channel name has been specified without a connection name, it throws an error.
 3. If both a connection name and a channel name have been specified:
 - If NMQ_MQ_LIB is set to mqic32xa.dll, it connects as an XA client.
 - If NMQ_MQ_LIB is set to managed, it connects as a managed client.
 - Otherwise it connects as a non-XA client.
 4. If NMQ_MQ_LIB is specified, it connects according to the value of NMQ_MQ_LIB.
 5. If a WebSphere MQ server is installed, it connects as a server.
 6. If a WebSphere MQ MQI client is installed, it connects as a non-XA client.
 7. Otherwise, it connects as a managed client.

Configuration files for WebSphere MQ classes for .NET

A .NET client application can use a WebSphere MQ MQI client configuration file and, if you are using the managed connection type, a .NET application configuration file. Settings in the application configuration file have priority.

Client configuration file

A WebSphere MQ classes for .NET client application can use a client configuration file in the same way as any other WebSphere MQ MQI client. This file is typically called mqclient.ini, but you can specify a different file name. For more information about the client configuration file, see [Configuring a client using a configuration file](#) WebSphere MQ MQI client configuration file.

Only the following attributes in a WebSphere MQ MQI client configuration file are relevant to WebSphere MQ classes for .NET. If you specify other attributes, it has no effect.

Stanza	Attribute
CHANNELS	CCSID
CHANNELS	ChannelDefinitionDirectory
CHANNELS	ChannelDefinitionFile
CHANNELS	ServerConnectionParms
ClientExitPath	ExitsDefaultPath
ClientExitPath	ExitsDefaultPath64
MessageBuffer	MaximumSize
MessageBuffer	PurgeTime
MessageBuffer	UpdatePercentage
TCP	CIntRcvBufSize
TCP	CIntSndBufSize
TCP	IPAddressVersion
TCP	KeepAlive

You can override any of these attributes using the appropriate environment variable.

Application configuration file

If you are running with the managed connection type, you can also override the WebSphere MQ client configuration file and the equivalent environment variables using the .NET application configuration file.

The .NET application configuration file settings are only acted upon when running with the managed connection type, and are ignored for other connection types.

The .NET application configuration file and its format are defined by Microsoft for general use within the .NET framework, but the particular section names, keys and values mentioned in this documentation are specific to Websphere MQ.

The format of the .NET application configuration file is a number of *sections*. Each section contains one or more *keys*, and each key has an associated *value*. The following example shows the sections, keys, and values used in a .NET application configuration file to control the TCP/IP KeepAlive property:

```
<configuration>
  <configSections>
    <section name="TCP" type="System.Configuration.NameValueSectionHandler"/>
  </configSections>
  <TCP>
```



```

    <add key="KeepAlive" value="true"></add>
  </TCP>
</configuration>

```

The keywords used in the .NET application configuration file section names and keys exactly match the keywords for the Stanzas and Attributes defined in the client configuration file.

See your Microsoft documentation for further information.

Example code fragment

The following C# code fragment demonstrates an application that performs three actions:

1. Connect to a queue manager
2. Put a message onto SYSTEM.DEFAULT.LOCAL.QUEUE
3. Get the message back

It also shows how to change the connection type.

```

// =====
// Licensed Materials - Property of IBM
// 5724-H72
// (c) Copyright IBM Corp. 2003, 2024
// =====
using System;
using System.Collections;

using IBM.WMQ;

class MQSample
{
    // The type of connection to use, this can be:-
    // MQC.TRANSPORT_MQSERIES_BINDINGS for a server connection.
    // MQC.TRANSPORT_MQSERIES_CLIENT for a non-XA client connection
    // MQC.TRANSPORT_MQSERIES_XACLIENT for an XA client connection
    // MQC.TRANSPORT_MQSERIES_MANAGED for a managed client connection
    const String connectionType = MQC.TRANSPORT_MQSERIES_CLIENT;

    // Define the name of the queue manager to use (applies to all connections)
    const String qManager = "your_Q_manager";

    // Define the name of your host connection (applies to client connections only)
    const String hostName = "your_hostname";

    // Define the name of the channel to use (applies to client connections only)
    const String channel = "your_channelname";

    /// <summary>
    /// Initialise the connection properties for the connection type requested
    /// </summary>
    /// <param name="connectionType">One of the MQC.TRANSPORT_MQSERIES_ values</param>
    static Hashtable init(String connectionType)
    {
        Hashtable connectionProperties = new Hashtable();

        // Add the connection type
        connectionProperties.Add(MQC.TRANSPORT_PROPERTY, connectionType);

        // Set up the rest of the connection properties, based on the
        // connection type requested
        switch(connectionType)
        {
            case MQC.TRANSPORT_MQSERIES_BINDINGS:
                break;
            case MQC.TRANSPORT_MQSERIES_CLIENT:
            case MQC.TRANSPORT_MQSERIES_XACLIENT:
            case MQC.TRANSPORT_MQSERIES_MANAGED:
                connectionProperties.Add(MQC.HOST_NAME_PROPERTY, hostName);
                connectionProperties.Add(MQC.CHANNEL_PROPERTY, channel);
                break;
        }

        return connectionProperties;
    }
}

```

```

}
/// <summary>
/// The main entry point for the application.
/// </summary>
[STAThread]
static int Main(string[] args)
{
    try
    {
        Hashtable connectionProperties = init(connectionType);

        // Create a connection to the queue manager using the connection
        // properties just defined
        MQQueueManager qMgr = new MQQueueManager(qManager, connectionProperties);

        // Set up the options on the queue we want to open
        int openOptions = MQC.MQOO_INPUT_AS_Q_DEF | MQC.MQOO_OUTPUT;

        // Now specify the queue that we want to open, and the open options
        MQQueue system_default_local_queue =
            qMgr.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE", openOptions);

        // Define a WebSphere MQ message, writing some text in UTF format
        MQMessage hello_world = new MQMessage();
        hello_world.WriteUTF("Hello World!");

        // Specify the message options
        MQPutMessageOptions pmo = new MQPutMessageOptions(); // accept the defaults,
                                                                // same as MQPMO_DEFAULT

        // Put the message on the queue
        system_default_local_queue.Put(hello_world, pmo);

        // Get the message back again

        // First define a WebSphere MQ message buffer to receive the message
        MQMessage retrievedMessage = new MQMessage();
        retrievedMessage.MessageId = hello_world.MessageId;

        // Set the get message options
        MQGetMessageOptions gmo = new MQGetMessageOptions(); //accept the defaults
                                                                //same as MQGMO_DEFAULT

        // Get the message off the queue
        system_default_local_queue.Get(retrievedMessage, gmo);

        // Prove we have the message by displaying the UTF message text
        String msgText = retrievedMessage.ReadUTF();
        Console.WriteLine("The message is: {0}", msgText);

        // Close the queue
        system_default_local_queue.Close();

        // Disconnect from the queue manager
        qMgr.Disconnect();
    }

    //If an error has occurred, try to identify what went wrong.

    //Was it a WebSphere MQ error?
    catch (MQException ex)
    {
        Console.WriteLine("A WebSphere MQ error occurred: {0}", ex.ToString());
    }

    catch (System.Exception ex)
    {
        Console.WriteLine("A System error occurred: {0}", ex.ToString());
    }

    return 0;
} //end of start
} //end of sample

```

Operations on queue managers

This section describes how to connect to, and disconnect from, a queue manager using WebSphere MQ classes for .NET.

Setting up the WebSphere MQ environment

Before you use the client connection to connect to a queue manager, you must set up the WebSphere MQ environment.

Note: This step is not necessary when using WebSphere MQ classes for .NET in server bindings mode.

The .NET programming interface allows you to use the `NMQ_MQ_LIB` customization value but also includes a class `MQEnvironment`. This class allows you to specify details that are to be used during the connection attempt, such as those in the following list:

- Channel name
- Host name
- Port number
- Channel exits
- SSL parameters
- User ID and password

For full information about the `MQEnvironment` class, see [MQEnvironment .NET class](#)

To specify the channel name and host name, use the following code:

```
MQEnvironment.Hostname = "host.domain.com";  
MQEnvironment.Channel = "client.channel";
```

By default, the clients attempt to connect to a WebSphere MQ listener at port 1414. To specify a different port, use the code:

```
MQEnvironment.Port = nnnn;
```

Connecting to a queue manager

You are now ready to connect to a queue manager by creating a new instance of the `MQQueueManager` class:

```
MQQueueManager queueManager = new MQQueueManager("QMGrName");
```

To disconnect from a queue manager, call the `Disconnect` method on the queue manager:

```
queueManager.Disconnect();
```

You must have inquire (`inq`) authority on the queue manager when attempting to connect to the queue manager. Without inquire authority, the connection attempt fails.

If you call the `Disconnect` method, all open queues and processes that you have accessed through that queue manager are closed. However, it is good programming practice to close these resources explicitly when you finish using them. To close the resources, use the `Close` method on the object associated with each resource.

The `Commit` and `Backout` methods on a queue manager replace the `MQCMIT` and `MQBACK` calls that are used with the procedural interface.

Accessing queues and topics

You can access queues and topics using methods of `MQQueueManager` or appropriate constructors.

To access queues, use the methods of the `MQQueueManager` class. The MQOD (object descriptor structure) is collapsed into the parameters of these methods. For example, to open a queue on a queue manager represented by an `MQQueueManager` object called `queueManager`, use the following code:

```
MQQueue queue = queueManager.AccessQueue("qName",
                                          MQC.MQOO_OUTPUT,
                                          "qMgrName",
                                          "dynamicQName",
                                          "altUserId");
```

The *options* parameter is the same as the *Options* parameter in the `MQOPEN` call.

The `AccessQueue` method returns a new object of class `MQQueue`.

When you have finished using the queue, use the `Close()` method to close it, as in the following example:

```
queue.Close();
```

With WebSphere MQ .NET, you can also create a queue by using the `MQQueue` constructor. The parameters are exactly the same as for the `accessQueue` method, with the addition of a queue manager parameter specifying the instantiated `MQQueueManager` object to use. For example:

```
MQQueue queue = new MQQueue(queueManager,
                             "qName",
                             MQC.MQOO_OUTPUT,
                             "qMgrName",
                             "dynamicQName",
                             "altUserId");
```

Constructing a queue object in this way enables you to write your own subclasses of `MQQueue`.

Similarly, you can also access topics using the methods of the `MQQueueManager` class. Use an `AccessTopic()` method to open a topic. This returns a new object of class `MQTopic`. When you have finished using the topic, use the `Close()` method of the `MQTopic` to close it.

You can also create a topic by using an `MQTopic` constructor. There are a number of constructors for topics; for more information see [MQTopic .NET class](#).

Handling messages

Messages are handled using the methods of the queue or topic classes. To build a new message, create a new `MQMessage` object.

Put messages onto queues or topics using the `Put()` method of the `MQQueue` or `MQTopic` class. Get messages from queues or topics using the `Get()` method of the `MQQueue` or `MQTopic` class. Unlike the procedural interface, where `MQPUT` and `MQGET` put and get arrays of bytes, the WebSphere MQ classes for .NET put and get instances of the `MQMessage` class. The `MQMessage` class encapsulates the data buffer that contains the actual message data, together with all the MQMD (message descriptor) parameters that describe that message.

To build a new message, create a new instance of the `MQMessage` class and use the `WriteXXX` methods to put data into the message buffer.

When the new message instance is created, all the MQMD parameters are automatically set to their default values, as defined in [Initial values and language declarations for MQMD](#). The `Put()` method of `MQQueue` also takes an instance of the `MQPutMessageOptions` class as a parameter. This class represents the MQPMO structure. The following example creates a message and puts it onto a queue:

```
// Build a new message containing my age followed by my name
MQMessage myMessage = new MQMessage();
myMessage.WriteInt(25);
```

```
String name = "Charlie Jordan";
myMessage.WriteUTF(name);

// Use the default put message options...
MQPutMessageOptions pmo = new MQPutMessageOptions();

// put the message!
queue.Put(myMessage, pmo);
```

The `Get()` method of `MQQueue` returns a new instance of `MQMessage`, which represents the message just taken from the queue. It also takes an instance of the `MQGetMessageOptions` class as a parameter. This class represents the `MQGMO` structure.

You do not need to specify a maximum message size, because the `Get()` method automatically adjusts the size of its internal buffer to fit the incoming message. Use the `ReadXXX` methods of the `MQMessage` class to access the data in the returned message.

The following example shows how to get a message from a queue:

```
// Get a message from the queue
MQMessage theMessage = new MQMessage();
MQGetMessageOptions gmo = new MQGetMessageOptions();
queue.Get(theMessage, gmo); // has default values

// Extract the message data
int age = theMessage.ReadInt();
String name1 = theMessage.ReadUTF();
```

You can alter the number format that the read and write methods use by setting the *encoding* member variable.

You can alter the character set to use for reading and writing strings by setting the *characterSet* member variable.

See [MQMessage .NET class](#) for more details.

Note: The `WriteUTF()` method of `MQMessage` automatically encodes the length of the string as well as the Unicode bytes it contains. When your message will be read by another .NET program (using `ReadUTF()`), this is the simplest way to send string information.

Handling message properties

Message properties allow you to select messages, or to retrieve information about a message without accessing its headers. The `MQMessage` class contains methods to get and set properties.

You can use message properties to allow an application to select messages to process, or to retrieve information about a message without accessing `MQMD` or `MQRFH2` headers. They also facilitate communication between WebSphere MQ and JMS applications. For more information about message properties in WebSphere MQ, see [Message properties](#).

The `MQMessage` class provides a number of methods to get and set properties, according to the data type of the property. The get methods have names of the format `Get*Property`, and the set methods have names of the format `Set*Property`, where the asterisk (*) represents one of the following strings:

- Boolean
- Byte
- Bytes
- Double
- Float
- Int
- Int2
- Int4
- Int8

- Long
- Object
- Short
- String

For example, to get the WebSphere MQ property `myproperty` (a character string), use the call `message.GetStringProperty('myproperty')`. You can optionally pass a property descriptor, which WebSphere MQ will complete.

Handling errors

Handle errors arising from WebSphere MQ classes for .NET using `try` and `catch` blocks.

Methods in the .NET interface do not return a completion code and reason code. Instead, they throw an exception whenever the completion code and reason code resulting from a WebSphere MQ call are not both zero. This simplifies the program logic so that you do not have to check the return codes after each call to WebSphere MQ. You can decide at which points in your program you want to deal with the possibility of failure. At these points, you can surround your code with `try` and `catch` blocks, as in the following example:

```
try
{
    myQueue.Put(messageA, PutMessageOptionsA);
    myQueue.Put(messageB, PutMessageOptionsB);
}
catch (MQException ex)
{
    // This block of code is only executed if one of
    // the two put methods gave rise to a non-zero
    // completion code or reason code.
    Console.WriteLine("An error occurred during the put operation:" +
        "CC = " + ex.CompletionCode +
        "RC = " + ex.ReasonCode);
    Console.WriteLine("Cause exception:" + ex );
}
```

Getting and setting attribute values

The classes `MQManagedObject`, `MQQueue`, and `MQQueueManager` contain methods that allow you to get and set their attribute values. Note that for `MQQueue`, the methods work only if you specify the appropriate inquire and set flags when you open the queue.

For common attributes, the `MQQueueManager` and `MQQueue` classes inherit from a class called `MQManagedObject`. This class defines the `Inquire()` and `Set()` interfaces.

When you create a new queue manager object by using the `new` operator, it is automatically opened for inquire. When you use the `AccessQueue()` method to access a queue object, that object is *not* automatically opened for either inquire or set operations, this could cause problems with some types of remote queues. To use the `Inquire` and `Set` methods and to set properties on a queue, you must specify the appropriate inquire and set flags in the `openOptions` parameter of the `AccessQueue()` method.

The inquire and set methods take three parameters:

- selectors array
- `intAttrs` array
- `charAttrs` array

You do not need the `SelectorCount`, `IntAttrCount`, and `CharAttrLength` parameters that are found in `MQINQ`, because the length of an array is always known. The following example shows how to make an inquiry on a queue:

```
//inquire on a queue
int [ ] selectors = new int [2] ;
```

```
int [ ] intAttrs = new int [1] ;
byte [ ] charAttrs = new byte [MQC.MQ_Q_DESC_LENGTH];
selectors [0] = MQC.MQIA_DEF_PRIORITY;
selectors [1] = MQC.MQCA_Q_DESC;
queue.Inquire(selectors,intAttrs,charAttrs);
ASCIIEncoding enc = new ASCIIEncoding();
String s1 = "";
s1 = enc.GetString(charAttrs);
```

All attributes of these objects can be inquired on. A subset of attributes is exposed as the properties of an object. For a list of object attributes, see [Attributes of objects](#). For object properties, see the appropriate class description.

Multithreaded programs

The .NET runtime environment is inherently multithreaded. WebSphere MQ classes for .NET allows a queue manager object to be shared across multiple threads but ensures that all access to the target queue manager is synchronized.

Consider a simple program that connects to a queue manager and opens a queue at startup. The program displays a single button on the screen. When a user clicks that button, the program fetches a message from the queue. In this situation, the application initialization occurs in one thread, and the code that executes in response to the button press executes in a separate thread (the user interface thread).

The implementation of WebSphere MQ .NET ensures that, for a particular connection (MQQueueManager object instance), all access to the target WebSphere MQ queue manager is synchronized. The default behaviour is that a thread that wants to issue a call to a queue manager is blocked until all other calls in progress for that connection are complete. If you require simultaneous access to the same queue manager from multiple threads within your program, create a new MQQueueManager object for each thread that requires concurrent access. (This is equivalent to issuing a separate MQCONN call for each thread.)

If the default connection options are overridden by MQC.MQCNO_HANDLE_SHARE_NONE or MQC.MQCNO_SHARE_NO_BLOCK then the queue manager is no longer synchronized.

Using a client channel definition table with .NET

You can use a client channel definition table (CCDT) with the .NET classes for WebSphere MQ. You specify the location of the CCDT in different ways, depending on whether you are using a managed or unmanaged connection.

Non-XA or XA unmanaged client connection type

With an unmanaged connection type, you can specify the location of the CCDT in two ways:

- Using the environment variables MQCHLLIB to specify the directory where the table is located, and MQCHLTAB to specify the file name of the table.
- Using the client configuration file. In the CHANNELS stanza, use the attributes ChannelDefinitionDirectory to specify the directory where the table is located, and ChannelDefinitionFile to specify the file name.

If the location is specified both in the client configuration file and by using environment variables, the environment variables take priority. You can use this feature to specify a standard location in the client configuration file and override it using environment variables when necessary.

Managed client connection type

With a managed connection type, you can specify the location of the CCDT in three ways:

- Using the .NET application configuration file. In the CHANNELS section, use the keys ChannelDefinitionDirectory to specify the directory where the table is located, and ChannelDefinitionFile to specify the file name.

- Using the environment variables MQCHLLIB to specify the directory where the table is located, and MQCHLTAB to specify the file name of the table.
- Using the client configuration file. In the CHANNELS stanza, use the attributes ChannelDefinitionDirectory to specify the directory where the table is located, and ChannelDefinitionFile to specify the file name.

If the location is specified in more than one of these ways, the environment variables take priority over the client configuration file, and the .NET Application Configuration File takes priority over both other methods. You can use this feature to specify a standard location in the client configuration file and override it using environment variables or the application configuration file when necessary.

How a .NET application determines what channel definition to use

In the WebSphere MQ .NET client environment, the channel definition to be used can be specified in a number of different ways. Multiple specifications of the channel definition can exist. An application derives the channel definition from one or more sources.

If more than one channel definition exists, the one used is selected in the following priority order:

1. Properties specified on the MQQueueManager constructor, either explicitly or by including *MQC.CHANNEL_PROPERTY* in the properties hashtable
2. A property *MQC.CHANNEL_PROPERTY* in the MQEnvironment.properties hashtable
3. The property *Channel* in MQEnvironment
4. The .NET application configuration file, section name CHANNELS, key ServerConnectionParms (applies to managed connections only)
5. The *MQSERVER* environment variable
6. The client configuration file, stanza CHANNELS, Attribute ServerConnectionParms
7. The client channel definition table (CCDT). The location of the CCDT is specified in the .NET application configuration file (applies to managed connections only)
8. The client channel definition table (CCDT). The location of the CCDT is specified using the environment variables *MQCHLLIB* and *MQCHLTAB*
9. The client channel definition table (CCDT). The location of the CCDT is specified using the client configuration file

For items 1-3, the channel definition is built up field by field from values provided by the application. These values can be provided using different interfaces and multiple values can exist for each one. Field values are added to the channel definition following the priority order given:

1. The value of *connName* on the MQQueueManager constructor
2. Values of properties from the MQQueueManager.properties hashtable
3. Values of properties from the MQEnvironment.properties hashtable
4. Values set as MQEnvironment fields (for example, MQEnvironment.Hostname, MQEnvironment.Port)

For items 4-6, the entire channel definition is supplied as the value. Unspecified fields on the channel definition take the system defaults. No values from other methods of defining channels and their fields are merged with these specifications.

For items 7-9, the entire channel definition is taken from the CCDT. Fields which were not specified explicitly when the channel was defined take the system defaults. No values from other methods of defining channels and their fields are merged with these specifications.

Using channel exits in IBM WebSphere MQ .NET

If you use client bindings, you can use channel exits as for any other client connection. If you use managed bindings, you must write an exit program that implements an appropriate interface.

Client bindings

If you use client bindings, you can use channel exits as described in [Channel exits](#). You cannot use channel exits written for managed bindings.

Managed bindings

If you use a managed connection, to implement an exit, you define a new .NET class that implements the appropriate interface. Three exit interfaces are defined in the WebSphere MQ package:

- MQSendExit
- MQReceiveExit
- MQSecurityExit

Note: User exits written using these interfaces are not supported as channel exits in the unmanaged environment.

The following sample defines a class that implements all three:

```
class MyMQExits : MQSendExit, MQReceiveExit, MQSecurityExit
{
    // This method comes from the send exit
    byte[] SendExit(MQChannelExit channelExitParms,
                    MQChannelDefinition channelDefinition,
                    byte[] dataBuffer,
                    ref int dataOffset,
                    ref int dataLength,
                    ref int dataMaxLength)
    {
        // complete the body of the send exit here
    }

    // This method comes from the receive exit
    byte[] ReceiveExit(MQChannelExit channelExitParms,
                      MQChannelDefinition channelDefinition,
                      byte[] dataBuffer,
                      ref int dataOffset,
                      ref int dataLength,
                      ref int dataMaxLength)
    {
        // complete the body of the receive exit here
    }

    // This method comes from the security exit
    byte[] SecurityExit(MQChannelExit channelExitParms,
                       MQChannelDefinition channelDefParms,
                       byte[] dataBuffer,
                       ref int dataOffset,
                       ref int dataLength,
                       ref int dataMaxLength)
    {
        // complete the body of the security exit here
    }
}
```

Each exit is passed an MQChannelExit and an MQChannelDefinition object instance. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

The data to be sent by a send exit, and the data received in a security or receive exit is specified using the exit's parameters.

On entry, the data at offset *dataOffset* with length *dataLength* in the byte array *dataBuffer* is the data that is about to be sent by a send exit, and the data received in a security or receive exit. The parameter

dataMaxLength gives the maximum length (from *dataOffset*) available to the exit in *dataBuffer*. Note: For a security exit, it is possible for the *dataBuffer* to be null, if this is the first time the exit is called or the partner end elected to send no data.

On return, the value of *dataOffset* and *dataLength* should be set to point to the offset and length within the returned byte array that the .NET classes should then use. For a send exit, this indicates the data that it should send, and for a security or receive exit, the data that should be interpreted. The exit should normally return a byte array; exceptions are a security exit which could elect to send no data, and any exit called with the INIT or TERM reasons. The simplest form of exit that can be written therefore is one which does nothing more than return *dataBuffer*:

The simplest possible exit body is:

```
{  
    return dataBuffer;  
}
```

MQChannelDefinition class

V7.5.0.6 From Version 7.5.0, Fix Pack 6, the userid and password that are specified with the managed .NET client application are set in the IBM WebSphere MQ .NET MQChannelDefinition class that is passed to the client security exit. The security exit copies the userid and password into the MQCD.RemoteUserIdentifier and MQCD.RemotePassword fields (see [“Writing a security exit” on page 388](#)).

Specifying channel exits (managed client)

If you specify a channel name and connection name when creating your MQQueueManager object (either in the MQEnvironment or on the MQQueueManager constructor) you can specify channel exits in two ways.

In order of precedence, these are:

1. Passing hashtable properties MQC.SECURITY_EXIT_PROPERTY, MQC.SEND_EXIT_PROPERTY or MQC.RECEIVE_EXIT_PROPERTY on the MQQueueManager constructor.
2. Setting the MQEnvironment SecurityExit, SendExit or ReceiveExit properties.

If you do not specify a channel name and connection name, the channel exits to use come from the channel definition picked up from a client channel definition table (CCDT). It is not possible to override the values stored in the channel definition. See [Client channel definition table](#) and [“Using a client channel definition table with .NET” on page 559](#) for more information about channel definition tables.

In each case, the specification takes the form of a string with the following format:

```
Assembly_name(Class_name)
```

Class_name is the fully qualified name, including namespace specification, of a .NET class that implements the IBM.WMQ.MQSecurityExit, IBM.WMQ.MQSendExit or IBM.WMQ.MQReceiveExit interface (as appropriate). *Assembly_name* is the fully qualified location, including file extension, of the assembly that houses the class. The length of the string is limited to 999 characters if you use the properties of MQEnvironment or MQQueueManager. However, if the channel exit name is specified in the CCDT, it is limited to 128 characters. When necessary, the .NET client code loads and creates an instance of the specified class by parsing the string specification.

Specifying channel exit user data (managed client)

Channel exits can have user data associated with them. If you specify a channel name and connection name when creating your MQQueueManager object (either in the MQEnvironment or on the MQQueueManager constructor) you can specify the user data in two ways.

In order of precedence, these are:

1. Passing hashtable properties MQC.SECURITY_USERDATA_PROPERTY, MQC.SEND_USERDATA_PROPERTY or MQC.RECEIVE_USERDATA_PROPERTY on the MQQueueManager constructor.

2. Setting the MQEnvironment SecurityUserData, SendUserData or ReceiveUserData properties.

If you do not specify a channel name and connection name, the exit user data values to use come from the channel definition picked up from the client channel definition table (CCDT). It is not possible to override the values stored in the channel definition. See [Client channel definition table](#) and [“Using a client channel definition table with .NET”](#) on page 559 for more information about channel definition tables.

In each case, the specification is a string, limited to 32 characters.

Automatic client reconnection in .NET

You can make your client reconnect automatically to a queue manager during an unexpected connection break.

A client can unexpectedly become disconnected from a queue manager if, for example, the queue manager stops, or the network or server fails.

Without automatic client reconnection, an error is produced when the connection fails. You can use the error code to help you reestablish the connection.

A client that uses the automatic client reconnection facility is called a reconnectable client. To create a reconnectable client, specify certain options called reconnect options while connecting to the queue manager.

If the client application is a WebSphere MQ .NET client, it can opt to get automatic client reconnection by specifying an appropriate value for CONNECT_OPTIONS_PROPERTY when you use the MQQueueManager class to create a queue manager. See [Reconnection options](#) for details of CONNECT_OPTIONS_PROPERTY values.

You can select whether the client application always connects and reconnects to a queue manager of the same name, to the same queue manager, or to any set of queue managers that are defined with the same QMNAME in the client connection table (See [Queue Manager Groups in CCDT](#) for details).

Secure Sockets Layer (SSL) support

The following section does not apply to the managed client.

WebSphere MQ classes for .NET client applications support Secure Sockets Layer (SSL) encryption. SSL provides communication encryption, authentication, and message integrity. It is typically used to secure communications between any two peers on the Internet or within an intranet.

Enabling SSL

SSL is supported only for client connections. To enable SSL, you must specify the CipherSpec to use when communicating with the queue manager, and this must match the CipherSpec set on the target channel.

To enable SSL, specify the CipherSpec using the SSLCipherSpec static member variable of MQEnvironment. The following example attaches to a SVRCONN channel named SECURE.SVRCONN.CHANNEL, which has been set up to require SSL with a CipherSpec of NULL_MD5:

```
MQEnvironment.Hostname      = "your_hostname";
MQEnvironment.Channel       = "SECURE.SVRCONN.CHANNEL";
MQEnvironment.SSLCipherSpec = "NULL_MD5";
MQEnvironment.SSLKeyRepository = "C:\mqm\key";
MQQueueManager qmgr = new MQQueueManager("your_q_manager");
```

See [Specifying CipherSpecs](#) for a list of CipherSpecs.

The SSLCipherSpec property can also be set using the MQC.SSL_CIPHER_SPEC_PROPERTY in the hash table of connection properties.

To successfully connect using SSL, the client keystore must be set up with Certificate Authority root certificates chain from which the certificate presented by the queue manager can be authenticated. Similarly, if SSLClientAuth on the SVRCONN channel has been set to MQSSL_CLIENT_AUTH_REQUIRED, the client keystore must contain an identifying personal certificate that is trusted by the queue manager.

Using the Distinguished Name of the queue manager

The queue manager identifies itself using an SSL certificate, which contains a *Distinguished Name* (DN).

A WebSphere MQ .NET client application can use this DN to ensure that it is communicating with the correct queue manager. A DN pattern is specified using the sslPeerName variable of MQEnvironment. For example, setting:

```
MQEnvironment.SSLPeerName = "CN=QMGR.*, OU=IBM, OU=WEBSPPHERE";
```

allows the connection to succeed only if the queue manager presents a certificate with a Common Name beginning QMGR., and at least two Organizational Unit names, the first of which must be IBM and the second WEBSPPHERE.

The SSLPeerName property can also be set using the MQC.SSL_PEER_NAME_PROPERTY in the hash table of connection properties. For more information about Distinguished Names and rules for setting peer names, refer to [Security](#).

If SSLPeerName is set, connections succeed only if it is set to a valid pattern and the queue manager presents a matching certificate.

Error handling when using SSL

The following reason codes can be issued by WebSphere MQ classes for .NET when connecting to a queue manager using SSL:

MQRC_SSL_NOT_ALLOWED

The SSLCipherSpec property was set, but bindings connect was used. Only client connect supports SSL.

MQRC_SSL_PEER_NAME_MISMATCH

The DN pattern specified in the SSLPeerName property did not match the DN presented by the queue manager.

MQRC_SSL_PEER_NAME_ERROR

The DN pattern specified in the SSLPeerName property was not valid.

Using the .NET Monitor

See [Features that can be used only with the primary installation on Windows](#) for important information.

The .NET Monitor is an application similar to a WebSphere MQ trigger monitor. You can create .NET components which are instantiated whenever a message is received on a monitored queue, and which then process that message. The .NET Monitor is started by the runmqdnm command and stopped by the endmqdnm command. For details of these commands, see [runmqdnm](#) and [endmqdnm](#).

To use the .NET Monitor, you write a component that implements the IMQObjectTrigger interface, which is defined in amqmdnm.dll.

Components can be either transactional or non-transactional. A transactional component must inherit from System.EnterpriseServices.ServicedComponent and be registered as either RequiresTransaction or SupportsTransaction. It must not be registered as RequiresNew as the .NET Monitor has already initiated a transaction.

The component receives MQQueueManager, MQQueue, and MQMessage objects from runmqdnm. It may also receive a User Parameter string if one was specified, using the -u command-line option, when runmqdnm was started. Note that your component receives the contents of a message that arrived on the monitored queue in an MQMessage object. It does not have to connect to the queue manager, open

the queue, or get the message itself. The component must then process the message as appropriate and return control to the .NET Monitor.

If your component has been written as a transactional component, it registers to commit or roll back the transaction using the facilities provided by `System.EnterpriseServices.ServicedComponent`.

As the component receives `MQQueueManager` and `MQQueue` objects as well as the message, it has complete context information for that message and can, for example, open another queue on the same queue manager without needing to separately connect to WebSphere MQ.

Example code fragments

This topic contains two examples of components which obtain a message from the .NET Monitor and print it, one using transactional processing and the other non-transactional processing. A third example shows common utility routines, applicable to both the first two examples. All the examples are in C#.

Example 1: Transactional processing

```
/* **** */
/* Licensed materials, property of IBM */
/* 63H9336 */
/* (C) Copyright IBM Corp. 2005, 2024. */
/* **** */
using System;
using System.EnterpriseServices;

using IBM.WMQ;
using IBM.WMQMonitor;

[assembly: ApplicationName("dnmsamp")]

// build:
//
// csc -target:library -reference:amqmdnet.dll;amqmdnm.dll TranAssembly.cs
//
// run (with dotnet monitor)
//
// runmqdmn -m <QMNAME> -q <QNAME> -a dnmsamp.dll -c Tran

namespace dnmsamp
{
    [TransactionAttribute(TransactionOption.Required)]
    public class Tran : ServicedComponent, IMQObjectTrigger
    {
        Util util = null;

        [AutoComplete(true)]
        public void Execute(MQQueueManager qmgr, MQQueue queue,
            MQMessage message, string param)
        {
            util = new Util("Tran");

            if (param != null)
                util.Print("PARAM: '" + param.ToString() + "'");

            util.PrintMessage(message);

            //System.Console.WriteLine("SETTING ABORT");
            //ContextUtil.MyTransactionVote = TransactionVote.Abort;

            System.Console.WriteLine("SETTING COMMIT");
            ContextUtil.SetComplete();
            //ContextUtil.MyTransactionVote = TransactionVote.Commit;
        }
    }
}
```

Example 2: Non-transactional processing

```
/* **** */
/* Licensed materials, property of IBM */
/* 63H9336 */
/* **** */
```

```

/* (C) Copyright IBM Corp. 2005, 2024. */
/*****/

using System;

using IBM.WMQ;
using IBM.WMQMonitor;

// build:
//
// csc -target:library -reference:amqmdnet.dll;amqmdnm.dll NonTranAssembly.cs
//
// run (with dotnet monitor)
//
// runmqdmn -m <QMNAME> -q <QNAME> -a dnmsamp.dll -c NonTran
namespace dnmsamp
{
    public class NonTran : IMQObjectTrigger
    {
        Util util = null;

        public void Execute(MQQueueManager qmgr, MQQueue queue,
            MQMessage message, string param)
        {
            util = new Util("NonTran");

            try
            {
                util.PrintMessage(message);
            }

            catch (Exception ex)
            {
                System.Console.WriteLine(">>> NonTran\n{0}", ex.ToString());
            }
        }
    }
}

```

Example 3: Common routines

```

/*****/
/* Licensed materials, property of IBM */
/* 63H9336 */
/* (C) Copyright IBM Corp. 2005, 2024. */
/*****/

using System;

using IBM.WMQ;

namespace dnmsamp
{
    /// <summary>
    /// Summary description for Util.
    /// </summary>
    public class Util
    {
        /* ----- */
        /* Default prefix string of the namespace. */
        /* ----- */
        private string prefixText = "dnmsamp";

        /* ----- */
        /* Constructor that takes the replacement prefix string to use. */
        /* ----- */
        public Util(String text)
        {
            prefixText = text;
        }

        /* ----- */
        /* Display an arbitrary string to the console. */
        /* ----- */
        public void Print(String text)
        {
            System.Console.WriteLine("{0} {1}\n", prefixText, text);
        }
    }
}

```

```

    }

    /* ----- */
    /* Display the content of the message passed to the console. */
    /* ----- */
    public void PrintMessage(MQMessage message)
    {
        if (message.Format.CompareTo(MQC.MQFMT_STRING) == 0)
        {
            try
            {
                string messageText = message.ReadString(message.MessageLength);

                Print(messageText);
            }

            catch(Exception ex)
            {
                Print(ex.ToString());
            }
        }
        else
        {
            Print("UNRECOGNISED FORMAT");
        }
    }

    /* ----- */
    /* Convert the byte array into a hex string. */
    /* ----- */
    static public string ToHexString(byte[] byteArray)
    {
        string hex = "0123456789ABCDEF";

        string retString = "";

        for(int i = 0; i < byteArray.Length; i++)
        {
            int h = (byteArray[i] & 0xF0)>>4;
            int l = (byteArray[i] & 0x0F);

            retString += hex.Substring(h,1) + hex.Substring(l,1);
        }

        return retString;
    }
}

```

Compiling WebSphere MQ .NET programs

Specimen commands to compile .NET applications written in various languages.

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

To build a C# application using WebSphere MQ classes for .NET, use the following command:

```
csc /t:exe /r:System.dll /r:amqmdnet.dll /lib:MQ_INSTALLATION_PATH\bin /out:MyProg.exe MyProg.cs
```

To build a Visual Basic application using WebSphere MQ classes for .NET, use the following command:

```
vbc /r:System.dll /r:MQ_INSTALLATION_PATH\bin\amqmdnet.dll /out:MyProg.exe MyProg.vb
```

To build a Managed C++ application using WebSphere MQ classes for .NET, use the following command:

```
cl /clr MQ_INSTALLATION_PATH\bin Myprog.cpp
```

For other languages, see the documentation supplied by the language vendor.

Tracing WebSphere MQ .NET programs

In WebSphere MQ .NET, you start and control the trace facility as in WebSphere 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.

IBM WebSphere MQ custom channel for Microsoft Windows Communication Foundation (WCF)

The Microsoft Windows Communication Foundation (WCF) custom channel for IBM WebSphere MQ sends and receives messages between WCF clients and services.

Related concepts

[“Introduction to the use of WebSphere MQ custom channel for WCF with .NET 3” on page 568](#)

Overview of the information available for programmers using the WebSphere MQ custom channel for Windows Communication Foundation (WCF) with .NET 3.

[“Using WebSphere MQ custom channels for WCF” on page 572](#)

Overview of the information available for programmers using WebSphere MQ V7 custom channels for Windows Communication Foundation (WCF).

[“Using the WCF samples” on page 588](#)

The Windows Communication Foundation (WCF) samples provide some simple examples of how the WebSphere MQ custom channel can be used.

[“Problem determination on the WCF custom channel for WebSphere MQ” on page 595](#)

You can use WebSphere MQ trace to collect detailed information about what various parts of the WebSphere MQ code is doing. When using Windows Communication Foundation (WCF), a separate trace output is generated for the WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.

Introduction to the use of WebSphere MQ custom channel for WCF with .NET 3

Overview of the information available for programmers using the WebSphere MQ custom channel for Windows Communication Foundation (WCF) with .NET 3.

What is the WebSphere MQ custom channel for WCF?

The custom channel for WebSphere MQ is a transport channel using the Microsoft Windows Communication Foundation (WCF) unified programming model.

The Microsoft Windows Communication Foundation framework, introduced in Microsoft .NET 3, enables .NET applications and services to be developed independently from the transport and protocols used to connect them, enabling alternative transports or configurations to be used according to the environment that the service or application is deployed in.

Connections are managed at run time by WCF by building a channel stack containing the required combination of:

- Protocol elements: An optional set of elements where none, one, or more can be added to support protocols such as the WS-* standards.
- Message encoder: A mandatory element in the stack controlling the serializing of the message into its wire format.
- Transport channel: A mandatory element in the stack responsible for transporting the serialized message to its endpoint.


The custom channel for WebSphere MQ is a transport channel, and as such must be paired with a message encoder and optional protocols as required by the application using a WCF custom binding. In this way, applications which have been developed to use WCF can use the custom channel for WebSphere MQ to send and receive data in the same way as they use the built-in transports provided by Microsoft, enabling simple integration with the asynchronous, scalable, and reliable messaging functions of WebSphere MQ. For a full list of supported functions, see: [“WCF Custom channel features and capabilities”](#) on page 572.

When and why do I use the WebSphere MQ custom channel for WCF?

The WebSphere MQ custom channel can be used to send and receive messages between WCF clients and services in the same way as the built-in transports provided by Microsoft, enabling applications to access the features of WebSphere MQ within the WCF unified programming model.

A typical usage pattern scenario of the WebSphere MQ custom channel for WCF is as an interface to web services hosted over WebSphere MQ (SOAP/JMS)

Messages are carried using the SOAP over JMS message format of WebSphere MQ, enabling WCF clients and services to also call or be called by other WebSphere MQ applications or hosting environments which are compatible with this format, including web services and clients running in WebSphere Application Server, CICS, Axis v1 (Java), and .asmx (.NET), as shown in the following diagram:



For details on SOAP over JMS, see: [“WebSphere MQ transport for SOAP”](#) on page 905

An example of a typical scenario from the diagram would be:

1. A web Service hosted within WebSphere Application Server and exposed over WebSphere MQ using the support for SOAP over JMS within WebSphere Application Server
2. The WSDL document describing the service can then be used by the WCF tool to generate a client proxy and configuration which would then create an appropriate WCF channel stack, including the custom channel.
3. The client application can then use the proxy to start the web service in the same way as any other web service.

The channel would typically be used with a WCF text/SOAP message encoder, but the channel can be paired with other WCF message encoders if required. Using alternative encoders can also provide limited integration with native WebSphere MQ applications which do not support SOAP over JMS, but this is not the primary role of the channel.

The key benefits of using the custom channel in a WCF environment are:

- Asynchronous invocation: Supporting fire and forget client operations where the client is decoupled from the availability of the service and features, such as rerouting of responses and multi-hop.
- Reliable scaling characteristics: Queue based messaging allows capacity to be predictably added to a system.

- Quality of service: Messages are tangible and traceable, and can be easily managed and administered.

Software requirements and installation instructions for the WebSphere MQ custom channel for WCF

This topic outlines the software requirements and installation information for the WebSphere MQ custom channel for WCF.

The WebSphere MQ custom channel for WCF can only connect to WebSphere MQ V7 or higher queue managers.

Software requirements for the WCF custom channel for WebSphere MQ

This information lists the software requirements for the WCF custom channel for WebSphere MQ.

Runtime environment

- Microsoft .NET Framework v3.0 or higher must be installed on the host machine.
- *Java and .NET Messaging and Web Services* is installed by default as part of the WebSphere MQ 7.0.1 installer. Installs the .NET assemblies needed for the custom channel into the Global Assembly Cache.

Note: If the Microsoft .NET Framework v2.0 or higher is not installed before installing WebSphere MQ V7.0.1, then the WebSphere MQ product installation continues without error, but the WebSphere MQ custom channel is unavailable. If the .NET Framework is installed after installing WebSphere MQ 7.0.1, then the WebSphere MQ custom channel must be activated by running the *WMQInstallDir\bin\amqiRegisterdotNet.cmd* script, where *WMQInstallDir* is the directory where WebSphere MQ 7.0.1 is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of *amqi*.log* files recording the actions taken are created in the %TEMP% directory. It is not necessary to rerun the *amqiRegisterdotNet.cmd* script if .NET is upgraded to v3.0 or higher from an earlier version, for example, from .NET v2.0.

Development environment

- Microsoft Visual Studio 2008 or Windows Software Development Kit for .NET 3.0 or later.
- Microsoft .NET Framework V3.5 or higher must be installed on the host machine in order to build the sample solution files.

Note: If the Microsoft .NET Framework v2.0 or higher is not installed before installing WebSphere MQ V7.0.1, then the WebSphere MQ product installation continues without error, but the WebSphere MQ custom channel is unavailable. If the .NET Framework is installed after installing WebSphere MQ 7.0.1, then the WebSphere MQ custom channel must be activated by running the *WMQInstallDir\bin\amqiRegisterdotNet.cmd* script, where *WMQInstallDir* is the directory where WebSphere MQ 7.0.1 is installed. This script installs the required assemblies in the Global Assembly Cache (GAC). A set of *amqi*.log* files recording the actions taken are created in the %TEMP% directory. It is not necessary to rerun the *amqiRegisterdotNet.cmd* script if .NET is upgraded to v3.0 or higher from an earlier version, for example, from .NET v2.0.

WebSphere MQ custom channel for WCF: What's installed?

The custom channel for WebSphere MQ is a transport channel using the Microsoft Windows Communication Foundation (WCF) unified programming model. The custom channel is installed by default as part of the WebSphere MQ 7.0.1 installation.

WebSphere MQ custom channel for WCF

The WebSphere MQ custom channel for WCF is installed by default as part of the WebSphere MQ 7.0.1 installation; The custom channel and its dependencies are contained within the *Java and .NET Messaging and Web Services* component, which is installed by default. When upgrading to WebSphere MQ 7.0.1 from an earlier version, the update will install the WebSphere MQ custom channel

for WCF by default if the Java and .NET Messaging and Web Services component was previously installed in an earlier installation.

The Java and .NET Messaging and Web Services component contains the IBM.XMS.WCF.dll file, and the IBM.XMS.WCF.dll file is the main custom channel assembly, which contains the WCF interface classes. This file is installed in the Global Assembly Cache (GAC) and is also available in the following directory: *MQ_INSTALLATION_PATH*\bin where *MQ_INSTALLATION_PATH* is the directory that WebSphere MQ 7.0.1 is installed in.

The key classes required for using the custom channel are in the *Namespace: IBM.XMS.WCF* and:

Transport Binding Name	IBM.XMS.WCF.SoapJmsIbmTransportBindingElement
Transport Binding Importer	IBM.XMS.WCF.SoapJmsIbmTransportBindingElementImporter
Transport Binding Config	IBM.XMS.WCF.SoapJmsIbmTransportBindingElementConfig

WebSphere MQ custom channel samples

The samples provide some simple examples of how the WebSphere MQ custom channel for WCF can be used. The samples and their associated files are located in the *MQ_INSTALLATION_PATH*\tools\wcf\samples\ directory, where *MQ_INSTALLATION_PATH* is the installation directory for WebSphere MQ. For more information on the WebSphere MQ custom channel samples, see: [“Using the WCF samples” on page 588](#)


svcutil.exe.config

The svcutil.exe.config is an example of the configuration settings required to enable the Microsoft WCF svcutil client proxy generation tool to recognize the custom channel. The svcutil.exe.config file is located in the *MQ_INSTALLATION_PATH*\tools\wcf\docs\examples\ directory, where *MQ_INSTALLATION_PATH* is the installation directory for WebSphere MQ. For more information on using the svcutil.exe.config, see: [“Generating a WCF client proxy and application configuration files using the svcutil tool with metadata from a running service” on page 586.](#)

WCF architecture

The WebSphere MQ custom channel for WCF is integrated on top of the IBM Message Service Client for .NET (XMS .NET) API.

The WCF architecture is as shown in the following diagram:



All the required components are installed by default with the WebSphere MQ V7.0.1 installation.

The three connections are: Managed client connections, Unmanaged server connections, and Unmanaged client connections. For more information about these connections, see [“WCF Connection options”](#) on page 576.

Using WebSphere MQ custom channels for WCF

Overview of the information available for programmers using WebSphere MQ V7 custom channels for Windows Communication Foundation (WCF).

The Microsoft Windows Communication Foundation underpins the web services and messaging support in the Microsoft .NET Framework 3. WebSphere MQ V7 can now be used as a custom channel within WCF in the .NET Framework 3 in the same manner as the built-in channels offered by Microsoft.

Messages transported across the custom channel are formatted according to the SOAP over JMS implementation of WebSphere MQ V7. Applications can then communicate with services hosted by WCF or by the WebSphere SOAP over JMS service infrastructure. For details on SOAP over JMS, see: [“WebSphere MQ transport for SOAP”](#) on page 905

WCF Custom channel features and capabilities

Use the following topics for information regarding WCF custom channel features and capabilities.

WCF custom channel shapes

Overview of the custom channel shapes that WebSphere MQ can be used as within the Microsoft Windows Communication Foundation (WCF) custom channels.

The WebSphere MQ custom channel for WCF supports two channel shapes:

- One-way
- Request-reply

WCF automatically selects the channel shape according to the service contract being hosted.

Contracts that include methods that only use the **IsOneWay** parameter are serviced by the one-way channel shape, for example:

```
[OperationContract(IsOneWay = true)]  
void printString(String text);
```

Contracts that include either a mixture of one-way and request-reply methods, or all request-reply methods, are serviced by the request-reply channel shape. For example:

```
[OperationContract]  
int subtract(int a, int b);  
  
[OperationContract(IsOneWay = true)]  
void printString(string text);
```

Note: When mixing one-way and request-reply methods in the same contract, you must ensure that the behavior is as intended, especially when working within a mixed environment because one-way methods wait until they receive a null reply from the service.

One-way channel

The WebSphere MQ one-way custom channel for WCF is used, for example, to send messages from a WCF client using a one-way channel shape. The channel can send messages in one direction only, for example; from a client queue manager to a queue on a WCF service.

Request-reply channel

The WebSphere MQ request-reply custom channel for WCF is used, for example, to send messages in two directions asynchronously; The same client instance must be used for asynchronous messaging. The channel can send messages in one direction, for example; from a client queue manager to a queue on a WCF service, and then send a reply message from the WCF to a queue on the client queue manager.

WCF URI parameter names and values

connectionFactory

The `connectionFactory` parameter is required. For the syntax of this parameter, see [URI syntax and parameters for Web service deployment](#).

initialContextFactory

The `initialContextFactory` parameter is required and must be set to "com.ibm.mq.jms.Nojndi" for compatibility with WebSphere Application Server and other products (see ["Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP" on page 960](#)).

WCF custom channel assured delivery

Assured Delivery guarantees that a service request or reply is actioned and not lost.

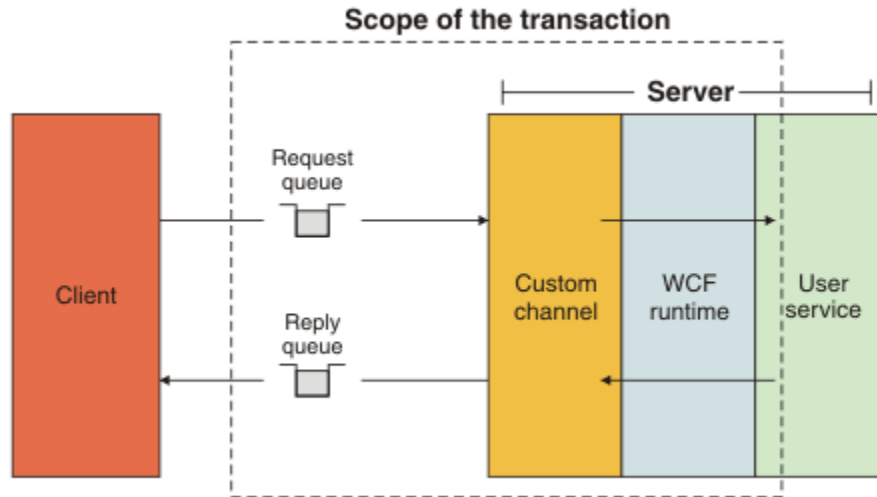
A request message is received and any reply message is sent under a local transaction sync point, which can be rolled back in the case of runtime failure. Examples of these failures are: An unhandled exception thrown by the service, failure to dispatch the message to the service, or failure to deliver the reply message.

`AssuredDelivery` is the assured delivery attribute which can be specified on a service contract to guarantee that any request messages received by a service, and any reply message sent from a service, is not lost in the event of a runtime failure.

To ensure that messages are also preserved in the event of system failure or power outage, messages must be sent as persistent. To use persistent messages the client application must have this option specified on its endpoint URI. For further information about setting URI properties, see: [URI syntax and parameters for Web service deployment](#).

Distributed transactions are not supported, and the scope of the transaction does not extend beyond the request and reply message processing performed by WebSphere MQ. Any work performed within

the service might get rerun as a result of a failure which causes the message to be received again. The following diagram shows the scope of the transaction:



Assured delivery is enabled by applying the `AssuredDelivery` attribute to the service class as shown in the following example:

```
[AssuredDelivery]
class TestCalculatorService : IWMQSampleCalculatorContract
{
    public int add(int a, int b)
    {
        int ans = a + b;
        return ans;
    }
}
```

When using the `AssuredDelivery` attribute, you must be aware of the following points:

- When a channel determines that a failure is likely to recur if a message was rolled-back and received again, the message is treated as a poison message and is not returned to the request queue for reprocessing. For example: If the received message is not correctly formatted or cannot be dispatched to a service. Unhandled exceptions thrown from a service operation are always resent until the message has been redelivered the maximum number of times specified by the backout threshold property of the request queue. For more information, see: [“WCF custom channel poison messages” on page 575](#)
- The channel performs the reading, processing, and replying of each request message as an atomic operation using a single thread of execution to enforce transactional integrity. To enable service operations to run concurrently, the channel enables WCF to create multiple instances of the channel. The number of channel instances available for processing requests is controlled by the binding property `MaxConcurrentCalls`. For more information, see: [“WCF binding configuration options” on page 582](#)
- The assured delivery function uses both the `IOperationInvoker` and the `IErrorHandler` WCF extensibility points. If these extensibility points are used externally by an application, the application must ensure that any previously registered extensibility points are called. Failure to do so for `IErrorHandler` can result in errors going unreported. Failure to do so for `IOperationInvoker` can cause WCF to stop responding.

WCF custom channel security

The WebSphere MQ custom channel for WCF supports the use of SSL only for unmanaged client connections to the queue manager.

SSL can be specified in one of two ways:

- Specify SSL directly on the SOAP over JMS URI. For a full description of the SSL options, see [SSL and the WebSphere MQ transport for SOAP](#)
- Specify SSL using an entry in the client channel definition table (CCDT). For more information about CCDTs, see [Client channel definition table](#)

WCF client channel definition tables (CCDT)

The WebSphere MQ custom channel for WCF supports the use of client channel definition tables (CCDT) to configure the connection information for client connections.

CCDTs are controlled through these two environment variables:

- *MQCHLLIB* specifies the directory where the table is located.
- *MQCHLTAB* specifies the file name of the table.

You cannot specify the channel definition table directly in the SOAP over JMS URI. If these environment variables are defined, then they take priority over any client connection details specified in the URI.

For more information about client channel definition tables, see: [Client channel definition table](#).

Related concepts

[Client channel definition table](#)

WCF custom channel poison messages

When a service fails to process a request message, or fails to deliver a reply message to a reply queue, then the message is treated as a poison message.

Poison request messages

If a request message cannot be processed, then it is treated as a poison message. This action prevents the service from receiving the same unprocessable message again. For an unprocessable request message to be treated as a poison message, one of the following situations must be true:

- The messages backout count exceeded the backout threshold specified on the request queue, which only occurs if assured delivery was specified for the service. For more information about assured delivery, see: [“WCF custom channel assured delivery” on page 573](#)
- The message was not formatted correctly and could not be interpreted as a SOAP over JMS message.

Poison reply messages

If a service fails to deliver a reply message to the reply queue, then the reply message is treated as a poison message. For reply messages, this action enables the reply messages to be retrieved later to aid problem determination.

Poison message handling

The action taken for a poison message depends on the queue manager configuration and the values set in the report options of the message. For SOAP over JMS, the following report options are set on request messages by default and are not configurable:


- *MQRO_EXCEPTION_WITH_FULL_DATA*
- *MQRO_EXPIRATION_WITH_FULL_DATA*
- *MQRO_DISCARD_MSG*

For SOAP over JMS, the following report option is set on reply messages by default and is not configurable:

- *MQRO_DEAD_LETTER_Q*

If messages come from a non-WCF source, then refer to the documentation for that source.

The following diagram shows the possible actions and the steps taken if a poison message handling fails:



WCF Connection options

There are three modes of connecting a WebSphere MQ custom channel for WCF to a queue manager. Consider which type of connection best suits your requirements.

For more information about connection options, see: [“Connection differences” on page 549](#)

For more information about WCF architecture, see: [“WCF architecture” on page 571](#)

Unmanaged client connection

A connection made in this mode connects as a WebSphere MQ client to a WebSphere MQ server running either on the local machine or on a remote machine.

To use the WebSphere MQ custom channel for WCF as a WebSphere MQ client, you can install it, with the WebSphere MQ MQI client, either on the WebSphere MQ server, or on a separate machine.

Unmanaged server connection

When used in server bindings mode, the WebSphere MQ custom channel for WCF uses the queue manager API, rather than communicating through a network. Using bindings connections provides better performance for WebSphere MQ applications than using network connections.

To use the bindings connection, you must install the WebSphere MQ custom channel for WCF on the WebSphere MQ server.

Managed client connection

A connection made in this mode connects as a WebSphere MQ client to a WebSphere MQ server running either on the local machine or on a remote machine.

The WebSphere MQ custom channel classes for .NET 3 connecting in this mode remain in .NET managed code and make no calls to native services. For more information about managed code, see Microsoft documentation.

There are a number of limitations to using the managed client. For more information about these limitations, see [“Managed client connections” on page 549](#).

Creating and configuring the WebSphere MQ custom channel for WCF

The WebSphere MQ V7 custom channels for WCF work in the same manner as transport WCF channels offered by Microsoft. The WebSphere MQ custom channel for WCF can be created in one of two ways.

About this task

The WebSphere MQ custom channel integrates with WCF as a WCF transport channel, and as such must be paired with a message encoder and optional protocol channels, so it can create a complete channel stack that can be used by an application. Two elements are required for a complete channel stack to be created successfully:

1. A binding definition: Specifies which elements are required to build the applications channel stack, including transport channel, message encoder, and any protocols, plus any general configuration settings. For the custom channel, the binding definition must be created in the form of a WCF custom binding.
2. An endpoint definition: Links the service contract with the binding definition, and also provides the actual connection URI which describes where the application can connect. For the custom channel, the URI is in the form of a SOAP over JMS URI.

These definitions can be created in one of two different ways:

- Administratively; The definitions are created by providing the details in an application configuration file (for example: `app.config`).
- Programmatically; The definitions are created directly from the application code.

The decision over which method to use to create the definitions must be based on the requirements of the application as follows:

- The Administrative method for configuration provides the flexibility to alter the details of the service and client post-deployment without rebuilding the application.
- The Programmatic method for configuration provides greater protection from configuration errors, and the ability to dynamically generate a configuration at run time.

Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file

The WebSphere MQ custom channel for WCF is a transport level WCF channel. An endpoint and binding must be defined to use the custom channel, and these definitions can be done by supplying the binding and endpoint information in an application configuration file.

To configure and use the WebSphere MQ custom channel for WCF, which is a transport level WCF channel, a binding and an endpoint definition must be defined. The binding holds the configuration information for the channel, and the endpoint definition holds the connection details. These definitions can be created in two ways:

- Programmatically directly from the application code, as described here: [“Creating a WCF custom channel by supplying binding and endpoint information programmatically” on page 580](#)
- Administratively, by providing the details in an application configuration file, as described in the following procedure.

The client or service application configuration file is commonly named *yourappname.exe.config* where *yourappname* is the name of your application. The application configuration file is most easily modified by using the Microsoft service configuration editor tool called *SvcConfigEditor.exe* in the following way:

- Start the *SvcConfigEditor.exe* configuration editor tool. The default installation location for the tool is: *Drive:\Program Files\Microsoft SDKs\Windows\v6.0\Bin\SvcConfigEditor.exe* where *Drive:* is the name of the installation drive.

Step 1: Add a binding element extension to enable WCF to locate the custom channel

1. Right-click **Advanced** > **Extension** > **binding element** to open the menu, and select **New**
2. Complete the fields as shown in this table:

Table 73. New binding element fields	
Field	Value
Name	IBM.XMS.WCF.SoapJmsIbmTransportChannel
Type	Navigate to IBM.XMS.WCF.dll in the Global Assembly Cache (GAC) and select IBM.XMS.WCFSoapJmsIbmTransportBindingElementConfig

Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder

1. Right-click **Bindings** to open the menu, and select **New Binding Configuration**
2. Complete the fields as shown in this table:

Table 74. New binding configuration fields	
Field	Value
Name	CustomBinding_WMQ
BindingElement 1	textMessageEncoding (MessageVersion: Soap11)
BindingElement 2	IBM.XMS.WCF.SoapJmsIbmTransportChannel

Step 3: Specify the binding properties

1. Select the *IBM.XMS.WCF.SoapJmsIbmTransportChannel* transport binding from the binding you created in: [“Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder” on page 578](#)
2. Make any required changes to the default values of the properties as described in: [“WCF binding configuration options” on page 582](#)

Step 4: Create an endpoint definition

Create an endpoint definition which references the custom binding you created in: [“Step 2: Create a custom binding definition which pairs the custom channel with a WCF message encoder” on page 578](#) and provides the connection details of the service. The way this information is specified is dependent on whether the definition is for a client application or a service application.

For a client application, add an endpoint definition to the client section as follows:

1. Right-click **Client** > **Endpoints** to open the menu, and select **New Client Endpoint**
2. Complete the fields as shown in this table:

Table 75. New client endpoint fields	
Field	Value
Name	Endpoint_WMQ
Address	The SOAP/JMS URI describing the WMQ connection details required to access the service. For further details see: “WebSphere MQ custom channel for WCF endpoint URI address format” on page 581
Binding	customBinding
BindingConfiguration	CustomBinding_WMQ
Contract	The name of your service contract interface

For a service application, add a service definition to the services section as follows:

1. Right-click **Services** to open the menu, and select **New Service**, then select the service class to be hosted.
2. Add an endpoint definition to the **Endpoints** section for your new service, and complete the fields as shown in this table:

Table 76. New service endpoint fields	
Field	Value
Name	Endpoint_WMQ
Address	The SOAP/JMS URI describing the WMQ connection details required to access the service. For further details see: “WebSphere MQ custom channel for WCF endpoint URI address format” on page 581
Binding	customBinding
BindingConfiguration	CustomBinding_WMQ
Contract	The name of your service implementation class

Creating a WCF custom channel by supplying binding and endpoint information programmatically

The WebSphere MQ custom channel for WCF is a transport level WCF channel. An endpoint and binding must be defined to use the custom channel, and these definitions can be done programmatically directly from the application code.

To configure and use the WebSphere MQ custom channel for WCF, which is a transport level WCF channel, a binding and an endpoint definition must be defined. The binding holds the configuration information for the channel, and the endpoint definition holds the connection details. For more information, see: [“Using the WCF samples” on page 588](#)

These definitions can be created in two ways:

- Administratively, by providing the details in an application configuration file, as described here: [“Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file” on page 578](#)
- Programmatically directly from the application code, as described in the following example.

Step 1: Create an instance of the transport binding element of the channel

Add the following code to your application:

```
SoapJmsIbmTransportBindingElement transportBindingElement = new  
SoapJmsIbmTransportBindingElement();
```

Step 2: Set binding properties

Set any required binding properties, for example, by adding the following code to your application to set the `ClientConnectionMode`.

```
transportBindingElement.ClientConnectionMode = XmsWCFBindingProperty.AS_URI;
```

Step 3: Create a custom binding that pairs the transport channel with a message encoder

Create a custom binding by adding the following code to your application:

```
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(),  
transportBindingElement);
```

Step 4: Creating the SOAP/JMS URI

The SOAP/JMS URI which describes the WebSphere MQ connection details required to access the service, must be provided as the endpoint address. This depends on whether the channel is being used for a service application or a client application.

For client applications, the SOAP/JMS URI must be created as an `EndpointAddress` as follows:

```
EndpointAddress address = new EndpointAddress("jms:/queue?  
destination=SampleQ@QM1&connectionFactory=connectQueueManager(QM1)&initialContextFactory=com.ibm  
.mq.jms.Nojndi");
```

For service applications, the SOAP/JMS URI must be created as a `Uri` as follows:

```
Uri address = new Uri("jms:/queue?  
destination=SampleQ@QM1&connectionFactory=connectQueueManager(QM1)&initialContextFactory=com.ibm  
.mq.jms.Nojndi");
```

For more information about the endpoint address, see: [“WebSphere MQ custom channel for WCF endpoint URI address format” on page 581](#)

WebSphere MQ custom channel for WCF endpoint URI address format

A Universal Resource Identifier (URI) provides location and connection details to specify a web service. This URI format permits a comprehensive degree of control over SOAP/ WebSphere MQ-specific parameters and options when accessing target services.

A web service is specified using a Universal Resource Identifier (URI). This section specifies the URI format that is supported in WebSphere MQ transport for SOAP. This URI format permits a comprehensive degree of control over SOAP/WebSphere MQ-specific parameters and options when accessing target services. This format is compatible with WebSphere Application Server (WAS) and with CICS facilitating the integration of WebSphere MQ with both those products.

The URI syntax is as follows:

```
jms:/queue?name=value&name=value...
```

where name is a parameter name and value is an appropriate value, and the name=value element can be repeated any number of times with the second and subsequent occurrences being preceded by an ampersand (&).

For further information about setting URI properties, see: [URI syntax and parameters for Web service deployment](#)

Parameter names are case-sensitive, as are names of WebSphere MQ objects. If any parameter is specified more than once, the final occurrence of the parameter takes effect meaning client applications can override parameter values by appending to the URI. If any additional unrecognised parameters are included, they are ignored.

If you store a URI in an XML string, you must represent the ampersand character as "&". Similarly, if a URI is coded in a script, take care to escape characters such as & which would otherwise be interpreted by the shell.

This is an example of a simple URI for an Axis service:

```
jms:/queue?destination=myQ&connectionFactory=()
&initialContextFactory=com.ibm.mq.jms.Nojndi
```

Here is an example of a simple URI for a .NET service:

```
jms:/queue?destination=myQ&connectionFactory=()&targetService=MyService.asmx
&initialContextFactory=com.ibm.mq.jms.Nojndi
```

Only the required parameters are supplied (targetService is required for .NET services only), and connectionFactory is given no options.

In this Axis example, connectionFactory contains a number of options:

```
jms:/queue?destination=myQ@myRQM&connectionFactory=connectQueueManager(myconnQM)
binding(client)clientChannel(myChannel)clientConnection(myConnection)
&initialContextFactory=com.ibm.mq.jms.Nojndi
```

In this Axis example, the sslPeerName option of connectionFactory has also been specified. The value of sslPeerName itself contains name value pairs and significant embedded blanks:

```
jms:/queue?destination=myQ@myRQM&connectionFactory=connectQueueManager(myconnQM)
binding(client)clientChannel(myChannel)clientConnection(myConnection)
sslPeerName(CN=MQ Test 1,O=IBM,S=Hampshire,C=GB)
&initialContextFactory=com.ibm.mq.jms.Nojndi
```

WCF binding configuration options

This topic describes how configuration options can be applied to the custom channels binding information, and lists the options that are available.

The binding configuration options can be set in one of two different ways:

1. Administratively: The binding property settings must be specified in the transport section of the custom binding definition in the applications configuration file, for example: `app.config`
2. Programmatically: The application code must be modified to specify the property during initialization of the custom binding.

Setting the binding properties administratively

The binding property settings can also be specified in the application config file, for example: `app.config`. The config file is generated by **svcutil**, for example:

```
<customBinding>
...
  <IBM.XMS.WCF.SoaPjmsIbmTransportChannel maxBufferPoolSize="524288"
    maxMessageSize="4000000" clientConnectionMode="0" maxConcurrentCalls="16"/>
...
</customBinding>
```

Setting the binding properties programmatically

To add a WCF binding property to specify the client connection mode, you must modify the service code to specify the property during initialization of the custom binding.

Use the following example to specify unmanaged client connection mode:

```
SoapJmsIbmTransportBindingElement transportBindingElement = new
SoapJmsIbmTransportBindingElement();
transportBindingElement.ClientConnectionMode = XmsWCFBindingProperty.CLIENT_UNMANAGED;

Binding sampleBinding = new CustomBinding(new TextMessageEncodingBindingElement(),
                                           transportBindingElement);
```

WCF binding properties

Property name	Client or Service application	Administrative value	Programmatic value	Description
maxBufferPoolSize	Both	0 to 64 bit signed integer	0 to 64 bit signed integer	Specifies the maximum size of the memory that can be used to store WCF message buffers for an instance of the channel.
maxMessageSize	Both	1 to 32 bit signed integer	1 to 32 bit signed integer	Specifies the maximum memory that can be used for an individual WCF message.

Property name	Client or Service application	Administrative value	Programmatic value	Description
clientConnectionMode	Both	0 (Default value) 1	AS_URI (Default value) CLIENT_UNMANAGED	Specifies the client connection mode of the transport channel. 0 means that the client connection mode is as specified in the URI. Only used if the client connection is used. Specifies that the client connection mode is as specified in the URI. 0 is the default value if no client connection mode is set. 1 means that the client connection mode is an unmanaged client. Only used if the client connection is used.
MaxConcurrentCalls	Client	The range is 0 - 2 147 483 647 16 is the default value	The range is 0 - 2 147 483 647 16 is the default value	This property defines the maximum number of concurrent operations that can take place on an individual client proxy at any one time. If more operations are started, they are queued until an in-progress operation either completes or times out. This setting can be used to control the maximum threads and resources which can be consumed by an individual proxy. 0 removes this limit, enabling all operations to be attempted concurrently.

Property name	Client or Service application	Administrative value	Programmatic value	Description
MaxConcurrentCalls	Service	The range is 1 - 2 147 483 647 16 is the default value	The range is 1 - 2 147 483 647 16 is the default value	This property is only used if the assured delivery feature is enabled (For more information about assured delivery, see “WCF custom channel assured delivery” on page 573). It specifies the maximum number of concurrent operations that can be in progress at the same time for the given endpoint. Care is needed when changing this setting. Each concurrent operation requires additional resources, in particular a new instance of the custom channel and the associated threads from the thread pool to action the requests. Over-allocating can be counter productive and affect performance severely. Appropriate configuration of the thread pool must be made to support this property.

Building and hosting services for WCF

Overview of Microsoft Windows Communication Foundation (WCF) services explaining how to create and configure WCF services.

The IBM WebSphere MQ custom channel for WCF and the WCF services which use it, can be hosted by the following methods:

- Self-hosting
- Windows Service

The IBM WebSphere MQ custom channel for WCF cannot be hosted in Windows Process Activation Service.

The following topics provide some simple self-hosting examples to demonstrate the steps involved. The Microsoft WCF online documentation, which contains further information and the latest details, can be found on the Microsoft MSDN website at <https://msdn.microsoft.com>.

Building WCF service applications using method 1: Self-hosting administratively using an application configuration file

Having created an application configuration file, open an instance of the service and add the specified code to your application.

Before you begin

Create or edit an application configuration file for the service, as described in: [“Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file” on page 578](#)

About this task

1. Instantiate and open an instance of the service in the service host. The service type must be the same as the service type specified in the service configuration file.
2. Add the following code to your application:

```
ServiceHost service = new ServiceHost(typeof(MyService));
service.Open();
...
service.Close();
```

Building WCF service applications using method 2: Self-hosting programmatically directly from the application

Add the binding properties, create the service host with an instance of the required service class and open the service.

Before you begin

1. Add a reference to the custom channel `IBM.XMS.WCF.dll` file to the project. The `IBM.XMS.WCF.dll` is in the `WMQInstallDir\bin` where `WMQInstallDir` is the directory that WebSphere MQ 7 is installed in.
2. Add a `using` statement to the `IBM.XMS.WCF` namespace, for example: `using IBM.XMS.WCF`
3. Create an instance of the channels binding element and endpoint as described in: [“Creating a WCF custom channel by supplying binding and endpoint information programmatically” on page 580](#)

About this task

If changes to the binding properties of the channel are required, then complete the following steps:

1. Add the binding properties to `transportBindingElement` as shown in the following example:

```
SoapJmsIbmTransportBindingElement transportBindingElement = new
SoapJmsIbmTransportBindingElement();
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(),
transportBindingElement);
Uri address = new Uri("jms:/queue?destination=SampleQ@QM1&connectionFactory=
connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.Nojndi");
```

2. Create the service host with an instance of the required service class:

```
ServiceHost service = new ServiceHost(typeof(MyService));
```

3. Open the service:

```
service.AddServiceEndpoint(typeof(IMyServiceContract), binding, address);
service.Open();
...
service.Close();
```

Exposing metadata using an HTTP endpoint

Instructions for exposing the metadata of a service which is configured to use the WebSphere MQ custom channel for WCF.

About this task

If the services metadata must be exposed (so that tools such as `svcutil` can access it directly from the running service rather than from an offline WSDL file for example) it must be done by exposing the services metadata with an HTTP endpoint. The following steps can be used to add this additional endpoint.

1. Add the base address of where the metadata must be exposed to the `ServiceHost`, for example:

```
ServiceHost service = new ServiceHost(typeof(TestService),
    new Uri("http://localhost:8000/MyService"));
```

2. Add the following code to the ServiceHost before the service is opened:

```
ServiceMetadataBehavior metadataBehavior = new ServiceMetadataBehavior();
metadataBehavior.HttpGetEnabled = true;
service.Description.Behaviors.Add(metadataBehavior);
service.AddServiceEndpoint(typeof(IMetadataExchange),
    MetadataExchangeBindings.CreateMexHttpBinding(), "mex");
```

Results

The metadata is now available at the following address: <http://localhost:8000/MyService>

Building client applications for WCF

Overview of generating and building Microsoft Windows Communication Foundation (WCF) client applications.

A client application can be created for a WCF service; client applications are typically generated by using the Microsoft ServiceModel Metadata Utility Tool (Svcutil.exe) to create the required configuration and proxy files which can be used directly by the application.

Generating a WCF client proxy and application configuration files using the svcutil tool with metadata from a running service

Instructions for using the Microsoft svcutil.exe tool to generate a client for a service which is configured to use the WebSphere MQ custom channel for WCF.

Before you begin

There are three prerequisites for using the svcutil tool to create the required configuration and proxy files which can be used directly by the application:

- The WCF service must be running before the svcutil tool is started.
- The WCF service must expose its metadata using an HTTP port in addition to the WebSphere MQ custom channel endpoint references to generate a client directly from a running service.
- The custom channel must be registered in the configuration data for svcutil.

About this task

The following steps explain how to generate a client for a service which is configured to use the WebSphere MQ custom channel, but also exposes its metadata at run time via a separate HTTP port:

1. Start the WCF service (The service must be running before the svcutil tool is started).
2. Add the details from the svcutil.exe config file from the root of the installation, into the active svcutil configuration file, typically `C:\Program Files\Microsoft SDKs\Windows\v6.0A\bin\svcutil.exe.config` so svcutil recognizes the WebSphere MQ custom channel.
3. Run svcutil from a command prompt, for example:

```
svcutil /language:C# /r:<installlocation>\bin\IBM.XMS.WCF.dll
/config:app.config http://localhost:8000/IBM.XMS.WCF/samples
```

4. Copy the generated app.config and YourService.cs files to the Microsoft Visual studio client project.

What to do next

If the services metadata cannot be directly retrieved, svcutil can be used to generate the client files from wsdl instead. For more information see: [“Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL” on page 587](#)

Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL

Instructions for generating WCF clients from WSDL if the metadata of the service is unavailable.

If the metadata of the service cannot be directly retrieved to generate a client from the metadata from a running service, then, svcutil can be used to generate the client files from WSDL instead. The following modifications must be made to the WSDL to specify that the WebSphere MQ custom channel is to be used:

1. Add the following namespace definitions and policy information:

```
<wsdl:definitions
  xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd">

  <wsp:Policy wsu:Id="CustomBinding_IWMQSampleContract_policy">
    <wsp:ExactlyOne>
      <wsp:All>
        <xms:xms xmlns:xms="http://sample.schemas.ibm.com/policy/xms" />
      </wsp:All>
    </wsp:ExactlyOne>
  </wsp:Policy>

  ...

</wsdl:definitions>
```

2. Modify the bindings section to refer to the new policy section and remove any transport definition from the underlying binding element:

```
<wsdl:definitions ...>

  <wsdl:binding ...>
    <wsp:PolicyReference URI="#CustomerBinding_IWMQSampleContract_policy" />
    <[soap]:binding ... transport="" />
  </wsdl:binding>

</wsdl:definitions>
```

3. Run svcutil from a command prompt, for example:

```
svcutil /language:C# /r:MQ_INSTALLATION_PATH\bin\IBM.XMS.WCF.dll
        /config:app.config
MQ_INSTALLATION_PATH\src\samples\WMQAxis\default\service\soap.server.stockQuoteAxis_Wmq.wsdl
```

Where `MQ_INSTALLATION_PATH` is the installation directory of WebSphere MQ.

Building WCF client applications using a client proxy with an application configuration file

Before you begin

Create or edit an application configuration file for the client, as described in: [“Creating a WCF custom channel administratively by supplying binding and endpoint information in an application configuration file” on page 578](#)

About this task

Instantiate and open an instance of the client proxy. The parameter passed to the generated proxy must be the same as the endpoint name specified in the client configuration file, for example `Endpoint_WMQ`:

```
MyClientProxy myClient = new MyClientProxy("Endpoint_WMQ");
try {
    myClient.myMethod("HelloWorld!");
    myClient.Close();
}
catch (TimeoutException e) {
    Console.Out.WriteLine(e);
    myClient.Abort();
}
```

```

    }
    catch (CommunicationException e) {
        Console.Out.WriteLine(e);
        myClient.Abort();
    }
    catch (Exception e) {
        Console.Out.WriteLine(e);
        myClient.Abort();
    }
}

```

Building WCF client applications using a client proxy with programmatic configuration

Before you begin

1. Add a reference to the custom channel `IBM.XMS.WCF.dll` file to the project. The `IBM.XMS.WCF.dll` is in the `WMQInstallDir\bin` directory where `WMQInstallDir` is the directory that WebSphere MQ 7 is installed in.
2. Add a `using` statement to the `IBM.XMS.WCF` namespace, for example: `using IBM.XMS.WCF`
3. Create an instance of the binding element and endpoint of the channel as described in: [“Creating a WCF custom channel by supplying binding and endpoint information programmatically” on page 580](#)

About this task

If changes to the binding properties of the channel are required, then complete the following steps:

1. Add the binding properties to `transportBindingElement` as shown in the following figure:

```

SoapJmsIbmTransportBindingElement transportBindingElement = new
SoapJmsIbmTransportBindingElement();
Binding binding = new CustomBinding(new TextMessageEncodingBindingElement(),
transportBindingElement);
EndpointAddress address =
    new EndpointAddress("jms:/queue?destination=SampleQ@QM1&connectionFactory=
connectQueueManager(QM1)&initialContextFactory=com.ibm.mq.jms.NoJndi");

```

2. Create the client proxy as shown in the following figure, where *binding* and *endpoint address* are the binding and endpoint address configured in step [“1” on page 588](#) and passed in:

```

MyClientProxy myClient = new MyClientProxy(binding, endpoint address);
try {
    myClient.myMethod("HelloWorld!");
    myClient.Close();
}
catch (TimeoutException e) {
    Console.Out.WriteLine(e);
    myClient.Abort();
}
catch (CommunicationException e) {
    Console.Out.WriteLine(e);
    myClient.Abort();
}
catch (Exception e) {
    Console.Out.WriteLine(e);
    myClient.Abort();
}
}

```

Using the WCF samples

The Windows Communication Foundation (WCF) samples provide some simple examples of how the WebSphere MQ custom channel can be used.

To build the sample projects, either the Microsoft .NET 3.5 SDK, or Microsoft Visual Studio 2008 is needed.

Simple one-way client and server WCF sample

This sample demonstrates the WebSphere MQ custom channel being used to start a Windows Communication foundation (WCF) service from a WCF client using a one-way channel shape.

About this task

The service implements a single method which outputs a string to the console. The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in [“Generating a WCF client proxy and application configuration files using the svcutil tool with metadata from a running service”](#) on page 586

The sample has been configured with specific resource names as described in the following procedure. If you must change the resource names, then you must also change the corresponding value on the client application in the

`MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\client\app.config` file, and on the service application in the

`MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\service\TestServices.cs` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM WebSphere MQ. For more information about formatting the JMS endpoint URI, see *WebSphere MQ Transport for SOAP* in the WebSphere MQ product documentation. If you need to modify the sample solution and source, then you need an IDE, for example, Microsoft Visual Studio 8 or higher.

Procedure

1. Create a queue manager called *QM1*
2. Create a queue destination called *SampleQ*
3. Start the service so the listener is waiting for messages: Run the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\service\bin\Release\TestService.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM WebSphere MQ.
4. Run the client once: Run the `MQ_INSTALLATION_PATH\tools\dotnet\samples\cs\wcf\samples\WCF\oneway\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for IBM WebSphere MQ.

The client application loops five times sending five messages to *SampleQ*

Results

The service application gets the messages from *SampleQ* and displays Hello World on the screen five times.

What to do next

Simple request-reply client and server WCF sample

This sample demonstrates the WebSphere MQ custom channel being used to start a Windows Communication foundation (WCF) service from a WCF client using a request-reply channel shape.

About this task

This service provides some simple calculator methods to add and subtract two numbers, and then return the result. The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in [“Generating a WCF client proxy and application configuration files using the svcutil tool with metadata from a running service”](#) on page 586

The sample has been configured with specific resource names as in the following procedure described. If you need to change the resource names, then you also need to change the corresponding value on the client application in the

`MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\service\RequestReplyService.cs` file, where `MQ_INSTALLATION_PATH` is the installation directory for WebSphere MQ. For more information about formatting the JMS endpoint URI, see *WebSphere MQ Transport for SOAP* in the WebSphere MQ product documentation. If you need to modify the sample solution and source, then you need an IDE, for example, Microsoft Visual Studio 8 or higher.

Procedure

1. Create a queue manager called *QM1*
2. Create a queue destination called *SampleQ*
3. Create a queue destination called *SampleReplyQ*
4. Start the service so the listener is waiting for messages: Run the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\service\bin\Release\SimpleRequestReply_Service.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for WebSphere MQ.
5. Run the client once: Run the `MQ_INSTALLATION_PATH\Tools\wcf\samples\WCF\requestreply\client\bin\Release\SimpleRequestReply_Client.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for WebSphere MQ.

Results

When the client has been run, the following process is started and repeats four times so a total of five messages are sent each way:

1. The client puts a request message on *SampleQ* and waits for a response.
2. The service gets the request message from *SampleQ*.
3. The service adds and subtracts some values using the contents of the message.
4. The service then puts the results into a message on *SampleReplyQ*, and waits for the client to put a new message.
5. The client gets the message from *SampleReplyQ* and displays the results on the screen.

What to do next

WCF client to a .NET service hosted by WebSphere MQ sample

Sample client applications and sample service proxy applications are supplied for both .NET and Java. The samples are based on a Stock Quote service that takes a request for a stock quote and then provides the stock quote.

Before you begin

The sample requires that the .NET SOAP over JMS service hosting environment is correctly installed and configured in WebSphere MQ and is accessible from a local queue manager. For information about installing and configuring the environment, see: [“Installing WebSphere MQ Web transport for SOAP” on page 915](#)

When the .NET SOAP over JMS service hosting environment is correctly installed and configured in WebSphere MQ and is accessible from a local queue manager, additional configuration steps must be completed.

1. Set the `WMQSOAP_HOME` environment variable to the WebSphere MQ installation directory, for example: `C:\Program Files\IBM\WebSphere MQ`
2. Ensure that the Java compiler `javac` is available and on the `PATH`.

3. Copy the file `axis.jar` from the `prereqs/axis` directory of the WebSphere installation CD to the WebSphere MQ production directory, for example: `C:\Program Files\IBM\WebSphere MQ\java\lib\soap`
4. Add to the PATH: `MQ_INSTALLATION_PATH\Java\lib` where `MQ_INSTALLATION_PATH` represents the directory where WebSphere MQ is installed, for example: `C:\Program Files\IBM\WebSphere MQ`
5. Ensure that the location of .NET is specified correctly in `MQ_INSTALLATION_PATH\bin\amqwcallsdl.cmd` where `MQ_INSTALLATION_PATH` represents the directory where WebSphere MQ is installed, for example: `C:\Program Files\IBM\WebSphere MQ`. The location of .NET can be specified for example: `set msfwdir=%ProgramFiles%\Microsoft Visual Studio .NET 2003\SDK\v1.1\Bin`

When the previous steps are complete, test and run the service:

1. Navigate to your SOAP over JMS working directory.
2. Enter one of the following commands to run the verification test and leave the service listener running:
 - For .NET: `MQ_INSTALLATION_PATH\Tools\soap\samples\runivt dotnet hold` where `MQ_INSTALLATION_PATH` represents the directory where WebSphere MQ is installed.
 - For AXIS: `MQ_INSTALLATION_PATH\Tools\soap\samples\runivt Dotnet2AxisClient hold` where `MQ_INSTALLATION_PATH` represents the directory where WebSphere MQ is installed.

The `hold` argument keeps the listeners running after the test completes.

If errors are reported during this configuration, you can remove all the changes so that the procedure may be restarted in the following way:

1. Delete the generated SOAP over JMS directory.
2. Delete the queue manager.

About this task

This sample demonstrates a connection from a WCF client to the .NET SOAP over JMS sample service provided in WebSphere MQ using a one-way channel shape. The service implements a simple StockQuote example, which outputs a text string to the console.

The client has been generated by using WSDL to generate client files as described in [“Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL” on page 587](#)

The sample has been configured with specific resource names as described in the following procedure. If you need to change the resource names, then you must also change the corresponding value on the client application in the

`MQ_INSTALLATION_PATH\tools\wcf\samples\WMQNET\default\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WMQNET\default\service\WmqDefaultSample_StockQuoteDotNet.wsdl` file, where `MQ_INSTALLATION_PATH` represents the installation directory for WebSphere MQ. For more information about formatting the JMS endpoint URI, see *WebSphere MQ Transport for SOAP* in the WebSphere MQ product documentation.

Procedure

Run the client once: Run the

`MQ_INSTALLATION_PATH\tools\wcf\samples\WMQNET\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` represents the installation directory for WebSphere MQ.

The client application loops five times sending five messages to the sample queue.

Results

The service application gets the messages from the sample queue and displays Hello World five times on the screen.

WCF client to an Axis Java service hosted by WebSphere MQ sample

Sample client applications and sample service proxy applications are supplied for both Java and .NET. The samples are based on a Stock Quote service that takes a request for a stock quote and then provides the stock quote.

Before you begin

This sample requires that the .NET SOAP over JMS service hosting environment is correctly installed and configured in WebSphere MQ and is accessible from a local queue manager. For information about installing and configuring the environment, see: [“Installing WebSphere MQ Web transport for SOAP” on page 915](#)

When the .NET SOAP over JMS service hosting environment is correctly installed and configured in WebSphere MQ and is accessible from a local queue manager, additional configuration steps must be completed.

1. Set the WMQSOAP_HOME environment variable to the WebSphere MQ installation directory, for example: C:\Program Files\IBM\WebSphere MQ
2. Ensure that the Java compiler javac is available and on the PATH.
3. Copy the file axis.jar from the prereqs/axis directory of the WebSphere installation CD to the WebSphere MQ installation directory.
4. Add to the PATH: MQ_INSTALLATION_PATH\Java\lib where MQ_INSTALLATION_PATH represents the directory where WebSphere MQ is installed, for example: C:\Program Files\IBM\WebSphere MQ
5. Ensure that the location of .NET is specified correctly in MQ_INSTALLATION_PATH\bin\amqwcallsdls.cmd where MQ_INSTALLATION_PATH represents the directory where WebSphere MQ is installed, for example: C:\Program Files\IBM\WebSphere MQ. The location of .NET can be specified for example: set msfwdir=%ProgramFiles%\Microsoft Visual Studio .NET 2003\SDK\v1.1\Bin

When the previous steps are complete, test and run the service:

1. Navigate to your SOAP over JMS working directory.
2. Enter one of the following commands to run the verification test and leave the service listener running:
 - For .NET: MQ_INSTALLATION_PATH\Tools\soap\samples\runivt dotnet hold where MQ_INSTALLATION_PATH represents the directory where WebSphere MQ is installed.
 - For AXIS: MQ_INSTALLATION_PATH\Tools\soap\samples\runivt Dotnet2AxisClient hold where MQ_INSTALLATION_PATH represents the directory where WebSphere MQ is installed.

The hold argument keeps the listeners running after the test completes.

If errors are reported during this configuration, you can remove all the changes so that the procedure is restarted in the following way:

1. Delete the generated SOAP over JMS directory.
2. Delete the queue manager.

About this task

The sample demonstrates a connection from a WCF client to the Axis Java SOAP over JMS sample service provided in WebSphere MQ using a one-way channel shape. The service implements a simple StockQuote example, which outputs a text string to a file which is saved in the current directory.

The client has been generated by using WSDL to generate client files as described in [“Generating a WCF client proxy and application configuration files using the svcutil tool with WSDL” on page 587](#)

The sample has been configured with specific resource names as described in this paragraph. If you need to change the resource names, then you must also change the corresponding value on the client application in the MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\client\app.config file, and

on the service application in the
`MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\service\WmqDefaultSample_StockQuoteDotNet.wsdl` file, where `MQ_INSTALLATION_PATH` represents the installation directory for WebSphere MQ.

Procedure

Run the client once: Run the
`MQ_INSTALLATION_PATH\tools\wcf\samples\WMQAxis\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` represents the installation directory for WebSphere MQ.

The client application loops five times sending five messages to the sample queue.

Results

The service application gets the messages from the sample queue and adds Hello World five times to a file in the current directory.

Related reference

[“Handling different SOAP response element names” on page 601](#)

WCF expects the name of a returned value to be in a specific format by default, but a service might not return an element with its name in the expected format.

WCF client to Java service hosted by WebSphere Application Server sample

Sample client applications and sample service proxy applications are supplied for WebSphere Application Server (WAS) 6. A request-response service is also provided.

Before you begin

This sample requires that the following WebSphere MQ configuration is used:

Table 77. WebSphere MQ required configuration	
Object	Required name
Queue manager	QM1
Local queue	HelloWorld
Local queue	HelloWorldReply

This sample also requires that a WebSphere Application Server V6 hosting environment is correctly installed and configured. WebSphere Application Server V6 uses a bindings mode connection to connect to WebSphere MQ by default. Therefore WebSphere Application Server V6 must be installed on the same machine as the queue manager.

After the WAS environment is configured, the following additional configuration steps must be completed:

1. Create the following JNDI objects in the WebSphere Application Server JNDI repository:

- a. A JMS queue destination called HelloWorld
 - Set the JNDI name to `jms/HelloWorld`
 - Set the queue name to HelloWorld
- b. A JMS queue connection factory called HelloWorldQCF
 - Set the JNDI name to `jms/HelloWorldQCF`
 - Set the queue manager name to QM1
- c. A JMS queue connection factory called WebServicesReplyQCF
 - Set the JNDI name to `jms/WebServicesReplyQCF`
 - Set the queue manager name to QM1

2. Create a Message Listener Port called HelloWorldPort in WebSphere Application Server with the following configuration:
 - Set the connection factory JNDI name to `jms/HelloWorldQCF`
 - Set the destination JNDI name to `jms/HelloWorld`
3. Install the web service HelloWorldEJBear.ear application to your WebSphere Application Server as follows:
 - a. Click **Applications > New Application > New Enterprise Application**.
 - b. Navigate to `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\HelloWorldsEJBear.ear` where `MQ_INSTALLATION_PATH` is the installation directory of WebSphere MQ.
 - c. Do not change any of the default option in the wizard, and restart the application server after the application has been installed.

When the WAS configuration is complete, test the service by running it once:

1. Navigate to your Soap over JMS working directory.
2. Enter this command to run the sample:
`MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\TestClient.exe` where `MQ_INSTALLATION_PATH` is the installation directory of WebSphere MQ.

About this task

The sample demonstrates a connection from a WCF client to the WebSphere Application Server SOAP over JMS sample service provided in the WCF samples included in WebSphere MQ V7, using a request-response channel shape. Messages flow between WCF and the WebSphere Application Server using WebSphere MQ queues. The service implements the `HelloWorld(...)` method, which takes a string and returns a greeting to the client.

The client has been generated by using the `svcutil` tool to retrieve the service metadata from a separately exposed HTTP endpoint as described in [“Generating a WCF client proxy and application configuration files using the svcutil tool with metadata from a running service”](#) on page 586

The sample has been configured with specific resource names as described in the following procedure. If you need to change the resource names, then you must also change the corresponding value on the client application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\default\client\app.config` file, and on the service application in the `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\HelloWorldsEJBear.ear` where `MQ_INSTALLATION_PATH` is the installation directory of WebSphere MQ. For more information about formatting the JMS endpoint URI, see [URI syntax and parameters for Web service deployment](#).

The service and client are based upon the service and client outlined in the IBM Developer article *Building a JMS Web service using SOAP over JMS and WebSphere Studio*. If you want to learn more about developing SOAP over JMS Web services which are compatible with the WebSphere MQ WCF custom channel, the relevant article can be found at: https://www.ibm.com/developerworks/websphere/library/techarticles/0402_du/0402_du.html.

Procedure

Run the client once: Run the `MQ_INSTALLATION_PATH\tools\wcf\samples\WAS\default\client\bin\Release\TestClient.exe` file, where `MQ_INSTALLATION_PATH` is the installation directory for WebSphere MQ.

The client application starts both of the service methods at the same time, sending two messages to the sample queue.

Results

The service application gets the messages from the sample queue and provides a response to the `HelloWorld(...)` method call which the client application outputs to the console.

Problem determination on the WCF custom channel for WebSphere MQ

You can use WebSphere MQ trace to collect detailed information about what various parts of the WebSphere MQ code is doing. When using Windows Communication Foundation (WCF), a separate trace output is generated for the WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.

Fully enabling trace for the WCF custom channel produces two output files:

1. The WCF custom channel trace integrated with the Microsoft WCF infrastructure trace.
2. The WCF custom channel trace integrated with XMS .NET.

By having two trace outputs, problems can be tracked at each interface using the appropriate tools, for example:

- WCF problem determination using suitable Microsoft tooling.
- WebSphere MQ MQI client issues using the XMS trace format.

To simplify trace enablement, the .NET 3 TraceSource and XMS .NET trace stack are both controlled using a single interface as described in: [“WCF trace configuration and trace file names”](#) on page 595.

WCF custom channel exception hierarchy

The exceptions types thrown by the custom channel are consistent with WCF and are typically a `TimeoutException` or `CommunicationException` (or a subclass of `CommunicationException`).

Further details of the error condition, where available, are provided using linked or inner exceptions. The following exceptions are typical examples, and each layer in the architecture of the channel contributes an additional linked exception, for example; `CommunicationsException` has a linked `XMSEException`, which has a linked `MQException`:

1. `System.ServiceModel.CommunicationsExceptions`
2. `IBM.XMS.XMSEException`
3. `IBM.WMQ.MQException`

Key information is captured and provided in the data collection of the highest `CommunicationException` in the hierarchy. This capture and provision of data prevents the need for the applications to link to each layer in the architecture of the channel in order to interrogate the linked exceptions, and any additional information they might contain. The following key names are defined:

- `IBM.XMS.WCF.ErrorCode`: The error message code of the current custom channel exception.
- `IBM.XMS.ErrorCode`: The error message of the first XMS exception in the stack.
- `IBM.WMQ.ReasonCode`: The underlying WebSphere MQ reason code.
- `IBM.WMQ.CompletionCode`: The underlying WebSphere MQ completion code.

WCF trace configuration and trace file names

When trace is fully enabled, it produces two output files, one for diagnosing WCF problems, and one detailed file for internal trace diagnostic material. To simplify trace enablement, both the .NET 3 `TraceSource` and XMS .NET trace stacks use a single interface.

Two different trace methods are available for the WCF custom channel, the two trace methods are activated independently or together. Each method produces its own trace file, so when both trace methods have been activated, two trace output files are generated.

To keep the configuration and enablement as simple as possible, the same interface is used to control both trace methods. The `app.config` file must be edited to include the relevant trace configuration as described in the following section. Users can then add their own equivalent sections to combine the output with trace from their own application.

WCF custom channel tracing is not enabled by default. You must first create a trace listener, then set the required trace level for the selected trace source in the `app.config` file.

Configuring WCF custom channel with WCF infrastructure trace

Add the following section of code to the <system.diagnostics><sources> section in the app.config file:

```
<source name="IBM.XMS.WCF" switchValue="Verbose,ActivityTracing">
  <listeners>
    <remove name="Default"/>
    <add name="NewListener"/>
  </listeners>
</source>
```

The preceding piece of code makes the channel trace using the .NET 3 TraceSource. All invocations of the configuration files associated with the executable files are controlled by this piece of code.

Configuring WCF custom channel with XMS .NET trace

Configuring the XMS .NET trace requires that you add a section of code to the <system.diagnostics><sources> section in the app.config file. However, the piece of code is added to the extensible <source> element shown in the [Configuring WCF custom channel with WCF infrastructure trace](#) section. So although the WCF infrastructure trace code must be present for the XMS .NET trace to work, the WCF infrastructure trace can be disabled if it is not required, as described in the [Enabling WCF trace](#) section.

```
<source name="IBM.XMS.WCF" switchValue="Verbose, ActivityTracing"
  xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path"
  xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced">
  <listeners>
    <remove name="Default"/>
    <add name="NewListener"/>
  </listeners>
</source>
```

WCF trace configuration variables

Table 78. WCF trace configuration variables	
Variable	Description
name	Specify the name as: IBM.XMS.WCF
switchValue	The switchValue controls the trace level. When switchValue is set to Off, the WCF infrastructure TraceSource is not generated. Any other value, such as Verbose, generates TraceSource. For detailed trace level information from Microsoft, consult your WCF documentation or go to the Microsoft WCF Tracing web page: https://msdn.microsoft.com/en-us/library/ms733025(vs.85).aspx

Table 78. WCF trace configuration variables (continued)

Variable	Description
<code>xmsTraceSpecification=ComponentName=type=state</code>	<p><i>ComponentName</i> is the name of the class that you want to trace. You can use a * wildcard character in this name. For example:</p> <pre>*=all=enabled</pre> <p>specifies that you want to trace all classes, and</p> <pre>IBM.XMS.impl.*=all=enabled</pre> <p>specifies that you require API trace only.</p> <p><i>type</i> can be any of the following trace types:</p> <ul style="list-style-type: none"> • all • debug • event • EntryExit <p><i>state</i> can be either enabled or disabled.</p>
<code>xmsTraceFilePath="filename"</code>	<p>If you do not specify an <code>xmsTraceFilePath</code>, or if the <code>xmsTraceFilePath</code> is present but contains an empty string, then the trace file is placed in the current directory. To store the trace file in a named directory, specify the directory name in the <code>xmsTraceFilePath</code>, for example:</p> <pre>xmsTraceFilePath="c:\somepath"</pre>
<code>xmsTraceFileSize="size"</code>	<p>The maximum allowed size of the trace file. When a file reaches this size, it is archived and renamed. The default maximum is 20 KB, which is specified as:</p> <pre>xmsTraceFileSize="20000000".</pre>
<code>xmsTraceFileNumber="number"</code>	<p>The number of trace files that are to be retained. The default is 4 (one active file and three archive files). The minimum number allowed is two.</p>
<code>xmsTraceFormat="format"</code>	<p>There are two levels of <code>xmsTraceFormat</code>: basic and advanced. The default trace format is basic if you do not specify an <code>xmsTraceFormat</code>, or if the <code>xmsTraceFormat</code> is present but contains an empty string. Trace files are produced in this format if you specify:</p> <pre>xmsTraceFormat="basic"</pre> <p>If you require trace that is compatible with trace analyzer tools, you must specify:</p> <pre>traceFormat="advanced"</pre>

Enabling WCF trace

There are four combinations for enabling and disabling the two different trace methods. The four combinations require editing the values of the sections of coded described in the preceding sections.

There is also an environment variable that can be set; for more information see [“Enabling WCF trace with the WCF_TRACE_ON environment variable”](#) on page 599.

This table and the values shown are dependent on the pieces of code demonstrated earlier have already been added to the app.config file.

Table 79. WCF trace enablement combinations.		
Trace type	Value changed	Example
XMS trace enabled. WCF TraceSource enabled	The switchValue is not set to Off	<pre><source name="IBM.XMS.WCF" switchValue="Verbose, ActivityTracing" xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"> <listeners> <remove name="Default"/> <add name="NewListener"/> </listeners> </source></pre>
XMS trace enabled. WCF TraceSource disabled	The switchValue is set to Off and an xmsTraceSpecification has been given	<pre><source name="IBM.XMS.WCF" switchValue="Off, ActivityTracing" xmsTraceSpecification="*=all=enabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"> <listeners> <remove name="Default"/> <add name="NewListener"/> </listeners> </source></pre>
XMS trace disabled. WCF TraceSource enabled	<p>There are two ways to achieve this result:</p> <ul style="list-style-type: none"> • The switchValue variable is not set to Off and an xmsTraceSpecification has not been added • The switchValue variable is not set to Off and the xmsTraceSpecification has been set to disabled 	<pre><source name="IBM.XMS.WCF" switchValue="Verbose, ActivityTracing" xmsTraceSpecification="*=all=disabled" xmsTraceFilePath="path" xmsTraceFileSize="2000000" xmsTraceFileNumber="4" xmsTraceFormat="advanced"> <listeners> <remove name="Default"/> <add name="NewListener"/> </listeners> </source></pre>

Table 79. WCF trace enablement combinations. (continued)		
Trace type	Value changed	Example
XMS trace disabled. WCF TraceSource disabled	<p>There are three ways to achieve this result:</p> <ul style="list-style-type: none"> No <source> element in the app.config file The switchValue variable is set to Off and anxmstraceSpecification has not been added The switchValue variable is set to Off and the xmstraceSpecification has been set to disabled 	<pre><source name="IBM.XMS.WCF" switchValue="Off, ActivityTracing" xmstraceSpecification="*=all=disabled" xmstraceFilePath="path" xmstraceFileSize="2000000" xmstraceFileNumber="4" xmstraceFormat="advanced"> <listeners> <remove name="Default"/> <add name="NewListener"/> </listeners> </source></pre>

Enabling WCF trace with the WCF_TRACE_ON environment variable

As well as the preceding methods described to enable WCF trace, the XMS .NET trace can also be enabled using the WCF_TRACE_ON environment variable.

Setting the WCF_TRACE_ON environment variable to any non-null value is the equivalent of setting the xmstraceSpecification to *=all=enabled, for example: "set WCF_TRACE_ON=true"

However, if the xmstraceSpecification is explicitly set in the app.config file, then the WCF_TRACE_ON environment variable is overridden.

WCF trace output files and file names

XMS trace files are traditionally named using the base name and process ID format of: xms_trace_pid.log, where pid is the process ID.

As XMS trace files can still be produced in parallel with WCF custom channel trace files, the WCF custom channel trace integrated with XMS .NET trace output files have the following format to avoid confusion: wcfxms_trace_pid.log, where pid is the process ID.

The trace output file is created in the current working directory by default, but this destination can be redefined if necessary.

WCF XMS First Failure Support Technology (FFST)

You can collect detailed information about what various parts of the WebSphere MQ code is doing by using WebSphere MQ trace. XMS FFST has its own configuration and output files for the WCF custom channel.

XMS FFST trace files are traditionally named using the base name and process ID format of: xmsffdcpid_date.txt, where pid is the process ID and date is the time and date.

As XMS FFST trace files can still be produced in parallel with WCF custom channel XMS FFST files, the WCF custom channel XMS FFST output files have the following format to avoid confusion: wcfffdcpid_date.txt, where pid is the process ID and date is the time and date.

This trace output file is created in the current working directory by default, but this destination can be redefined if necessary.

The WCF custom channel with XMS .NET trace header is similar to the following example:

```

***** Start Display XMS WCF Environment *****
Product Name :- value
WCF Version :- value
Level :- value
***** End Display XMS WCF Environment *****

```

The FFST trace files are formatted in the standard way, without any formatting that is specific to the custom channel.

WCF version information

WCF version information aids with problem determination and is included in the assembly metadata of the custom channel.

The WebSphere MQ custom channel for WCF version metadata can be retrieved in one of three ways:

- Using the WebSphere MQ utility dspmqver. For information about how to use dspmqver, see: [dspmqver](#)
- Using the Windows Explorer properties dialog: In the Windows Explorer, right-click **IBM.XMS.WCF.dll** > **Properties** > **Version**.
- From the header information of any of the channels FFST or trace files. For more information about the FFST header information, see: [“WCF XMS First Failure Support Technology \(FFST\)” on page 599](#)

WCF hints and tips

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 might save you time if they are relevant to the work that you are doing.

Externalizing exceptions from the WCF service host

For services hosted using the WCF service host; any unhandled exceptions thrown by the service, WCF internals, or channel stack are not externalized by default. To be informed of these exceptions, an error handler must be registered.

The following code provides an example of defining the error handler service behavior which can be applied as an attribute of a service:

```

using System.ServiceModel.Dispatcher;
using System.Collections.ObjectModel;
....
public class ErrorHandlerBehaviorAttribute : Attribute, IServiceBehavior, IErrorHandler
{
    //
    // IServiceBehavior Interface
    //
    public void AddBindingParameters(ServiceDescription serviceDescription,
        ServiceHostBase serviceHostBase, Collection<ServiceEndpoint> endpoints,
        BindingParameterCollection bindingParameters)
    {
    }
    public void ApplyDispatchBehavior(ServiceDescription serviceDescription,
        ServiceHostBase serviceHostBase)
    {
        foreach (ChannelDispatcher channelDispatcher in serviceHostBase.ChannelDispatchers)
        {
            channelDispatcher.ErrorHandlers.Add(this);
        }
    }
    public void Validate(ServiceDescription serviceDescription, ServiceHostBase
serviceHostBase)
    {
    }

    //
    // IErrorHandler Interface
    //
    public bool HandleError(Exception e)
    {
        // Process the exception in the required way, in this case just outputting to the
console
    }
}

```



```

        Console.Out.WriteLine(e);

        // Always return false to allow any other error handlers to run
        return false;
    }
    public void ProvideFault(Exception error, MessageVersion version, ref Message fault)
    {
    }
}

```

Handling different SOAP response element names

WCF expects the name of a returned value to be in a specific format by default, but a service might not return an element with its name in the expected format.

WCF has the convention of expecting the returned value to be named in the following format: *methodNameResult* where *methodName* is the name of the service operation. For example, for a service called *getQuote*, WCF expects the response to be called: *getQuoteResult*.

However, the service can return an element with a name that does not conform to this format.

When running the *scvutil* tool to generate a proxy client, if the WSDL specifies a different name, then the proxy interface adds parameters to instruct WCF with the name to look for. For example:

```

[System.ServiceModel.OperationContractAttribute(Action = "", ReplyAction = "*")]
[System.ServiceModel.XmlSerializerFormatAttribute(Style =
System.ServiceModel.OperationFormatStyle.Rpc,
                                Use =
System.ServiceModel.OperationFormatUse.Encoded)]
[return: System.ServiceModel.MessageParameterAttribute(Name = "getQuoteReturn")]
float getQuote(string in0);

```

If you create your own interface (for example, by adding a request-reply method to an existing proxy interface), then you must ensure that you add the same parameters to the interface if the service returns a different name. If you do not do so, then the most common issue is that a call to the service method always returns a null value; if an object is returned, then the method returns null, but if a numeric value such as an integer is returned, then the method returns 0.

Using C++

WebSphere MQ provides C++ classes equivalent to WebSphere MQ objects and some additional classes equivalent to the array data types. It provides a number of features not available through the MQI.

As of WebSphere MQ Version 7.0, enhancements to the WebSphere MQ programming interfaces will not be applied to the C++ classes.

WebSphere MQ C++ provides the following features:

- Automatic initialization of WebSphere MQ data structures.
- Just-in-time queue manager connection and queue opening.
- Implicit queue closure and queue manager disconnection.
- Dead-letter header transmission and receipt.
- IMS bridge header transmission and receipt.
- Reference message header transmission and receipt.
- Trigger message receipt.
- CICS bridge header transmission and receipt.
- Work header transmission and receipt.
- Client channel definition.

The following Booch class diagrams show that all the classes are broadly parallel to those WebSphere MQ entities in the procedural MQI (for example using C) that have either handles or data structures. All classes inherit from the *ImqError* class (see [ImqError C++ class](#)), which allows an error condition to be associated with each object.

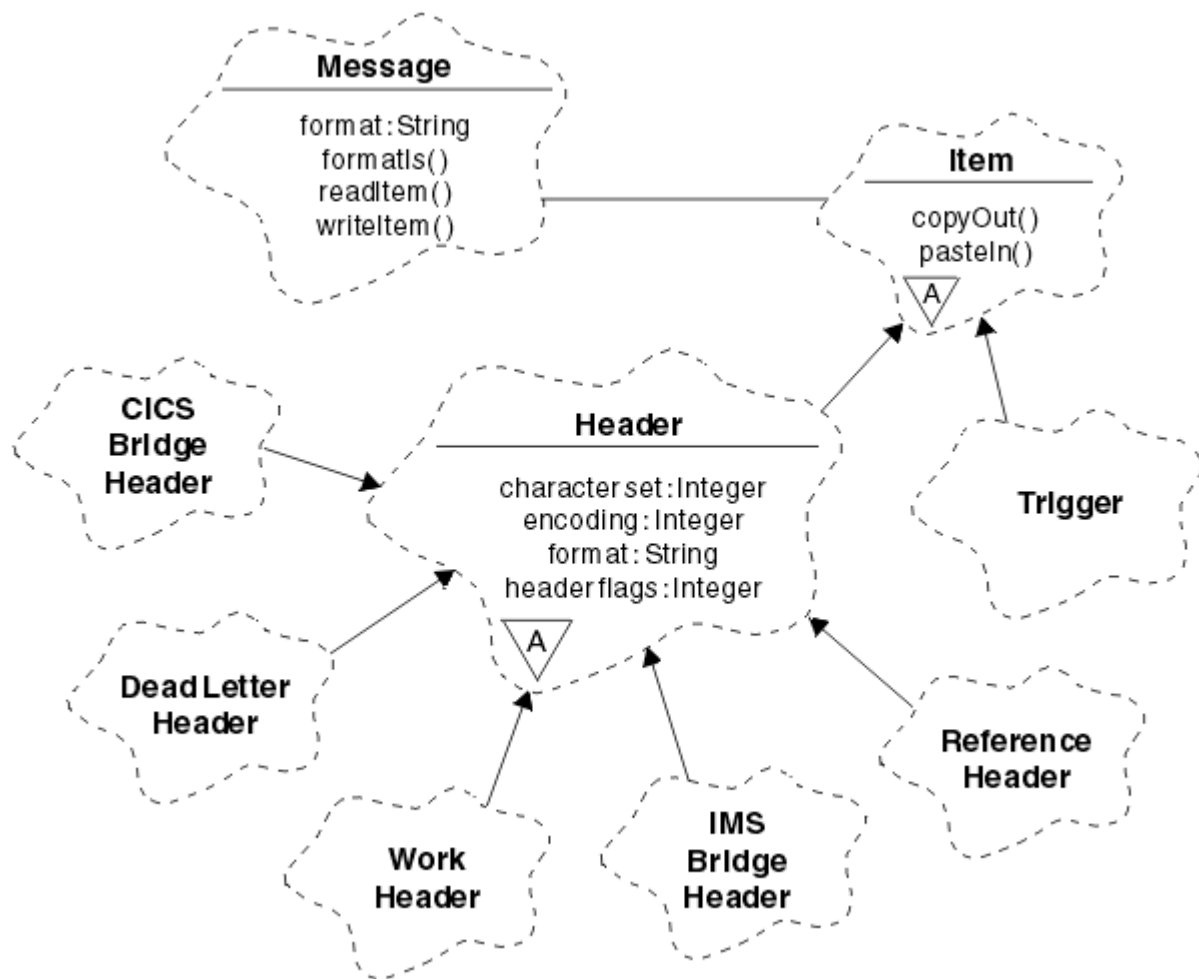


Figure 120. WebSphere MQ C++ classes (item handling)




Figure 121. WebSphere MQ C++ classes (queue management)

To interpret Booch class diagrams correctly, be aware of the following conventions:

- Methods and noteworthy attributes are shown below the *class* name.
- A small triangle within a cloud denotes an *abstract class*.
- *Inheritance* is denoted by an arrow to the parent class.
- An undecorated line between clouds denotes a *cooperative relationship* between classes.
- A line decorated with a number denotes a *referential relationship* between two classes. The number indicates the number of objects that can participate in a particular relationship at any one time.

The following classes and data types are used in the C++ method signatures of the queue management classes (see [Figure 121 on page 603](#)) and the item handling classes (see [Figure 120 on page 602](#)):

- The `ImqBinary` class (see `ImqBinary C++ class`), which encapsulates byte arrays such as `MQBYTE24`.
- The `ImqBoolean` data type, which is defined as **`typedef unsigned char ImqBoolean`**.

- The `ImqString` class (see [ImqString C++ class](#)), which encapsulates character arrays such as `MQCHAR64`.

Entities with data structures are subsumed within appropriate object classes. Individual data structure fields (see [C++ and MQI cross-reference](#)) are accessed with methods.

Entities with handles come under the `ImqObject` class hierarchy (see [ImqObject C++ class](#)) and provide encapsulated interfaces to the MQI. Objects of these classes exhibit intelligent behavior that can reduce the number of method invocations required relative to the procedural MQI. For example, you can establish and discard queue manager connections as required, or you can open a queue with appropriate options, then close it.

The `ImqMessage` class (see [ImqMessage C++ class](#)) encapsulates the MQMD data structure and also acts as a holding point for user data and *items* (see [“Reading messages in C++”](#) on page 613) by providing cached buffer facilities. You can provide fixed-length buffers for user data and use the buffer many times. The amount of data present in the buffer can vary from one use to the next. Alternatively, the system can provide and manage a buffer of flexible length. Both the size of the buffer (the amount available for receipt of messages) and the amount actually used (either the number of bytes for transmission or the number of bytes actually received) become important considerations.

Related concepts

[Technical overview](#)

[“C++ sample programs”](#) on page 604

Four sample programs are supplied, to demonstrate getting and putting messages.

[“C++ language considerations”](#) on page 608

This collection of topics details the aspects of the C++ language usage and conventions that you must consider when writing application programs that use the Message Queue Interface (MQI).

[“Preparing message data in C++”](#) on page 612

Message data is prepared in a buffer, which can be supplied by the system or the application. There are advantages to either method. Examples of using a buffer are given.

[“Deciding which programming language to use”](#) on page 75

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Developing applications”](#) on page 7

IBM WebSphere MQ provides several ways in which you can develop applications to send and receive messages that you need to support your business processes. You can also develop applications to manage your queue managers and related resources.

Related reference

[“Building WebSphere MQ C++ programs”](#) on page 618

The URL of supported compilers is listed, together with the commands to use to compile, link and run C++ programs and samples on WebSphere MQ platforms.

[C++ and MQI cross reference](#)

[WebSphere MQ C++ classes](#)

C++ sample programs

Four sample programs are supplied, to demonstrate getting and putting messages.

The sample programs are:

- HELLO WORLD (`imqwrld.cpp`)
- SPUT (`imqsput.cpp`)
- SGET (`imqsget.cpp`)
- DPUT (`imqdput.cpp`)

The sample programs are located in the directories shown in [Table 80](#) on page 605.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Table 80. Location of sample programs

Environment	Directory containing source	Directory containing built programs
AIX	<code>MQ_INSTALLATION_PATH/samp</code>	<code>MQ_INSTALLATION_PATH/samp/bin/ia</code>
HP-UX	<code>MQ_INSTALLATION_PATH/samp</code>	<code>MQ_INSTALLATION_PATH/samp/bin/ah</code> (see note “2” on page 605)
Solaris	<code>MQ_INSTALLATION_PATH/samp</code>	<code>MQ_INSTALLATION_PATH/samp/bin/as</code>
Linux	<code>MQ_INSTALLATION_PATH/samp</code>	<code>MQ_INSTALLATION_PATH/samp/bin/</code>
Windows	<code>MQ_INSTALLATION_PATH\tools\cplus\samples</code>	<code>MQ_INSTALLATION_PATH\tools\cplus\samples\bin\vn</code> (see note “3” on page 605)

Notes:

1. Programs built using the ILE C++ compiler for IBM i are in the library QMQM. The include files are in `/QIBM/ProdData/mqm/inc`.
2. Programs built using the HP ANSI C++ compiler are found in directory `MQ_INSTALLATION_PATH/samp/bin/ah`. For further information, see “Building C++ programs on HP-UX” on page 619.
3. Programs built using the Microsoft Visual Studio are found in `MQ_INSTALLATION_PATH\tools\cplus\samples\bin\vn`. For further information about these compilers, see “Building C++ programs on Windows” on page 624.

Sample program HELLO WORLD (imqwrld.cpp)

This C++ sample program shows how to put and get a regular datagram (C structure) using the `ImqMessage` class.

This program shows how to put and get a regular datagram (C structure) using the `ImqMessage` class. This sample uses few method invocations, taking advantage of implicit method invocations such as **open**, **close**, and **disconnect**.

On all platforms except z/OS

If you are using a server connection to WebSphere MQ, follow one of the following procedures:

- To use the existing default queue, `SYSTEM.DEFAULT.LOCAL.QUEUE`, run the program **imqwrlds** without passing any parameters
- To use a temporary dynamically assigned queue, run **imqwrlds** passing the name of the default model queue, `SYSTEM.DEFAULT.MODEL.QUEUE`.

If you are using a client connection to WebSphere MQ, follow one of the following procedures:

- Set up the `MQSERVER` environment variable (see [MQSERVER](#) for more information) and run **imqwrldc**, or
- Run **imqwrldc** passing as parameters the **queue-name**, **queue-manager-name**, and **channel-definition**, where a typical **channel-definition** might be `SYSTEM.DEF.SVRCONN/TCP/hostname(1414)`

Sample code

```
extern "C" {
#include <stdio.h>
}

#include <imqi.hpp> // WebSphere MQ C++

#define EXISTING_QUEUE "SYSTEM.DEFAULT.LOCAL.QUEUE"

#define BUFFER_SIZE 12

static char gpszHello[ BUFFER_SIZE ] = "Hello world" ;
int main ( int argc, char * * argv ) {
    ImqQueueManager manager ;
    int iReturnCode = 0 ;

    // Connect to the queue manager.
    if ( argc > 2 ) {
        manager.setName( argv[ 2 ] );
    }
    if ( manager.connect( ) ) {
        ImqQueue * pqueue = new ImqQueue ;
        ImqMessage * pmsg = new ImqMessage ;

        // Identify the queue which will hold the message.
        pqueue -> setConnectionReference( manager );
        if ( argc > 1 ) {
            pqueue -> setName( argv[ 1 ] );

            // The named queue can be a model queue, which will result in
            // the creation of a temporary dynamic queue, which will be
            // destroyed as soon as it is closed. Therefore we must ensure
            // that such a queue is not automatically closed and reopened.
            // We do this by setting open options which will avoid the need
            // for closure and reopening.
            pqueue -> setOpenOptions( MQOO_OUTPUT | MQOO_INPUT_SHARED |
                                     MQOO_INQUIRE );
        } else {
            pqueue -> setName( EXISTING_QUEUE );

            // The existing queue is not a model queue, and will not be
            // destroyed by automatic closure and reopening. Therefore we
            // will let the open options be selected on an as-needed basis.
            // The queue will be opened implicitly with an output option
            // during the "put", and then implicitly closed and reopened
            // with the addition of an input option during the "get".
        }

        // Prepare a message containing the text "Hello world".
        pmsg -> useFullBuffer( gpszHello , BUFFER_SIZE );
        pmsg -> setFormat( MQFMT_STRING );

        // Place the message on the queue, using default put message
        // Options.
        // The queue will be automatically opened with an output option.
        if ( pqueue -> put( * pmsg ) ) {
            ImqString strQueue( pqueue -> name( ) );

            // Discover the name of the queue manager.
            ImqString strQueueManagerName( manager.name( ) );
            printf( "The queue manager name is %s.\n",
                   (char *)strQueueManagerName );

            // Show the name of the queue.
            printf( "Message sent to %s.\n", (char *)strQueue );

            // Retrieve the data message just sent ("Hello world" expected)
            // from the queue, using default get message options. The queue
            // is automatically closed and reopened with an input option
            // if it is not already open with an input option. We get the
            // message just sent, rather than any other message on the
            // queue, because the "put" will have set the ID of the message
            // so, as we are using the same message object, the message ID
            // acts as in the message object, a filter which says that we
            // are interested in a message only if it has this
            // particular ID.

            if ( pqueue -> get( * pmsg ) ) {
```

```

    int iDataLength = pmsg -> dataLength( );

    // Show the text of the received message.
    printf( "Message of length %d received, ", iDataLength );

    if ( pmsg -> formatIs( MQFMT_STRING ) ) {
        char * pszText = pmsg -> bufferPointer( );

        // If the last character of data is a null, then we can
        // assume that the data can be interpreted as a text
        // string.
        if ( ! pszText[ iDataLength - 1 ] ) {
            printf( "text is \"%s\".\n", pszText );
        } else {
            printf( "no text.\n" );
        }

    } else {
        printf( "non-text message.\n" );
    }

} else {
    printf( "ImqQueue::get failed with reason code %ld\n",
           pqueue -> reasonCode( ) );
    iReturnCode = (int)pqueue -> reasonCode( );
}

} else {
    printf( "ImqQueue::open/put failed with reason code %ld\n",
           pqueue -> reasonCode( ) );
    iReturnCode = (int)pqueue -> reasonCode( );
}

// Deletion of the queue will ensure that it is closed.
// If the queue is dynamic then it will also be destroyed.
delete pqueue ;
delete pmsg ;

} else {
    printf( "ImqQueueManager::connect failed with reason code %ld\n",
           manager.reasonCode( ) );
    iReturnCode = (int)manager.reasonCode( );
}

// Destruction of the queue manager ensures that it is
// disconnected. If the queue object were still available
// and open (which it is not), the queue would be closed
// prior to disconnection.

return iReturnCode ;
}

```

Sample programs SPUT (imqspu.cpp) and SGET (imqsget.cpp)

These C++ programs place messages to, and retrieve messages from, a named queue.

These samples show the use of the following classes:

- ImqError (see [ImqError C++ class](#))
- ImqMessage (see [ImqMessage C++ class](#))
- ImqObject (see [ImqObject C++ class](#))
- ImqQueue (see [ImqQueue C++ class](#))
- ImqQueueManager (see [ImqQueueManager C++ class](#))

Follow the appropriate instructions to run the programs.

On all platforms except z/OS

1. Run **imqspu** *queue-name*.
2. Type lines of text at the console. These lines are placed as messages onto the specified queue.
3. Enter a null line to end the input.

4. Run **imqsgets** *queue-name* to retrieve all the lines and display them at the console.

Sample program DPUT (imqdput.cpp)

This C++ sample program puts messages to a distribution list consisting of two queues.

DPUT shows the use of the `ImqDistributionList` class (see [ImqDistributionList C++ class](#)). This sample is not supported on z/OS.

1. Run **imqdputs** *queue-name-1 queue-name-2* to place messages on the two named queues.
2. Run **imqsgets** *queue-name-1* and **imqsgets** *queue-name-2* to retrieve the messages from those queues.

C++ language considerations

This collection of topics details the aspects of the C++ language usage and conventions that you must consider when writing application programs that use the Message Queue Interface (MQI).

C++ Header files

Header files are provided as part of the definition of the MQI, to help you write WebSphere MQ application programs in the C++ language.

These header files are summarized in the following table.

Table 81. C/C++ header files	
Filename	Contents
IMQI.HPP	C++ MQI Classes (includes CMQC.H and IMQTYPE.H)
IMQTYPE.H	Defines the ImqBoolean data type
CMQC.H	MQI data structures and manifest constants

To improve the portability of applications, code the name of the header file in lowercase on the **#include** preprocessor directive:

```
#include <imqi.hpp> // C++ classes
```

C++ methods and attributes

Method names are in mixed case. Various considerations apply to parameters and return values. Attributes are accessed using set and get methods as appropriate.

Parameters of methods that are *const* are for input only. Parameters with signatures including a pointer (*) or a reference (&) are passed by reference. Return values that do not include a pointer or a reference are passed by value; in the case of returned objects, these are new entities that become the responsibility of the caller.

Some method signatures include items that take a default if not specified. Such items are always at the end of signatures and are denoted by an equal sign (=); the value after the equal sign indicates the default value that applies if the item is omitted.

All method names in these classes are mixed case, beginning with lowercase. Each word, except the first within a method name, begins with a capital letter. Abbreviations are not used unless their meaning is widely understood. Abbreviations used include *id* (for identity) and *sync* (for synchronization).

Object attributes are accessed using set and get methods. A set method begins with the word *set*; a get method has no prefix. If an attribute is *read-only*, there is no set method.

Attributes are initialized to valid states during object construction, and the state of an object is always consistent.

Data types in C++

All data types are defined by the C **typedef** statement.

The type **ImqBoolean** is defined as **unsigned char** in IMQTYPE.H and can have the values TRUE and FALSE. You can use **ImqBinary** class objects in place of **MQBYTE** arrays, and **ImqString** class objects in place of **char ***. Many methods return objects instead of **char** or **MQBYTE** pointers to ease storage management. All return values become the responsibility of the caller, and, in the case of a returned object, the storage can be disposed of using delete.

Manipulating binary strings in C++

Strings of binary data are declared as objects of the **ImqBinary** class. Objects of this class can be copied, compared, and set using the familiar C operators. Example code is provided.

The following code sample shows operations on a binary string:

```
#include <imqi.hpp> // C++ classes

ImqMessage message ;
ImqBinary id, correlationId ;
MQBYTE24 byteId ;

correlationId.set( byteId, sizeof( byteId ) ); // Set.
id = message.id( );                          // Assign.
if ( correlationId == id ) {                  // Compare.
    ...
}
```

Manipulating character strings in C++

Character data is often returned in **ImqString** class objects which can be cast to **char *** using a conversion operator. The **ImqString** class contains methods to assist in the processing of character strings.

When character data is accepted or returned using MQI C++ methods, the character data is always null-terminated and can be of any length. However, certain limits are imposed by WebSphere MQ that might result in information being truncated. To ease storage management, character data is often returned in **ImqString** class objects. These objects can be cast to **char *** using the conversion operator provided, and used for *read-only* purposes in many situations where a **char *** is required.

Note: The **char *** conversion result from an **ImqString** class object might be null.

Although C functions can be used on the **char ***, there are special methods of the **ImqString** class that are preferable; **operator length()** is the equivalent of **strlen** and **storage()** indicates the memory allocated for the character data.

Initial state of objects in C++

All objects have a consistent initial state reflected by their attributes. The initial values are defined in the class descriptions.

Using C from C++

When you use C functions from a C++ program, include appropriate headers.

The following example shows **string.h** included in a C++ program:

```
extern "C" {
#include <string.h>
}
```

C++ notational conventions

This example shows how to invoke methods and declare parameters.

This code sample uses the methods and parameters **ImqBoolean ImqQueue::get(ImqMessage & msg)**

Declare and use the parameters as follows:

```
ImqQueueManager * pmanager ;    // Queue manager
ImqQueue * pqueue ;            // Message queue
ImqMessage msg ;               // Message
char szBuffer[ 100 ] ;         // Buffer for message data

pmanager = new ImqQueueManager ;
pqueue = new ImqQueue ;
pqueue -> setName( "myreplyq" );
pqueue -> setConnectionReference( pmanager );

msg.useEmptyBuffer( szBuffer, sizeof( szBuffer ) );

if ( pqueue -> get( msg ) ) {
    long lDataLength = msg.dataLength( );
} ...
```

Implicit operations in C++

Several operations can occur implicitly, *just in time*, to satisfy the prerequisite conditions for the successful execution of a method. These implicit operations are connect, open, reopen, close, and disconnect. You can control connect and open implicit behavior using class attributes.

Connect

An ImqQueueManager object is connected automatically for any method that results in any call to the MQI (see [C++ and MQI cross-reference](#)).

Open

An ImqObject object is opened automatically for any method that results in an MQGET, MQINQ, MQPUT, or MQSET call. Use the **openFor** method to specify one or more relevant **open option** values.

Reopen

An ImqObject is reopened automatically for any method that results in an MQGET, MQINQ, MQPUT, or MQSET call, where the object is already open, but the existing **open options** are not adequate to allow the MQI call to be successful. The object is temporarily closed using a temporary **close options** value of MQCO_NONE. Use the **openFor** method to add a relevant **open option**.

Reopen can cause problems in specific circumstances:

- A temporary dynamic queue is destroyed when it is closed and can never be reopened.
- A queue opened for exclusive input (either explicitly or by default) might be accessed by others in the window of opportunity during closure and reopening.
- A browse cursor position is lost when a queue is closed. This situation does not prevent closure and reopening, but prevents subsequent use of the cursor until MQGMO_BROWSE_FIRST is used again.
- The context of the last message retrieved is lost when a queue is closed.

If any of these circumstances occur or can be foreseen, avoid reopens by explicitly setting adequate **open options** before an object is opened (either explicitly or implicitly).

Setting the **open options** explicitly for complex queue-handling situations results in better performance and avoids the problems associated with the use of reopen.

Close

An ImqObject is closed automatically at any point where the object state is no longer viable, for example if an ImqObject connection reference is severed, or if an ImqObject object is destroyed.

Disconnect

An ImqQueueManager is disconnected automatically at any point where the connection is no longer viable, for example if an ImqObject connection reference is severed, or if an ImqQueueManager object is destroyed.

Binary and character strings in C++

The ImqString class encapsulates the traditional *char ** data format. The ImqBinary class encapsulates the binary byte array. Some methods that set character data might truncate the data.

Methods that set character (**char ***) data always take a copy of the data, but some methods might truncate the copy, because certain limits are imposed by WebSphere MQ.

The ImqString class (see [ImqString C++ class](#)) encapsulates the traditional **char *** and provides support for:

- Comparison
- Concatenation
- Copying
- Integer-to-text and text-to-integer conversion
- Token (word) extraction
- Uppercase translation

The ImqBinary class (see [ImqBinary C++ class](#)) encapsulates binary byte arrays of arbitrary size. In particular, it is used to hold the following attributes:

- **accounting token** (MQBYTE32)
- **connection tag** (MQBYTE128)
- **correlation id** (MQBYTE24)
- **facility token** (MQBYTE8)
- **group id** (MQBYTE24)
- **instance id** (MQBYTE24)
- **message id** (MQBYTE24)
- **message token** (MQBYTE16)
- **transaction instance id** (MQBYTE16)

Where these attributes belong to objects of the following classes:

- ImqCICSBridgeHeader (see [ImqCICSBridgeHeader C++ class](#))
- ImqGetMessageOptions (see [ImqGetMessageOptions C++ class](#))
- ImqIMSBridgeHeader (see [ImqIMSBridgeHeader C++ class](#))
- ImqMessageTracker (see [ImqMessageTracker C++ class](#))
- ImqQueueManager (see [ImqQueueManager C++ class](#))
- ImqReferenceHeader (see [ImqReferenceHeader C++ class](#))
- ImqWorkHeader (see [ImqWorkHeader C++ class](#))

The ImqBinary class also provides support for comparison and copying.

Unsupported functions in C++

The WebSphere MQ C++ classes and methods are independent of WebSphere MQ platform. They might therefore offer some functions that are not supported on certain platforms.

If you try to use a function on a platform on which it is not supported, the function is detected by WebSphere MQ but not by the C++ language bindings. WebSphere MQ reports the error to your program, like any other MQI error.

Messaging in C++

This collection of topics details how to prepare, read, and write messages in C++.

Preparing message data in C++

Message data is prepared in a buffer, which can be supplied by the system or the application. There are advantages to either method. Examples of using a buffer are given.

When you send a message, message data is first prepared in a buffer managed by an `ImqCache` object (see `ImqCache C++ class`). A buffer is associated (by inheritance) with each `ImqMessage` object (see `ImqMessage C++ class`): it can be supplied by the application (using either the **`useEmptyBuffer`** or **`useFullBuffer`** method) or automatically by the system. The advantage of the application supplying the message buffer is that no data copying is necessary in many cases because the application can use prepared data areas directly. The disadvantage is that the supplied buffer is of a fixed length.

The buffer can be reused, and the number of bytes transmitted can be varied each time, by using the **`setMessageLength`** method before transmission.

When supplied automatically by the system, the number of bytes available is managed by the system, and data can be copied into the message buffer using, for example, the `ImqCache` **`write`** method, or the `ImqMessage` **`writeItem`** method. The message buffer grows according to need. As the buffer grows, there is no loss of previously written data. A large or multipart message can be written in sequential pieces.

The following examples show simplified message sends.

1. Use prepared data in a user-supplied buffer

```
char szBuffer[ ] = "Hello world" ;  
  
msg.useFullBuffer( szBuffer, sizeof( szBuffer ) );  
msg.setFormat( MQFMT_STRING );
```

2. Use prepared data in a user-supplied buffer, where the buffer size exceeds the data size

```
char szBuffer[ 24 ] = "Hello world" ;  
  
msg.useEmptyBuffer( szBuffer, sizeof( szBuffer ) );  
msg.setFormat( MQFMT_STRING );  
msg.setMessageLength( 12 );
```

3. Copy data to a user-supplied buffer

```
char szBuffer[ 12 ];  
  
msg.useEmptyBuffer( szBuffer, sizeof( szBuffer ) );  
msg.setFormat( MQFMT_STRING );  
msg.write( 12, "Hello world" );
```

4. Copy data to a system-supplied buffer

```
msg.setFormat( MQFMT_STRING );
msg.write( 12, "Hello world" );
```

5. Copy data to a system-supplied buffer using objects (objects set the message format as well as content)

```
ImqString strText( "Hello world" );
msg.writeItem( strText );
```

Reading messages in C++

A buffer can be supplied by the application or the system. Data can be accessed directly from the buffer or read sequentially. There is a class equivalent to each message type. Sample code is given.

When receiving data, the application or the system can supply a suitable message buffer. The same buffer can be used for both multiple transmission and multiple receipt for a particular `ImqMessage` object. If the message buffer is supplied automatically, it grows to accommodate whatever length of data is received. However, a message buffer supplied by the application might not be big enough to hold the data received. Then either truncation or failure might occur, depending on the options used for message receipt.

Incoming data can be accessed directly from the message buffer, in which case the data length indicates the total amount of incoming data. Alternatively, incoming data can be read sequentially from the message buffer. In this case, the data pointer addresses the next byte of incoming data, and the data pointer and data length are updated each time data is read.

Items are pieces of a message, all in the user area of the message buffer, that need to be processed sequentially and separately. Apart from regular user data, an item might be a dead-letter header or a trigger message. Items are always associated with message formats; message formats are **not** always associated with items.

There is a class of object for each item that corresponds to a recognizable WebSphere MQ message format. There is one for a dead-letter header and one for a trigger message. There is no class of object for user data. That is, once the recognizable formats have been exhausted, processing the remainder is left to the application program. Classes for user data can be written by specializing the `ImqItem` class.

The following example shows a message receipt that takes account of a number of potential items that can precede the user data, in an imaginary situation. Non-item user data is defined as anything that occurs after items that can be identified. An automatic buffer (the default) is used to hold an arbitrary amount of message data.

```
ImqQueue queue ;
ImqMessage msg ;

if ( queue.get( msg ) ) {
    /* Process all items of data in the message buffer. */
    do while ( msg.dataLength( ) ) {
        ImqBoolean bFormatKnown = FALSE ;
        /* There remains unprocessed data in the message buffer. */

        /* Determine what kind of item is next. */

        if ( msg.formatIs( MQFMT_DEAD_LETTER_HEADER ) ) {
            ImqDeadLetterHeader header ;
            /* The next item is a dead-letter header. */
            /* For the next statement to work and return TRUE, */
            /* the correct class of object pointer must be supplied. */
            bFormatKnown = TRUE ;

            if ( msg.readItem( header ) ) {
                /* The dead-letter header has been extricated from the */
                /* buffer and transformed into a dead-letter object. */
            }
        }
    }
}
```

```

        /* The encoding and character set of the dead-letter object
        /* itself are MQENC_NATIVE and MQCCSI_Q_MGR.
        /* The encoding and character set from the dead-letter
        /* header have been copied to the message attributes
        /* to reflect any remaining data in the buffer.
        */

        /* Process the information in the dead-letter object.
        /* Note that the encoding and character set have
        /* already been processed.
        */
        ...
    }
    /* There might be another item after this,
    /* or just the user data.
    */
}
if ( msg.formatIs( MQFMT_TRIGGER ) ) {
    ImqTrigger trigger ;
    /* The next item is a trigger message.
    /* For the next statement to work and return TRUE,
    /* the correct class of object pointer must be supplied.
    bFormatKnown = TRUE ;
    if ( msg.readItem( trigger ) ) {

        /* The trigger message has been extricated from the
        /* buffer and transformed into a trigger object.
        /* Process the information in the trigger object.
        */
        ...
    }

    /* There is usually nothing after a trigger message.
    */
}

if ( msg.formatIs( FMT_USERCLASS ) ) {
    UClass object ;
    /* The next item is an item of a user-defined class.
    /* For the next statement to work and return TRUE,
    /* the correct class of object pointer must be supplied.
    bFormatKnown = TRUE ;

    if ( msg.readItem( object ) ) {
        /* The user-defined data has been extricated from the
        /* buffer and transformed into a user-defined object.
        */

        /* Process the information in the user-defined object.
        */
        ...
    }

    /* Continue looking for further items.
    */
}
if ( ! bFormatKnown ) {
    /* There remains data that is not associated with a specific
    /* item class.
    char * pszDataPointer = msg.dataPointer( );
    int iDataLength = msg.dataLength( );

    /* The encoding and character set for the remaining data are
    /* reflected in the attributes of the message object, even
    /* if a dead-letter header was present.
    ...
}
}
}

```

In this example, FMT_USERCLASS is a constant representing the 8-character format name associated with an object of class UClass, and is defined by the application.

UClass is derived from the ImqItem class (see [ImqItem C++ class](#)), and implements the virtual **copyOut** and **pasteIn** methods from that class.

The next two examples show code from the ImqDeadLetterHeader class (see [ImqDeadLetterHeader C++ class](#)). The first example shows custom-encapsulated message-writing code.

```

// Insert a dead-letter header.
// Return TRUE if successful.
ImqBoolean ImqDeadLetterHeader :: copyOut ( ImqMessage & msg ) {
    ImqBoolean bSuccess ;

```

```

if ( msg.moreBytes( sizeof( omqdlh ) ) ) {
    ImqCache cacheData( msg ); // Preserve original message content.
    // Note original message attributes in the dead-letter header.
    setEncoding( msg.encoding( ) );
    setCharacterSet( msg.characterSet( ) );
    setFormat( msg.format( ) );

    // Set the message attributes to reflect the dead-letter header.
    msg.setEncoding( MQENC_NATIVE );
    msg.setCharacterSet( MQCCSI_Q_MGR );
    msg.setFormat( MQFMT_DEAD_LETTER_HEADER );
    // Replace the existing data with the dead-letter header.
    msg.clearMessage( );
    if ( msg.write( sizeof( omqdlh ), (char *) & omqdlh ) ) {
        // Append the original message data.
        bSuccess = msg.write( cacheData.messageLength( ),
                               cacheData.bufferPointer( ) );
    } else {
        bSuccess = FALSE ;
    }
} else {
    bSuccess = FALSE ;
}
// Reflect and cache error in this object.
if ( ! bSuccess ) {
    setReasonCode( msg.reasonCode( ) );
    setCompletionCode( msg.completionCode( ) );
}

return bSuccess ;
}

```

The second example shows custom-encapsulated message-reading code.

```

// Read a dead-letter header.
// Return TRUE if successful.
ImqBoolean ImqDeadLetterHeader :: pasteIn ( ImqMessage & msg ) {
    ImqBoolean bSuccess = FALSE ;

    // First check that the eye-catcher is correct.
    // This is also our guarantee that the "character set" is correct.
    if ( ImqItem::structureIdIs( MQDLH_STRUC_ID, msg ) ) {
        // Next check that the "encoding" is correct, as the MQDLH
        // contains numeric data.
        if ( msg.encoding( ) == MQENC_NATIVE ) {

            // Finally check that the "format" is correct.
            if ( msg.formatIs( MQFMT_DEAD_LETTER_HEADER ) ) {
                char * pszBuffer = (char *) & omqdlh ;
                // Transfer the MQDLH from the message and move pointer on.
                if ( bSuccess = msg.read( sizeof( omdlh ), pszBuffer ) ) {
                    // Update the encoding, character set and format of the
                    // message to reflect the remaining data.
                    msg.setEncoding( encoding( ) );
                    msg.setCharacterSet( characterSet( ) );
                    msg.setFormat( format( ) );
                } else {

                    // Reflect the cache error in this object.
                    setReasonCode( msg.reasonCode( ) );
                    setCompletionCode( msg.completionCode( ) );
                }
            } else {
                setReasonCode( MQRC_INCONSISTENT_FORMAT );
                setCompletionCode( MQCC_FAILED );
            }
        } else {
            setReasonCode( MQRC_ENCODING_ERROR );
            setCompletionCode( MQCC_FAILED );
        }
    } else {
        setReasonCode( MQRC_STRUC_ID_ERROR );
        setCompletionCode( MQCC_FAILED );
    }
}

return bSuccess ;
}

```

With an automatic buffer, the buffer storage is *volatile*. That is, buffer data might be held at a different physical location after each **get** method invocation. Therefore, each time buffer data is referenced, use the **bufferPointer** or **dataPointer** methods to access message data.

You might want a program to set aside a fixed area for receiving message data. In this case, invoke the **useEmptyBuffer** method before using the **get** method.

Using a fixed, nonautomatic area limits messages to a maximum size, so it is important to consider the MQGMO_ACCEPT_TRUNCATED_MSG option of the ImqGetMessageOptions object. If this option is not specified (the default), the MQRC_TRUNCATED_MSG_FAILED reason code can be expected. If this option is specified, the MQRC_TRUNCATED_MSG_ACCEPTED reason code might be expected depending on the design of the application.

The next example shows how a fixed area of storage can be used to receive messages:

```
char * pszBuffer = new char[ 100 ];

msg.useEmptyBuffer( pszBuffer, 100 );
gmo.setOptions( MQGMO_ACCEPT_TRUNCATED_MSG );
queue.get( msg, gmo );

delete [ ] pszBuffer ;
```

In this code fragment, the buffer can always be addressed directly, with *pszBuffer*, as opposed to using the **bufferPointer** method. However, it is better to use the **dataPointer** method for general-purpose access. The application (not the ImqCache class object) must discard a user-defined (nonautomatic) buffer.

Attention: Specifying a null pointer and zero length with **useEmptyBuffer** does not nominate a fixed-length buffer of length zero as might be expected. This combination is interpreted as a request to ignore any previous user-defined buffer, and instead revert to the use of an automatic buffer.

Writing a message to the dead-letter queue in C++

Example program code for writing a message to the dead-letter queue.

A typical case of a multipart message is one containing a dead-letter header. The data from a message that cannot be processed is appended to the dead-letter header.

```
ImqQueueManager mgr ;           // The queue manager.
ImqQueue queueIn ;             // Incoming message queue.
ImqQueue queueDead ;           // Dead-letter message queue.
ImqMessage msg ;               // Incoming and outgoing message.
ImqDeadLetterHeader header ;    // Dead-letter header information.

// Retrieve the message to be rerouted.
queueIn.setConnectionReference( mgr );
queueIn.setName( MY_QUEUE );
queueIn.get( msg );

// Set up the dead-letter header information.
header.setDestinationQueueManagerName( mgr.name( ) );
header.setDestinationQueueName( queueIn.name( ) );
header.setPutApplicationName( /* ? */ );
header.setPutApplicationType( /* ? */ );
header.setPutDate( /* TODAY */ );
header.setPutTime( /* NOW */ );
header.setDeadLetterReasonCode( FB_APPL_ERROR_1234 );

// Insert the dead-letter header information. This will vary
// the encoding, character set and format of the message.
// Message data is moved along, past the header.
msg.writeItem( header );

// Send the message to the dead-letter queue.
queueDead.setConnectionReference( mgr );
queueDead.setName( mgr.deadLetterQueueName( ) );
queueDead.put( msg );
```


Writing a message to the IMS bridge in C++

Example program code for writing a message to the IMS bridge.

Messages sent to the WebSphere MQ-IMS bridge might use a special header. The IMS bridge header is prefixed to regular message data.

```
ImqQueueManager mgr;           // The queue manager.
ImqQueue queueBridge;         // IMS bridge message queue.
ImqMessage msg;               // Outgoing message.
ImqIMSBridgeHeader header;     // IMS bridge header.

// Set up the message.
//
// Here we are constructing a message with format
// MQFMT_IMS_VAR_STRING, and appropriate data.
//
msg.write( 2, /* ? */ );      // Total message length.
msg.write( 2, /* ? */ );      // IMS flags.
msg.write( 7, /* ? */ );      // Transaction code.
msg.write( /* ? */ , /* ? */ ); // String data.
msg.setFormat( MQFMT_IMS_VAR_STRING ); // The format attribute.

// Set up the IMS bridge header information.
//
// The reply-to-format is often specified.
// Other attributes can be specified, but all have default values.
//
header.setReplyToFormat( /* ? */ );

// Insert the IMS bridge header into the message.
//
// This will:
// 1) Insert the header into the message buffer, before the existing
//    data.
// 2) Copy attributes out of the message descriptor into the header,
//    for example the IMS bridge header format attribute will now
//    be set to MQFMT_IMS_VAR_STRING.
// 3) Set up the message attributes to describe the header, in
//    particular setting the message format to MQFMT_IMS.
//
msg.writeItem( header );

// Send the message to the IMS bridge queue.
//
queueBridge.setConnectionReference( mgr );
queueBridge.setName( /* ? */ );
queueBridge.put( msg );
```

Writing a message to the CICS bridge in C++

Example program code for writing a message to the CICS bridge.

Messages sent to WebSphere MQ for z/OS using the CICS bridge require a special header. The CICS bridge header is prefixed to regular message data.

```
ImqQueueManager mgr ;           // The queue manager.
ImqQueue queueIn ;              // Incoming message queue.
ImqQueue queueBridge ;          // CICS bridge message queue.
ImqMessage msg ;                // Incoming and outgoing message.
ImqCicsBridgeHeader header ;    // CICS bridge header information.

// Retrieve the message to be forwarded.
queueIn.setConnectionReference( mgr );
queueIn.setName( MY_QUEUE );
queueIn.get( msg );

// Set up the CICS bridge header information.
// The reply-to format is often specified.
// Other attributes can be specified, but all have default values.
header.setReplyToFormat( /* ? */ );

// Insert the CICS bridge header information. This will vary
// the encoding, character set and format of the message.
```

```
// Message data is moved along, past the header.
msg.writeItem( header );

// Send the message to the CICS bridge queue.
queueBridge.setConnectionReference( mgr );
queueBridge.setName( /* ? */ );
queueBridge.put( msg );
```

Writing a message with a work header in C++

Example program code for writing a message destined for a queue managed by the z/OS Workload Manager.

Messages sent to WebSphere MQ for z/OS, which are destined for a queue managed by the z/OS Workload Manager, require a special header. The work header is prefixed to regular message data.

```
ImqQueueManager mgr ;           // The queue manager.
ImqQueue queueIn ;             // Incoming message queue.
ImqQueue queueWLM ;           // WLM managed queue.
ImqMessage msg ;              // Incoming and outgoing message.
ImqWorkHeader header ;         // Work header information

// Retrieve the message to be forwarded.
queueIn.setConnectionReference( mgr );
queueIn.setName( MY_QUEUE );
queueIn.get( msg );

// Insert the Work header information. This will vary
// the encoding, character set and format of the message.
// Message data is moved along, past the header.
msg.writeItem( header );

// Send the message to the WLM managed queue.
queueWLM.setConnectionReference( mgr );
queueWLM.setName( /* ? */ );
queueWLM.put( msg );
```

Building WebSphere MQ C++ programs

The URL of supported compilers is listed, together with the commands to use to compile, link and run C++ programs and samples on WebSphere MQ platforms.

The compilers for each supported platform and version of WebSphere MQ are listed at the WebSphere MQ system requirements page at [System Requirements for IBM WebSphere MQ](#).

The command you need to compile and link your WebSphere MQ C++ program depends on your installation and requirements. The examples that follow show typical compile and link commands for some of the compilers using the default installation of WebSphere MQ on a number of platforms.

Building C++ programs on AIX

Build WebSphere MQ C++ programs on AIX using the XL C Enterprise Edition compiler.

Client

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

32-bit unthreaded application

```
xlc -o imqsputc_32 imqspcut.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -limqc23ia -limqb23ia -lmqic
```

32-bit threaded application

```
xlc_r -o imqsputc_32_r imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -limqc23ia_r -limqb23ia_r -lmqic_r
```

64-bit unthreaded application

```
xlc -q64 -o imqsputc_64 imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -limqc23ia -limqb23ia -lmqic
```

64-bit threaded application

```
xlc_r -q64 -o imqsputc_64_r imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -limqc23ia_r -limqb23ia_r -lmqic_r
```

Server

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

32-bit unthreaded application

```
xlc -o imqsputc_32 imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -limqs23ia -limqb23ia -lmqm
```

32-bit threaded application

```
xlc_r -o imqsputc_32_r imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -limqs23ia_r -limqb23ia_r -lmqm_r
```

64-bit unthreaded application

```
xlc -q64 -o imqsputc_64 imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -limqs23ia -limqb23ia -lmqm
```

64-bit threaded application

```
xlc_r -q64 -o imqsputc_64_r imqsputc.cpp -qchars=signed -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -limqs23ia_r -limqb23ia_r -lmqm_r
```

Building C++ programs on HP-UX

Build WebSphere MQ C++ programs on HP-UX using the aC++ or aCC compilers.

On HP-UX Itanium, WebSphere MQ supports only the Standard runtime. Use the aCC compiler.

- libimqi23bh.sl provides the WebSphere MQ C++ classes for the Standard runtime.
- For compatibility with earlier releases, a symbolic link is provided from libimqi23ah.sl to libimqi23bh.sl.

IA64 (IPF)

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Client: IA64 (IPF)

32-bit unthreaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsputc_32 imqsputc.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh -lmqic
```

32-bit threaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsputc_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh_r -lmqic_r -lpthread
```

64-bit unthreaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsputc_64 imqsput.cpp  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh -lmqic
```

64-bit threaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsputc_64_r imqsput.cpp  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh_r  
-lmqic_r  
-lpthread
```

Server: IA64 (IPF)

32-bit unthreaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsput_32 imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh -lmqm
```

32-bit threaded application

```
aCC -Wl,+b,: +e -D_HPUX_SOURCE -o imqsput_32_r imqsput.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -L/usr/lib/hpux32 -limqi23bh_r -lmqm_r -lpthread
```

64-bit unthreaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsput_64 imqsput.cpp  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh -lmqm
```

64-bit threaded application

```
aCC +DD64 +e -D_HPUX_SOURCE -o imqsput_64_r imqsput.cpp  
-IMQ_INSTALLATION_PATH/inc -LMQ_INSTALLATION_PATH/lib64 -L/usr/lib/hpux64 -limqi23bh_r  
-lmqm_r  
-lpthread
```

Building C++ programs on Linux

Build WebSphere MQ C++ programs on Linux using the GNU g++ compiler.

System p

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Client: System p

32-bit unthreaded application

```
g++ -m32 -o imqsputc_32 imqsput.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl  
-limqb23gl -lmqic
```

32-bit threaded application

```
g++ -m32 -o imqsputc_r32 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl_r  
-limqb23gl_r -lmqic_r
```

64-bit unthreaded application

```
g++ -m64 -o imqsputc_64 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64  
-limqc23gl -limqb23gl -lmqic
```

64-bit threaded application

```
g++ -m64 -o imqsputc_r64 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64  
-limqc23gl_r -limqb23gl_r -lmqic_r
```

Server: System p

32-bit unthreaded application

```
g++ -m32 -o imqspcut_32 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl  
-limqb23gl -lmqm
```

32-bit threaded application

```
g++ -m32 -o imqspcut_r32 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqs23gl_r  
-limqb23gl_r -lmqm_r
```

64-bit unthreaded application

```
g++ -m64 -o imqspcut_64 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64  
-limqs23gl -limqb23gl -lmqm
```

64-bit threaded application

```
g++ -m64 -o imqspcut_r64 imqspcut.cpp -fsigned-char -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64  
-limqs23gl_r -limqb23gl_r -lmqm_r
```

System z

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Client: System z

32-bit unthreaded application

```
g++ -m31 -fsigned-char -o imqspcut_32 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib  
-limqc23gl -limqb23gl -lmqic
```

32-bit threaded application

```
g++ -m31 -fsigned-char -o imqspu32_r imqspu32.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-lmqc23gl_r -limqb23gl_r -lmqic_r
-lpthread
```

64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqspu64 imqspu64.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl -limqb23gl -lmqic
```

64-bit threaded application

```
g++ -m64 -fsigned-char -o imqspu64_r imqspu64.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl_r -limqb23gl_r -lmqic_r -lpthread
```

Server: System z

32-bit unthreaded application

```
g++ -m31 -fsigned-char -o imqspu32 imqspu32.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-lmqc23gl -limqb23gl -lmqm
```

32-bit threaded application

```
g++ -m31 -fsigned-char -o imqspu32_r imqspu32.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib
-lmqc23gl_r -limqb23gl_r -lmqm_r -lpthread
```

64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqspu64 imqspu64.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl -limqb23gl -lmqm
```

64-bit threaded application

```
g++ -m64 -fsigned-char -o imqspu64_r imqspu64.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64
-lmqc23gl_r -limqb23gl_r -lmqm_r -lpthread
```

System x (32-bit)

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Client: System x (32-bit)

32-bit unthreaded application

```
g++ -m32 -fsigned-char -o imqspu32 imqspu32.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -LMQ_INSTALLATION_PATH/lib
-Wl,
-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl -limqb23gl -lmqic
```

32-bit threaded application

```
g++ -m32 -fsigned-char -o imqspu32_r imqspu32.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -LMQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqc23gl_r -limqb23gl_r
-lmqic_r -lpthread
```

64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqsputc_64 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -LMQ_INSTALLATION_PATH/
lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqc23gl -limqb23gl
-lmqic
```

64-bit threaded application

```
g++ -m64 -fsigned-char -o imqsputc_64_r imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -LMQ_INSTALLATION_PATH/
lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqc23gl_r -limqb23gl_r
-lmqic_r -lpthread
```

Server: System x (32-bit)

32-bit unthreaded application

```
g++ -m32 -fsigned-char -o imqspcut_32 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -LMQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl -limqb23gl -lmqm
```

32-bit threaded application

```
g++ -m32 -fsigned-char -o imqspcut_32_r imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -Wl,-rpath=MQ_INSTALLATION_PATH/lib -LMQ_INSTALLATION_PATH/lib
-Wl,-rpath=MQ_INSTALLATION_PATH/lib -Wl,-rpath=/usr/lib -limqs23gl_r -limqb23gl_r
-lmqm_r -lpthread
```

64-bit unthreaded application

```
g++ -m64 -fsigned-char -o imqspcut_64 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -LMQ_INSTALLATION_PATH/
lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqs23gl -limqb23gl -lmqm
```

64-bit threaded application

```
g++ -m64 -fsigned-char -o imqspcut_64_r imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -LMQ_INSTALLATION_PATH/
lib64
-Wl,-rpath=MQ_INSTALLATION_PATH/lib64 -Wl,-rpath=/usr/lib64 -limqs23gl_r -limqb23gl_r
-lmqm_r -lpthread
```

Building C++ programs on Solaris

Build WebSphere MQ C++ programs on Solaris using the Sun ONE compiler.

SPARC

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Client: SPARC

32-bit application

```
CC -xarch=v8plus -mt -o imqspcut_32 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as
-lmqic -lsocket -lnsl -ldl
```

64-bit application

```
CC -xarch=v9 -mt -o imqspcut_64 imqspcut.cpp -IMQ_INSTALLATION_PATH/inc
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as
```

```
-limqb23as  
-lmqic -lsocket -lnsl -ldl
```

Server: SPARC

32-bit application

```
CC -xarch=v8plus -mt -o imqspu32 imqspu32.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as  
-lmqm -lsocket -lnsl -ldl
```

64-bit application

```
CC -xarch=v9 -mt -o imqspu64 imqspu64.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as  
-limqb23as  
-lmqm -lsocket -lnsl -ldl
```

x86-64

MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.

Client: x86-64

32-bit application

```
CC -xarch=386 -mt -o imqspu32c imqspu32c.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqc23as -limqb23as  
-lmqic -lsocket -lnsl -ldl
```

64-bit application

```
CC -xarch=amd64 -mt -o imqspu64c imqspu64c.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqc23as  
-limqb23as  
-lmqic -lsocket -lnsl -ldl
```

Server: x86-64

32-bit application

```
CC -xarch=386 -mt -o imqspu32 imqspu32.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib -RMQ_INSTALLATION_PATH/lib -R/usr/lib/32 -limqs23as -limqb23as  
-lmqm -lsocket -lnsl -ldl
```

64-bit application

```
CC -xarch=amd64 -mt -o imqspu64 imqspu64.cpp -IMQ_INSTALLATION_PATH/inc  
-LMQ_INSTALLATION_PATH/lib64 -RMQ_INSTALLATION_PATH/lib64 -R/usr/lib/64 -limqs23as  
-limqb23as  
-lmqm -lsocket -lnsl -ldl
```

Building C++ programs on Windows

Build WebSphere MQ C++ programs on Windows using the Microsoft Visual Studio C++ compiler.

Library (.lib) files and dll files for use with 32-bit applications are installed in *MQ_INSTALLATION_PATH/Tools/Lib*, files for use with 64-bit applications are installed in *MQ_INSTALLATION_PATH/Tools/Lib64*. *MQ_INSTALLATION_PATH* represents the high-level directory in which WebSphere MQ is installed.

Client

```
cl -MD imqspu.c /Feimqspu.c imqb23vn.lib imqc23vn.lib
```

Server

```
cl -MD imqspu.c /Feimqspu.c imqb23vn.lib imqc23vn.lib
```

Using WebSphere MQ classes for Java

WebSphere MQ classes for Java enable you to use WebSphere MQ in a Java environment. A Java application can use either WebSphere MQ classes for Java or WebSphere MQ classes for JMS to access WebSphere MQ resources.

WebSphere MQ classes for Java allow a Java application to:

- Connect to WebSphere MQ as a WebSphere MQ client
- Connect directly to a WebSphere MQ queue manager

WebSphere MQ classes for Java encapsulate the Message Queue Interface (MQI), the native WebSphere MQ API.

WebSphere MQ classes for Java use a similar object model to the C++ and .NET interfaces to WebSphere MQ.

Why should I use WebSphere MQ classes for Java?

If the following points are significant in your installation, consider using WebSphere MQ classes for Java:

- WebSphere MQ classes for Java encapsulate the Message Queue Interface (MQI), the native WebSphere MQ API.
 - If you are familiar with the use of the MQI in procedural languages, you can transfer this knowledge to the Java environment.
 - You can exploit the full range of features of WebSphere MQ, beyond those available through JMS.
- WebSphere MQ classes for Java use a similar object model to the C++ and .NET interfaces to WebSphere MQ. If you are familiar with these interfaces, you can transfer this knowledge to the Java environment.

Note: Automatic client reconnection is not supported by WebSphere MQ classes for Java .

Getting started with WebSphere MQ classes for Java

This collection of topics gives an overview of WebSphere MQ classes for Java and their uses.

What are WebSphere MQ classes for Java?

WebSphere MQ classes for Java allow you to use WebSphere MQ in a Java environment.

WebSphere MQ classes for Java allow a Java application to:

- Connect to WebSphere MQ as a WebSphere MQ client
- Connect directly to a WebSphere MQ queue manager

WebSphere MQ classes for Java encapsulate the Message Queue Interface (MQI), the native WebSphere MQ API.

WebSphere MQ classes for Java use a similar object model to the C++ and .NET interfaces to WebSphere MQ.

Why should I use WebSphere MQ classes for Java?

A Java application can use either WebSphere MQ classes for Java or WebSphere MQ classes for JMS to access WebSphere MQ resources. There are a number of advantages to using WebSphere MQ classes for Java.

If the following points are significant in your installation, consider using Websphere MQ classes for Java:

- WebSphere MQ classes for Java encapsulate the Message Queue Interface (MQI), the native WebSphere MQ API.
 - If you are familiar with the use of the MQI in procedural languages, you can transfer this knowledge to the Java environment.
 - You can exploit the full range of features of WebSphere MQ, beyond those available through JMS.
- WebSphere MQ classes for Java use a similar object model to the C++ and .NET interfaces to WebSphere MQ. If you are familiar with these interfaces, you can transfer this knowledge to the Java environment.

Connection options for WebSphere MQ classes for Java

WebSphere MQ classes for Java can connect in client or bindings mode.

Programmable options allow WebSphere MQ classes for Java to connect to WebSphere MQ in either of the following ways:

- As a WebSphere MQ MQI client using Transmission Control Protocol/Internet Protocol (TCP/IP)
- In bindings mode, connecting directly to WebSphere MQ using the Java Native Interface (JNI)

Clients cannot be run on z/OS, but clients on other platforms can connect to a WebSphere MQ for z/OS queue manager if the Client Attach Facility is installed.

The following sections describe the client mode and bindings mode connection options in more detail.

Client connection

To connect to a queue manager in client mode, a WebSphere MQ classes for Java application can run on the same system on which the queue manager is running, or on a different system. In each case, WebSphere MQ classes for Java connects to the queue manager over TCP/IP.

A WebSphere MQ classes for Java application can connect to any supported queue manager using client mode.

For more information on how to write applications to use client mode connections, see [“WebSphere MQ classes for Java connection modes”](#) on page 639.

Bindings connection

When used in bindings mode, WebSphere MQ classes for Java uses the Java Native Interface (JNI) to call directly into the existing queue manager API, rather than communicating through a network. In most environments, connecting in bindings mode provides better performance for WebSphere MQ classes for Java applications than connecting in client mode, by avoiding the cost of TCP/IP communication.

Applications that use the WebSphere MQ classes for Java to connect in bindings mode must run on the same system as the queue manager to which they are connecting.

The Java Runtime Environment, that is being used to run the WebSphere MQ classes for Java application, must be configured to load the WebSphere MQ classes for Java libraries; see [The WebSphere MQ classes for Java libraries](#) for further information.

For more information on how to write applications to use bindings mode connections, see [“WebSphere MQ classes for Java connection modes”](#) on page 639.

Prerequisites for WebSphere MQ classes for Java

To use WebSphere MQ classes for Java, you need certain other software products.

For the latest information about the prerequisites for WebSphere MQ classes for Java, see the [WebSphere MQ README file](#).

To develop WebSphere MQ classes for Java applications, you need a Java Development Kit (JDK). Details of the JDKs supported with your operating system can be found on the WebSphere MQ system requirements page at [System Requirements for IBM WebSphere MQ](#).

To run WebSphere MQ classes for Java applications, you need the following software components:

- A WebSphere MQ queue manager, for applications that connect to a queue manager
- A Java Runtime Environment (JRE), for each system on which you run applications. A suitable JRE is supplied with WebSphere MQ.

If you require SSL connections to use cryptographic modules that have been FIPS 140-2 certified, you need the IBM Java JSSE FIPS provider (IBMJSEFIPS). Every IBM JDK and JRE at Version 1.4.2 or later contains IBMJSSEFIPS.

You can use Internet Protocol Version 6 (IPv6) addresses in your WebSphere MQ classes for Java applications if IPv6 is supported by your Java virtual machine (JVM) and the TCP/IP implementation on your operating system.

Installation and configuration of WebSphere MQ classes for Java

This section describes the directories and files that are created when you install WebSphere MQ classes for Java, and tells you how to configure WebSphere MQ classes for Java after installation.

What is installed for WebSphere MQ classes for Java

The latest version of WebSphere MQ classes for Java is installed with WebSphere MQ. You might need to override default installation options to make sure this is done.

For more information about installing WebSphere MQ see:

[Installing a WebSphere MQ server](#)

[Installing an IBM WebSphere MQ client](#)

WebSphere MQ classes for Java are contained in the Java archive (JAR) files, `com.ibm.mq.jar`, and `com.ibm.mq.jmqi.jar`.

Support for standard message headers, such as Programmable Command Format (PCF), is contained in the JAR file `com.ibm.mq.headers.jar`.

Support for Programmable Command Format (PCF) is contained in the JAR file `com.ibm.mq.pcf.jar`.

Installing and upgrading the WebSphere MQ classes for Java JAR files

The only supported way to get the WebSphere MQ classes for Java JAR files onto a system, is to install either the WebSphere MQ product, or the WebSphere MQ MQI client SupportPac, or to use a software management tool such as Apache Maven, for more information see [“IBM WebSphere MQ classes for Java and software management tools”](#) on page 635.

Do not move, or copy the WebSphere MQ classes for Java JAR files from other machines, unless you are using a software management tool.

- Fix packs cannot be applied to an "installation" where JAR files have been copied from another machine, and makes it much more difficult to ensure that all the JAR files are kept in step with each other, and are at compatible levels.
- Copying the WebSphere MQ classes for JMS JAR files between machines can also result in multiple copies of the files residing on the same machine, which can cause problems servicing the code, and debugging problems.

Do not include the WebSphere MQ classes for Java JAR files within application archives.

- Updates to the WebSphere MQ classes for Java cannot be applied using a WebSphere MQ Fix Pack.
- It is not possible for IBM Support to easily determine the version of the WebSphere MQ classes for Java that are being used by the application.
- Problems can arise if multiple applications running inside the same Java Runtime Environment include different versions of the WebSphere MQ classes for Java, as multiple versions of the WebSphere MQ classes for Java are loaded into the Java Runtime Environment at the same time.
- If an application uses the BINDINGS transport to connect to a queue manager, any major upgrades to the queue manager also require the application to be updated to include the corresponding level of the WebSphere MQ classes for Java.

For example, if a queue manager is upgraded to the WebSphere MQ Version 7.1 level, then any applications that connect to the queue manager using the BINDINGS transport also need to be updated to include the WebSphere MQ Version 7.1 classes for Java.

The following Java library is distributed with WebSphere MQ classes for Java:

- connector.jar (Version 1.0)

The sample application called Postcard is in the JAR file com.ibm.mq.postcard.jar.

The Javadoc tool has been used to generate the HTML pages containing the specifications of the WebSphere MQ classes for Java and WebSphere MQ classes for JMS APIs. The HTML pages are in the doc subdirectory of the WebSphere MQ classes for JMS installation directory. On UNIX, Linux, and Windows systems, the doc subdirectory contains the individual HTML pages.

When installation is complete, files and samples are installed in the locations shown in “[Installation directories for WebSphere MQ classes for Java](#)” on page 628.

After installation, on any platform other than Windows, you must update your environment variables as described in “[Environment variables relevant to WebSphere MQ classes for Java](#)” on page 629.

Installation directories for WebSphere MQ classes for Java

WebSphere MQ classes for Java files are installed in different locations according to platform.

Table 82 on page 628 shows where the WebSphere MQ classes for Java files are installed.

<i>Table 82. WebSphere MQ classes for Java installation directories</i>	
Platform	Directory
AIX	<i>MQ_INSTALLATION_PATH</i> /java/lib
HP-UX, Linux, and Solaris	<i>MQ_INSTALLATION_PATH</i> /java/lib
Windows	<i>MQ_INSTALLATION_PATH</i> \java\lib
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which WebSphere MQ is installed.	

Some sample applications, such as the Installation Verification Programs (IVPs), are supplied with WebSphere MQ. Table 83 on page 628 shows where the sample applications are installed. The WebSphere MQ classes for Java samples are in a subdirectory called wmqjava. The PCF samples are in a subdirectory called pc.f.

<i>Table 83. Samples directories</i>	
Platform	Directory
AIX	<i>MQ_INSTALLATION_PATH</i> /samp/wmqjava/
HP-UX, Linux, and Solaris	<i>MQ_INSTALLATION_PATH</i> /samp/wmqjava/
Windows	<i>MQ_INSTALLATION_PATH</i> \tools\wmqjava\
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which WebSphere MQ is installed.	

Environment variables relevant to WebSphere MQ classes for Java

If you want to run WebSphere MQ classes for Java applications, their class paths must include the WebSphere MQ classes for Java and samples directories.

For WebSphere MQ classes for Java applications to run, their class path must include the appropriate WebSphere MQ classes for Java directory. To run the sample applications, the class path must also include the appropriate samples directories. This information can be provided in the Java invocation command or in the CLASSPATH environment variable.

Table 84 on page 629 shows the appropriate CLASSPATH setting to use on each platform to run WebSphere MQ classes for Java applications, including the sample applications.

Table 84. CLASSPATH setting to run WebSphere MQ classes for Java applications, including the WebSphere MQ classes for Java sample applications	
Platform	CLASSPATH setting
AIX	CLASSPATH=MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jar: MQ_INSTALLATION_PATH/samp/wmqjava/samples:
HP-UX, Linux, and Solaris	CLASSPATH=MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jar: MQ_INSTALLATION_PATH/samp/wmqjava/samples:
Windows	CLASSPATH=MQ_INSTALLATION_PATH\Java\lib\com.ibm.mq.jar; MQ_INSTALLATION_PATH\tools\wmqjava\samples;
MQ_INSTALLATION_PATH represents the high-level directory in which WebSphere MQ is installed.	

If you compile using the -Xlint option, you might see a message warning you that com.ibm.mq.es.e.jar is not present. You can ignore the warning. This file is only present if you have installed IBM WebSphere MQ Advanced Message Security.

The scripts provided with WebSphere MQ classes for Java use the following environment variables:

MQ_JAVA_DATA_PATH

This environment variable specifies the directory for log and trace output.

MQ_JAVA_INSTALL_PATH

This environment variable specifies the directory where WebSphere MQ classes for Java are installed, as shown in [WebSphere MQ classes for Java installation directories](#).

MQ_JAVA_LIB_PATH

This environment variable specifies the directory where the WebSphere MQ classes for Java libraries are stored, as shown in [The location of the WebSphere MQ classes for Java libraries for each platform](#). Some scripts supplied with WebSphere MQ classes for Java, such as IVTRun, use this environment variable.

On Windows, all the environment variables are set automatically during installation. On any other platform, you must set them yourself. On a UNIX system, you can use the script **setjmsenv** (if you are using a 32-bit JVM) or **setjmsenv64** (if you are using a 64-bit JVM) to set the environment variables. On AIX, HP-UX, Linux, and Solaris, these scripts are in the MQ_INSTALLATION_PATH/java/bin directory.

The IBM WebSphere MQ classes for Java libraries

The location of the IBM WebSphere MQ classes for Java libraries varies according to platform. Specify this location when you start an application.

To specify the location of the Java Native Interface (JNI) libraries, start your application using a **java** command with the following format:

```
java -Djava.library.path=library_path application_name
```

where *library_path* is the path to the WebSphere MQ classes for Java libraries , which include the JNI libraries. Table 85 on page 630 shows the location of the WebSphere MQ classes for Java libraries for each platform.

<i>Table 85. The location of the WebSphere MQ classes for Java libraries for each platform</i>	
Platform	Directory containing the WebSphere MQ classes for Java libraries
AIX	<i>MQ_INSTALLATION_PATH</i> /java/lib (32 bit libraries) <i>MQ_INSTALLATION_PATH</i> /java/lib64 (64 bit libraries)
HP-UX Linux (POWER®, x86-64 and zSeries s390x platforms) Solaris (x86-64 and SPARC platforms)	<i>MQ_INSTALLATION_PATH</i> /java/lib (32 bit libraries) <i>MQ_INSTALLATION_PATH</i> /java/lib64 (64 bit libraries)
Linux (x86 platform)	<i>MQ_INSTALLATION_PATH</i> /java/lib
Windows	<i>MQ_INSTALLATION_PATH</i> \Java\lib (32 bit libraries) <i>MQ_INSTALLATION_PATH</i> \Java\lib64 (64 bit libraries)
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which WebSphere MQ is installed.	

Note:

1. On AIX, HP-UX, Linux (Power platform), or Solaris, use either the 32 bit libraries or the 64 bit libraries. Use the 64 bit libraries only if you are running your application in a 64 bit Java virtual machine (JVM) on a 64 bit platform. Otherwise, use the 32 bit libraries.
2. On Windows, you can use the PATH environment variable to specify the location of the WebSphere MQ classes for Java libraries instead of specifying their location on the **java** command.
3. To use WebSphere MQ classes for Java in bindings mode on IBM i, ensure that the library QMQMJAVA is in your library list.

Related tasks

[Using WebSphere MQ classes for Java](#)

Support for OSGi on IBM WebSphere MQ classes for Java

OSGi provides a framework that supports the deployment of applications as bundles. One OSGi bundle is supplied as part of the IBM WebSphere MQ classes for Java .

OSGi provides a general purpose, secure, and managedJava framework, which supports the deployment of applications that come in the form of bundles. OSGi-compliant devices can download and install bundles, and remove them when they are no longer required. The framework manages the installation and update of bundles in a dynamic and scalable fashion.

The IBM WebSphere MQ classes for Java. includes the following OSGi bundle.

com.ibm.mq.osgi.java_<version number>.jar

The JAR files to allow applications to use the IBM WebSphere MQ classes for Java.

where <version number> is the version number of WebSphere MQ that has been installed.

The bundle is installed into the `java/lib/OSGi` subdirectory of your IBM WebSphere MQ installation, or the `java\lib\OSGi` folder on Windows.

Nine other bundles are also installed into the `java/lib/OSGi` subdirectory of your IBM WebSphere MQ installation, or the `java\lib\OSGi` folder on Windows. These bundles are part of the IBM WebSphere MQ classes for JMS , and must not be loaded into an OSGi runtime environment that has the IBM

WebSphere MQ classes for Java bundle loaded. If the IBM WebSphere MQ classes for Java OSGi bundle is loaded into an OSGi runtime environment that also has the IBM WebSphere MQ classes for JMS bundles loaded, errors such as:

```
java.lang.ClassCastException: com.ibm.mq.MQException incompatible with com.ibm.mq.MQException
```

occur, when applications using either the IBM WebSphere MQ classes for Java bundle or the IBM WebSphere MQ classes for JMS bundles are run.

The OSGi bundle for the IBM WebSphere MQ classes for Java has been written to the OSGi Release 4 specification; it does not work in an OSGi Release 3 environment.

You must set your system path or library path correctly so that the OSGi runtime environment can find any required DLL files or shared libraries.

If you use the OSGi bundle for the IBM WebSphere MQ classes for Java, channel exit classes written in Java are not supported because of an inherent problem in loading classes in a multiple class loader environment such as OSGi. A user bundle can be aware of the IBM WebSphere MQ classes for Java bundle, but the IBM WebSphere MQ classes for Java bundle is not aware of any user bundle. As a result, the class loader used in a IBM WebSphere MQ classes for Java bundle cannot load a channel exit class that is in a user bundle.

For more information about OSGi, see the [OSGi alliance](#) website.

The IBM WebSphere MQ classes for Java configuration file

An IBM WebSphere MQ classes for Java configuration file specifies properties that are used to configure the IBM WebSphere MQ classes for Java.

The format of an IBM WebSphere MQ classes for Java configuration file is that of a standard Java properties file.

V7.5.0.9 From IBM WebSphere MQ Version 7.5.0, Fix Pack 9, a sample configuration file that is called `mqjava.config` is supplied in the `bin` subdirectory of the IBM WebSphere MQ classes for Java installation directory. This file documents all the supported properties and their default values.

Note: The sample configuration file is overwritten when the IBM WebSphere MQ installation is upgraded to a future Fix Pack. Therefore, it is recommended that you make a copy of the sample configuration file for use with your applications.

You can choose the name and location of an IBM WebSphere MQ classes for Java configuration file. When you start your application, use a **java** command with the following format:

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

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

The following example shows a **java** command:

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

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

An IBM WebSphere MQ classes for Java configuration file can be used with any of the supported transports between an application and a queue manager or broker.

Overriding properties specified in an IBM WebSphere MQ classes for Java configuration file

An IBM WebSphere MQ MQI client configuration file can also specify properties that are used to configure IBM WebSphere MQ classes for Java. However, properties that are specified in an IBM WebSphere MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in an IBM WebSphere MQ MQI client configuration file by specifying it as a property in an IBM WebSphere MQ classes for Java configuration file. To override an attribute in an IBM WebSphere MQ MQI client configuration file, use an entry with the following format in the IBM WebSphere MQ classes for Java configuration file:

```
com.ibm.mq.cfg.stanza.propName=propValue
```

The variables in the entry have the following meanings:

stanza

The name of the stanza in the IBM WebSphere MQ MQI client configuration file that contains the attribute.

propName

The name of the attribute as specified in the IBM WebSphere MQ MQI client configuration file.

propValue

The value of the property that overrides the value of the attribute that is specified in the IBM WebSphere MQ MQI client configuration file.

Alternatively, you can override an attribute in an IBM WebSphere MQ MQI client configuration file by specifying the property as a system property on the **java** command. Use the preceding format to specify the property as a system property.

Only the following attributes in an IBM WebSphere MQ MQI client configuration file are relevant to IBM WebSphere MQ classes for Java. If you specify or override other attributes, it has no effect. Specifically, note that the `ChannelDefinitionFile` and `ChannelDefinitionDirectory` in the `CHANNELS` stanza of the client configuration file are not used. See [“Using a client channel definition table with IBM WebSphere MQ classes for Java” on page 643](#) for details of how to use the CCDT with the IBM WebSphere MQ classes for Java.

Table 86. Which stanza of the client configuration file contains which attribute	
Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	CIntRcvBufSize
TCP stanza of the client configuration file	CIntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

For more information on the IBM WebSphere MQ MQI client configuration, see [Configuring a client using a configuration file](#).

Related tasks

[Tracing IBM WebSphere MQ classes for Java applications](#)

Java Standard Environment Trace stanza

You can use the Java Standard Environment Trace Settings stanza to configure the IBM WebSphere MQ classes for Java trace facility.

com.ibm.msg.client.commonservices.trace.outputName = traceOutputName

traceOutputName is the directory and file name to which trace output is sent.

The default name of the trace file depends on the version of the IBM WebSphere MQ classes for Java that are being used by an application:

- For IBM WebSphere MQ classes for Java for Version 7.5.0, Fix Pack 8 or earlier, *traceOutputName* defaults to a file named `mqjms_%PID%.trc` in the current working directory.
- **V7.5.0.9** From IBM WebSphere MQ classes for Java from Version 7.5.0, Fix Pack 9, *traceOutputName* defaults to a file named `mqjava_%PID%.trc` in the current working directory.

where `%PID%` is the current process ID. If a process ID is unavailable, a random number is generated and prefixed with the letter `f`. To include the process ID in a file name that you specify, use the string `%PID%`.

If you specify an alternative directory, it must exist, and you must have write permission for this directory. If you do not have write permission, the trace output is written to `System.err`.

com.ibm.msg.client.commonservices.trace.include = includeList

includeList is a list of packages and classes that are traced, or the special values ALL or NONE.

Separate package or class names with a semicolon (;). **includeList** defaults to ALL, and traces all packages and classes in IBM WebSphere MQ classes for Java.

Note: You can include a package but then exclude subpackages of that package. For example, if you include package `a.b` and exclude package `a.b.x`, the trace includes everything in `a.b.y` and `a.b.z`, but not `a.b.x` or `a.b.x.1`.

com.ibm.msg.client.commonservices.trace.exclude = excludeList

excludeList is a list of packages and classes that are not traced, or the special values ALL or NONE.

Separate package or class names with a semicolon (;). **excludeList** defaults to NONE, and therefore excludes no packages and classes in IBM WebSphere MQ classes for Java from being traced.

Note: You can exclude a package but then include subpackages of that package. For example, if you exclude package `a.b` and include package `a.b.x`, the trace includes everything in `a.b.x` and `a.b.x.1`, but not `a.b.y` or `a.b.z`.

Any package or class that is specified, at the same level, as both included and excluded is included.

com.ibm.msg.client.commonservices.trace.maxBytes = maxArrayBytes

maxArrayBytes is the maximum number of bytes that are traced from any byte arrays.

If **maxArrayBytes** is set to a positive integer, it limits the number of bytes in a byte-array that are written out to the trace file. It truncates the byte array after writing *maxArrayBytes* out. Setting **maxArrayBytes** reduces the size of the resulting trace file, and reduces the effect of tracing on the performance of the application.

A value of 0 for this property means that none of the contents of any byte arrays are sent to the trace file.

The default value is -1, which removes any limit on the number of bytes in a byte array that are sent to the trace file.

com.ibm.msg.client.commonservices.trace.limit = maxTraceBytes

maxTraceBytes is the maximum number of bytes that are written to a trace output file.

maxTraceBytes works with **traceCycles**. If the number of bytes of trace written is near to the limit, the file is closed, and a new trace output file is started.

A value of 0 means that a trace output file has zero length. The default value is -1, which means that the amount of data to be written to a trace output file is unlimited.

com.ibm.msg.client.commonservices.trace.count = traceCycles

traceCycles is the number of trace output files to cycle through.

If the current trace output file reaches the limit that is specified by **maxTraceBytes**, the file is closed. Further trace output is written to the next trace output file in sequence. Each trace output file is distinguished by a numeric suffix that is appended to the file name. The current or most recent trace output file has the suffix `.trc.0`, the next most recent trace output file ends with `.trc.1`, and so on. Older trace files follow the same numbering pattern up to the limit.

The default value of **traceCycles** is 1. If **traceCycles** is 1, when the current trace output file reaches its maximum size, the file is closed and deleted. A new trace output file with the same name is started. Therefore, only one trace output file exists at a time.

com.ibm.msg.client.commonservices.trace.parameter = traceParameters

traceParameters controls whether method parameters and return values are included in the trace.

traceParameters defaults to TRUE. If **traceParameters** is set to FALSE, only method signatures are traced.

com.ibm.msg.client.commonservices.trace.compress = compressedTrace

Set **compressedTrace** to TRUE to compress trace output.

The default value of **compressedTrace** is FALSE.

If **compressedTrace** is set to TRUE, trace output is compressed. The default trace output file name has the extension `.trz`. If compression is set to FALSE, the default value, the file has the extension `.trc` to indicate it is uncompressed. However, if the file name for the trace output is specified in **traceOutputName** that name is used instead and no suffix is applied to the file.

Compressed trace output is smaller than uncompressed. Because there is less I/O, it can be written out faster than uncompressed trace. Compressed tracing has less effect on the performance of the IBM WebSphere MQ classes for Java than uncompressed tracing.

If **maxTraceBytes** and **traceCycles** are set, multiple compressed trace files are created in place of multiple flat files.

If the IBM WebSphere MQ classes for Java ends in an uncontrolled manner, a compressed trace file might not be valid. For this reason, trace compression must be used only when the IBM WebSphere MQ classes for Java close down in a controlled manner. Use trace compression only if the problems that are being investigated do not cause the JVM itself to stop unexpectedly. Do not use trace compression when you are diagnosing problems that can result in `System.Halt()` shutdowns or abnormal, uncontrolled JVM terminations.

com.ibm.msg.client.commonservices.trace.level = traceLevel

traceLevel specifies a filtering level for the trace. The defined trace levels are as follows:

Table 87. What is traced for each trace level	
Value	What is traced
0	Trace is turned off
1	Exceptions
3	Exceptions Warnings

Table 87. What is traced for each trace level (continued)	
Value	What is traced
6	Exceptions Warnings Informational trace points
8	Exceptions Warnings Informational trace points Method entry and exit
9	Exceptions Warnings Informational trace points Method entry and exit Data that is sent between the IBM WebSphere MQ classes for Java and a queue manager.

Note: Always use the value 9 unless otherwise directed by IBM Support.

IBM WebSphere MQ classes for Java and software management tools

Software management tools such as Apache Maven can be used with the IBM WebSphere MQ classes for Java.

Many large development organizations use these tools to centrally manage repositories of third-party libraries.

The IBM WebSphere MQ classes for Java are composed of a number of JAR files. When you are developing Java language applications by using this API, an installation of either an IBM WebSphere MQ Server, IBM WebSphere MQ Client, or IBM WebSphere MQ Client SupportPac is required on the machine where the application is being developed.

If you want to use a software management tool and add the JAR files that make up the IBM WebSphere MQ classes for Java to a centrally managed repository, the following points must be observed:

- A repository or container must be made available only to developers within your organization. Any distribution outside of the organization is not permitted.
- The repository needs to contain a complete and consistent set of JAR files from a single IBM WebSphere MQ release or Fix Pack.
- You are responsible for updating the repository with any maintenance provided by IBM Support.

For IBM WebSphere MQ Version 7.5, the following JAR files need to be installed into the repository:

- com.ibm.mq.commonservices.jar
- com.ibm.mq.jar
- com.ibm.mq.jmqi.jar
- com.ibm.mq.pcf.jar
- com.ibm.mq.headers.jar
- connector.jar

Post installation setup for IBM WebSphere MQ applications

After installing IBM WebSphere MQ, you can configure your installation so that you run your own applications.

Remember to check the IBM WebSphere MQ README file for later, or more specific, information for your environment.

Before attempting to run an IBM WebSphere MQ classes for Java application in bindings mode, make sure that you have configured IBM WebSphere MQ as described in [Configuring](#).

Configuring your queue manager to accept client connections from WebSphere MQ classes for Java

To configure your queue manager to accept incoming connection requests from clients, define and permit use of a server connection channel and start a listener program.

See [“Preparing and running the sample programs”](#) on page 104 for details.

Running WebSphere MQ classes for Java applications under the Java Security Manager

WebSphere MQ classes for Java can run with the Java Security Manager enabled. To run applications successfully with the Security Manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy definition file.

The simplest way to do this is to change the policy file supplied with your JRE. On most systems, this file is stored in the path `lib/security/java.policy`, relative to your JRE directory. You can edit policy files using your preferred editor or the `policytool` program supplied with your JRE.

You must give authority to the `com.ibm.mq.jmqi.jar` file so that it can:

- Create sockets (in client mode)
- Load the native library (in bindings mode)
- Read various properties from the environment

The system property `os.name` must be available to the WebSphere MQ classes for Java when running under the Java Security Manager.

Here is an example of a policy file entry that allows WebSphere MQ classes for Java to run successfully under the default security manager:

```
grant codeBase "file:MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jmqi.jar" {
    //Required
    permission java.util.PropertyPermission "user.name","read";
    permission java.util.PropertyPermission "os.name","read";
    //Required if mqclient.ini/mqs.ini configuration files are used
    permission java.io.FilePermission "/var/mqm/mqclient.ini","read";
    permission java.io.FilePermission "/var/mqm/mqs.ini","read";
    //For the client transport type.
    permission java.net.SocketPermission "*", "connect";
    //For the bindings transport type.
    permission java.lang.RuntimePermission "loadLibrary.*";
    //For applications that use CCDT tables (access to the CCDT
    AMQCLCHL.TAB)
    permission java.io.FilePermission
"/var/mqm/qmgrs/QMGR/@ipcc/AMQCLCHL.TAB", "read";
    //For applications that use User Exits
    permission java.io.FilePermission "/var/mqm/exits/*", "read";
    permission java.lang.RuntimePermission "createClassLoader";
    //Required for the z/OS platform
    permission java.util.PropertyPermission
"com.ibm.vm.bitmode","read";
};
grant codeBase
"file:MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.commonservices.jar" {
    permission java.util.PropertyPermission "user.dir","read";
    permission java.util.PropertyPermission "line.separator","read";
    //tracing permissions
    permission java.util.PropertyPermission "com.ibm.mq.commonservices", "read";
    permission java.util.logging.LoggingPermission "control";
    //For access to the trace properties file.
```

```

permission java.io.FilePermission "/tmp/trace.properties", "read";
//For access to the trace output files.
permission java.io.FilePermission "/tmp/*", "read,write";
};

```

Notes:

- `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.
- This example of a policy file enables the WebSphere MQ classes for Java to work correctly under the security manager, but you might still need to enable your own code to run correctly before your applications work.
- To allow WebSphere MQ classes for Java to access the Java archive (JAR) files of an application, add the following permission to the first grant statement:

```

permission java.io.FilePermission "/path_to_your_app/-", "read";

```

- To use these grant statements in your policy configuration file, you might need to modify the path names depending on where you have installed WebSphere MQ classes for Java and where you store your applications.
- The sample code shipped with WebSphere MQ classes for Java has not been specifically enabled for use with the security manager; however the IVT tests run with this policy file and the default security manager in place.

Verifying the IBM WebSphere MQ classes for Java installation

An installation verification program, MQIVP, is supplied with IBM WebSphere MQ classes for Java. You can use this program to test all the connection modes of IBM WebSphere MQ classes for Java.

The program prompts for a number of choices and other data to determine which connection mode you want to verify. Use the following procedure to verify your installation:

1. If you are going to run the program in client mode, configure your queue manager as described in [“Preparing and running the sample programs” on page 104](#). The queue to use is `SYSTEM.DEFAULT.LOCAL.QUEUE`.
2. If you are going to run the program in client mode, see also [“Using WebSphere MQ classes for Java” on page 625](#).

Perform the remaining steps of this procedure on the system on which you are going to run the program.

3. Make sure that you have updated your `CLASSPATH` environment variable according to the instructions in [“Environment variables relevant to WebSphere MQ classes for Java” on page 629](#).
4. Change Directory to `MQ_INSTALLATION_PATH/mqm/VRM/java/samples/wmqjava`, where `MQ_INSTALLATION_PATH` is the path to your IBM WebSphere MQ installation and `VRM` is the version, release, and modification number of the product. Then at the command prompt, enter:

```

java -Djava.library.path=library_path MQIVP

```

where *library_path* is the path to the IBM WebSphere MQ classes for Java libraries (see [The WebSphere MQ classes for Java libraries](#)).

At the prompt marked (1):

- To use a TCP/IP connection, enter a IBM WebSphere MQ server host name.
- To use native connection (bindings mode), leave the field blank (do not enter a name).

The program tries to:

- 1. Connect to the queue manager
- 2. Open the queue `SYSTEM.DEFAULT.LOCAL.QUEUE`, put a message on the queue, get a message from the queue, and then close the queue

- 3. Disconnect from the queue manager
- 4. Return a message if the operations are successful

Here is an example of the prompts and responses you might see. The actual prompts and your responses depend on your IBM WebSphere MQ network.

```

Please enter the IP address of the MQ server           : ipaddress(1)
Please enter the port to connect to                   : (1414)(2)
Please enter the server connection channel name       : channelname(2)
Please enter the queue manager name                  : qmname
Success: Connected to queue manager.
Success: Opened SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Put a message to SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Got a message from SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Closed SYSTEM.DEFAULT.LOCAL.QUEUE
Success: Disconnected from queue manager

Tests complete -
SUCCESS: This MQ Transport is functioning correctly.
Press Enter to continue ...

```

Note:

1. If you choose server connection, you do not see the prompts marked ⁽²⁾.

Solving IBM WebSphere MQ problems

Initially, run the installation verification program. You might also have to use the trace facility.

If a program does not complete successfully, run the installation verification program, and follow the advice given in the diagnostic messages. This program is described in [“Verifying the IBM WebSphere MQ classes for Java installation” on page 637](#).

If the problems continue and you need to contact the IBM service team, you might be asked to turn on the trace facility. Do this as shown in the following example.

To trace the MQIVP program:

- Create a *com.ibm.mq.commonservices* properties file (see [Using com.ibm.mq.commonservices](#) .
- Enter the following command:

```

java -Dcom.ibm.mq.commonservices=commonservices_properties_file java
-Djava.library.path=library_path MQIVP -trace

```

where:

- *commonservices_properties_file* is the path (including the filename) to the *com.ibm.mq.commonservices* properties file.
- *library_path* is the path to the WebSphere MQ classes for Java libraries (see [The WebSphere MQ classes for Java libraries](#)).

For more information about how to use trace, see [Tracing IBM WebSphere MQ classes for Java applications](#).

Introduction for programmers

This collection of topics contains general information for programmers.

For more detailed information about writing programs, see [“Writing WebSphere MQ classes for Java applications” on page 639](#).

The WebSphere MQ classes for Java interface

The procedural WebSphere MQ application programming interface uses verbs, which act on objects. The Java programming interface uses objects, which you act upon by calling methods.

The procedural WebSphere MQ application programming interface is built around verbs such as these:

```
MQBACK, MQBEGIN, MQCLOSE, MQCONN, MQDISC,  
MQGET, MQINQ, MQOPEN, MQPUT, MQSET, MQSUB
```

These verbs all take, as a parameter, a handle to the WebSphere MQ object on which they are to operate. Because Java is object-oriented, the Java programming interface turns this round. Your program consists of a set of WebSphere MQ objects, which you act upon by calling methods on those objects.

When you use the procedural interface, you disconnect from a queue manager by using the call `MQDISC(Hconn, CompCode, Reason)`, where *Hconn* is a handle to the queue manager.

In the Java interface, the queue manager is represented by an object of class `MQQueueManager`. You disconnect from the queue manager by calling the `disconnect()` method on that class.

```
// declare an object of type queue manager  
MQQueueManager queueManager=new MQQueueManager();  
...  
// do something...  
...  
// disconnect from the queue manager  
queueManager.disconnect();
```

Writing WebSphere MQ classes for Java applications

This collection of topics provides information to assist with writing Java applications to interact with WebSphere MQ systems.

To use WebSphere MQ classes for Java to access WebSphere MQ queues, you write Java applications that contain calls that put messages onto, and get messages from, WebSphere MQ queues. For details of individual classes, see [WebSphere MQ classes for Java](#).

Note: Automatic client reconnection is not supported by WebSphere MQ classes for Java .

WebSphere MQ classes for Java connection modes

The way you program for WebSphere MQ classes for Java has some dependencies on the connection modes you want to use.

If you use client connections, there are a number of differences from the IBM WebSphere MQ MQI client but it is conceptually similar. If you use bindings mode you can use fastpath bindings and can issue the `MQBEGIN` command. You specify which mode to use by setting variables in the `MQEnvironment` class.

WebSphere MQ classes for Java client connections

When WebSphere MQ classes for Java is used as a client, it is like the IBM WebSphere MQ MQI client, but has a number of differences.

If you are programming for *WebSphere MQ classes for Java* for use as a client, be aware of the following differences:

- It supports only TCP/IP.
- It does not read any WebSphere MQ environment variables at startup.
- Information that would be stored in a channel definition and in environment variables can be stored in a class called `Environment`. Alternatively, this information can be passed as parameters when the connection is made.
- Error and exception conditions are written to a log specified in the `MQException` class. The default error destination is the Java console.

- Only the following attributes in a WebSphere MQ client configuration file are relevant to WebSphere MQ classes for Java. If you specify other attributes, they are ineffective.

Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	ClntRcvBufSize
TCP stanza of the client configuration file	ClntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

- If connecting to a queue manager that requires character data to be converted, then the V7 Java client is now capable of doing the conversion if queue manager is unable to do so. The client JVM must support the conversion between the CCSID of the client and that of the queue manager.
- Automatic client reconnect is not supported by WebSphere MQ classes for Java.

When used in client mode, *WebSphere MQ classes for Java* does not support the MQBEGIN call.

See [“Connection options for WebSphere MQ classes for Java” on page 626](#) for further information concerning supported environments.

WebSphere MQ classes for Java bindings mode

The bindings mode of WebSphere MQ classes for Java differs from the client mode in three main ways.

When used in bindings mode, WebSphere MQ classes for Java uses the Java Native Interface (JNI) to call directly into the existing queue manager API, rather than communicating through a network.

By default, applications that use the WebSphere MQ classes for Java in bindings mode connect to a queue manager using the *ConnectOption*, MQCNO_STANDARD_BINDINGS.

The WebSphere MQ classes for Java support the following *ConnectOptions*:

- MQCNO_FASTPATH_BINDING
- MQCNO_STANDARD_BINDING
- MQCNO_SHARED_BINDING
- MQCNO_ISOLATED_BINDING

For further information on *ConnectOptions*, see [“Connecting to a queue manager using the MQCONN call” on page 200](#).

Bindings mode supports the MQBEGIN call to initiate global units of work that are coordinated by the queue manager, on all platforms apart from WebSphere MQ for IBM i and WebSphere MQ for z/OS.

Most of the parameters provided by the MQEnvironment class are not relevant to bindings mode and are ignored.

See [“Connection options for WebSphere MQ classes for Java”](#) on page 626 for further information concerning supported environments.

Defining which WebSphere MQ classes for Java connection to use

The type of connection to use is determined by the setting of variables in the MQEnvironment class.

Two variables are used:

MQEnvironment.properties

The connection type is determined by the value associated with the key name CMQC.TRANSPORT_PROPERTY. Possible values are as follows:

CMQC.TRANSPORT_MQSERIES_BINDINGS

Connect in bindings mode

CMQC.TRANSPORT_MQSERIES_CLIENT

Connect in client mode

CMQC.TRANSPORT_MQSERIES

Connection mode is determined by the value of the *hostname* property

MQEnvironment.hostname

Set the value of this variable as follows:

- For client connections, set the value of this variable to the host name of the IBM WebSphere MQ server to which you want to connect
- For bindings mode, do not set this variable, or set it to null

Operations on queue managers

This collection of topics describes how to connect to, and disconnect from, a queue manager using WebSphere MQ classes for Java.

Setting up the WebSphere MQ environment for WebSphere MQ classes for Java

For an application to connect to a queue manager in client mode, the application must specify the channel name, host name, and port number.

Note: The information in this topic is relevant only if your application connects to a queue manager in client mode. It is *not* relevant if it connects in bindings mode. See: [“Connection modes for WebSphere MQ classes for JMS”](#) on page 738

You can specify the channel name, host name, and port number in one of two ways: either as fields in the MQEnvironment class or as properties of the MQQueueManager object.

If you set fields in the MQEnvironment class, they apply to your whole application, except where they are overridden by a properties hash table. To specify the channel name and host name in MQEnvironment, use the following code:

```
MQEnvironment.hostname = "host.domain.com";
MQEnvironment.channel = "java.client.channel";
```

This is equivalent to setting an **MQSERVER** environment variable:

```
"java.client.channel/TCP/host.domain.com".
```

By default, the Java clients attempt to connect to a WebSphere MQ listener at port 1414. To specify a different port, use the following code:

```
MQEnvironment.port = nnnn;
```

where nnnn is the required port number

If you pass properties to a queue manager object at its creation, they apply only to that queue manager. Create entries in a Hashtable object with keys of **hostname**, **channel**, and, optionally, **port**, and with appropriate values. To use the default port, 1414, you can omit the **port** entry. Create the MQQueueManager object by using a constructor that accepts the properties hash table.

Identifying a connection to the queue manager by setting an application name

An application can set a name that identifies its connection to the queue manager. This application name is shown by the **DISPLAY CONN MQSC/PCF** command (where the field is called **APPLTAG**) or in the WebSphere MQ Explorer **Application Connections** display (where the field is called **App name**).

Application names are limited to 28 characters and longer names are truncated to fit. If an application name is not specified, a default is provided. The default name is based on the invoking (main) class, but if this information is not available, the text WebSphere MQ Client for Java is used.

If the name of the invoking class is used, it is adjusted to fit by removing leading package names, if necessary. For example, if the invoking class is `com.example.MainApp`, the full name is used, but if the invoking class is `com.example.dictionaryAndThesaurus.multilingual.mainApp`, the name `multilingual.mainApp` is used, because it is the longest combination of class name and rightmost package name that fits into the available length.

If the class name itself is more than 28 characters long, it is truncated to fit. For example, `com.example.mainApplicationForSecondTestCase` becomes `mainApplicationForSecondTest`.

Note: Queue managers running on z/OS platforms do not support setting application names.

To set an application name in the MQEnvironment class, add the name to the MQEnvironment.properties hash table, with a key of **MQConstants.APPNAME_PROPERTY**, by using the following code:

```
MQEnvironment.properties.put(MQConstants.APPNAME_PROPERTY, "my_application_name");
```

To set an application name in the properties hash table that is passed to the MQQueueManager constructor, add the name to the properties hash table with a key of **MQConstants.APPNAME_PROPERTY**.

Overriding properties specified in a WebSphere MQ client configuration file

A WebSphere MQ client configuration file can also specify properties that are used to configure WebSphere MQ classes for Java. However, properties specified in a WebSphere MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in a WebSphere MQ configuration file in any of the following ways. The options are shown in order of precedence.

- Set a Java system property for the configuration property.
- Set the property in the MQEnvironment.properties map.
- On Java5 and later releases, set a system environment variable.

Only the following attributes in a WebSphere MQ client configuration file are relevant to WebSphere MQ classes for Java. If you specify or override other attributes, it has no effect.

Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize

Stanza	Attribute
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	CIntRcvBufSize
TCP stanza of the client configuration file	CIntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

Connecting to a queue manager in WebSphere MQ classes for Java

Connect to a queue manager by creating a new instance of the MQQueueManager class. Disconnect from a queue manager by calling the disconnect() method.

You are now ready to connect to a queue manager by creating a new instance of the MQQueueManager class:

```
MQQueueManager queueManager = new MQQueueManager("qMgrName");
```

To disconnect from a queue manager, call the disconnect() method on the queue manager:

```
queueManager.disconnect();
```

If you call the disconnect method, all open queues and processes that you have accessed through that queue manager are closed. However, it is good programming practice to close these resources explicitly when you finish using them. To do this, use the close() method on the relevant objects.

The commit() and backout() methods on a queue manager are equivalent to the MQCMIT and MQBACK calls that are used with the procedural interface.

Using a client channel definition table with IBM WebSphere MQ classes for Java

A IBM WebSphere MQ classes for Java client application can use client connection channel definitions stored in a client channel definition table (CCDT).

As an alternative to creating a client connection channel definition by setting certain fields and environment properties in the MQEnvironment class or passing them to an MQQueueManager in a properties hash table, a IBM WebSphere MQ classes for Java client application can use client connection channel definitions that are stored in a client channel definition table. These definitions are created by IBM WebSphere MQ Script (MQSC) commands or IBM WebSphere MQ Programmable Command Format (PCF) commands, or using the IBM WebSphere MQ Explorer.

When the application creates an MQQueueManager object, the IBM WebSphere MQ classes for Java client searches the client channel definition table for a suitable client connection channel definition, and uses the channel definition to start an MQI channel. For more information about client channel definition tables and how to construct one, see [Client channel definition table](#).

To use a client channel definition table, an application must first create a URL object. The URL object encapsulates a uniform resource locator (URL) that identifies the name and location of the file containing the client channel definition table and specifies how the file can be accessed.

For example, if the file ccdt1.tab contains a client channel definition table and is stored on the same system on which the application is running, the application can create a URL object in the following way:

```
java.net.URL chanTab1 = new URL("file:///home/admdata/ccdt1.tab");
```

As another example, suppose the file `ccdt2.tab` contains a client channel definition table and is stored on a system that is different from the one on which the application is running. If the file can be accessed using the FTP protocol, the application can create a URL object in the following way:

```
java.net.URL chanTab2 = new URL("ftp://ftp.server/admdata/ccdt2.tab");
```

After the application has created a URL object, the application can create an `MQQueueManager` object using one of the constructors that takes a URL object as a parameter. Here is an example:

```
MQQueueManager mairs = new MQQueueManager("MARS", chanTab2);
```

This statement causes the IBM WebSphere MQ classes for Java client to access the client channel definition table identified by the URL object `chanTab2`, search the table for a suitable client connection channel definition, and then use the channel definition to start an MQI channel to the queue manager called `MARS`.

Note the following points that apply if an application uses a client channel definition table:

- When the application creates an `MQQueueManager` object using a constructor that takes a URL object as a parameter, no channel name must be set in the `MQEnvironment` class, either as a field or as an environment property. If a channel name is set, the IBM WebSphere MQ classes for Java client throws an `MQException`. The field or environment property specifying the channel name is considered to be set if its value is anything other than null, an empty string, or a string containing all blank characters.
- The **queueManagerName** parameter on the `MQQueueManager` constructor can have one of the following values:
 - The name of a queue manager
 - An asterisk (*) followed by the name of a queue manager group
 - An asterisk (*)
 - Null, an empty string, or a string containing all blank characters

These are the same values that can be used for the **QMgrName** parameter on an `MQCONN` call issued by a client application that is using Message Queue Interface (MQI). For more information about the meaning of these values, see [“The Message Queue Interface overview” on page 187](#).

If your application uses connection pooling, see [“Controlling the default connection pool in WebSphere MQ classes for Java” on page 663](#).

- When the IBM WebSphere MQ classes for Java client finds a suitable client connection channel definition in the client channel definition table, it uses only the information extracted from this channel definition to start an MQI channel. Any channel related fields or environment properties that the application might have set in the `MQEnvironment` class are ignored.

In particular, note the following points if you are using Secure Sockets Layer (SSL):

- An MQI channel uses SSL only if the channel definition extracted from the client channel definition table specifies the name of a CipherSpec supported by the IBM WebSphere MQ classes for Java client.
- A client channel definition table also contains information about the location of Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs). The IBM WebSphere MQ classes for Java client uses only this information to access LDAP servers that hold CRLs.
- A client channel definition table can also contain the location of an Online Certificate Status Protocol (OCSP) responder. IBM WebSphere MQ classes for Java cannot use the OCSP information in a client channel definition table file. However, you can configure OCSP as described in the section [Using Online Certificate Protocol](#).

For more information about using SSL with a client channel definition table, see [Specifying that an MQI channel uses SSL](#).

Note also the following points if you are using channel exits:

- An MQI channel uses the channel exits and associated user data specified by the channel definition extracted from the client channel definition table in preference to channel exits and data specified using other methods.
- A channel definition extracted from a client channel definition table can specify channel exits that are written in Java, C, or C++. For more information about how to write a channel exit in Java, see [“Creating a channel exit in WebSphere MQ classes for Java” on page 657](#). For more information about how to write a channel exit in other languages, see [“Using channel exits not written in Java with WebSphere MQ classes for Java” on page 660](#).

Specifying a range of ports for IBM WebSphere MQ classes for Java client connections

You can specify a port, or a range of ports, that an application can bind to in either of two ways.

When a IBM WebSphere MQ classes for Java application attempts to connect to an IBM WebSphere MQ queue manager in client mode, a firewall might allow only those connections that originate from specified ports or range of ports. In this situation, you can specify a port, or a range of ports, that the application can bind to. You can specify the port(s) in the following ways:

- You can set the `localAddressSetting` field in the `MQEnvironment` class. Here is an example:

```
MQEnvironment.localAddressSetting = "192.0.2.0(2000,3000)";
```

- You can set the environment property `CMQC.LOCAL_ADDRESS_PROPERTY`. Here is an example:

```
(MQEnvironment.properties).put(CMQC.LOCAL_ADDRESS_PROPERTY,
                                "192.0.2.0(2000,3000)");
```

- When you can construct the `MQQueueManager` object, you can pass a properties hashtable containing a `LOCAL_ADDRESS_PROPERTY` with the value `"192.0.2.0(2000,3000)"`

In each of these examples, when the application later connects to a queue manager, the application binds to a local IP address and port number in the range 192.0.2.0(2000) to 192.0.2.0(3000).

In a system with more than one network interface, you can also use the `localAddressSetting` field, or the environment property `CMQC.LOCAL_ADDRESS_PROPERTY`, to specify which network interface must be used for a connection.

Connection errors might occur if you restrict the range of ports. If an error occurs, an `MQException` is thrown containing the IBM WebSphere MQ reason code `MQRC_Q_MGR_NOT_AVAILABLE` and the following message:

```
Socket connection attempt refused due to LOCAL_ADDRESS_PROPERTY restrictions
```

An error might occur if all the ports in the specified range are in use, or if the specified IP address, host name, or port number is not valid (a negative port number, for example).

Accessing queues, topics, and processes in WebSphere MQ classes for Java

To access queues, topics, and processes, use methods of the `MQQueueManager` class. The `MQOD` (object descriptor structure) is collapsed into the parameters of these methods.

Queues

To open a queue you can use the `accessQueue` method of the `MQQueueManager` class. For example, on a queue manager called `queueManager`, use the following code:

```
MQQueue queue = queueManager.accessQueue("qName", CMQC.MQOO_OUTPUT);
```

The `accessQueue` method returns a new object of class `MQQueue`.

When you have finished using the queue, use the `close()` method to close it, as in the following example:

```
queue.close();
```

You can also create a queue by using the `MQQueue` constructor. The parameters are exactly the same as for the `accessQueue` method, with the addition of a queue manager parameter. For example:

```
MQQueue queue = new MQQueue(queueManager,
                             "qName",
                             CMQC.MQOO_OUTPUT,
                             "qMgrName",
                             "dynamicQName",
                             "altUserID");
```

You can specify a number of options when you create queues. For details of these, see [Class.com.ibm.mq.MQQueue](#). Constructing a queue object in this way enables you to write your own subclasses of `MQQueue`.

Topics

Similarly, you can open a topic using the `accessTopic` method of the `MQQueueManager` class. For example, on a queue manager called `queueManager`, use the following code to create a subscriber and publisher:

```
MQTopic subscriber =
    queueManager.accessTopic("TOPICSTRING", "TOPICNAME",
                             CMQC.MQTOPIC_OPEN_AS_SUBSCRIPTION, CMQC.MQSO_CREATE);
```

```
MQTopic publisher =
    queueManager.accessTopic("TOPICSTRING", "TOPICNAME",
                             CMQC.MQTOPIC_OPEN_AS_PUBLICATION, CMQC.MQOO_OUTPUT);
```

When you have finished using the topic, use the `close()` method to close it.

You can also create a topic by using the `MQTopic` constructor. The parameters are exactly the same as for the `accessTopic` method, with the addition of a queue manager parameter. For example:

```
MQTopic subscriber = new
    MQTopic(queueManager, "TOPICSTRING", "TOPICNAME",
            CMQC.MQTOPIC_OPEN_AS_SUBSCRIPTION, CMQC.MQSO_CREATE);
```

You can specify a number of options when you create topics. For details of these, see [Class.com.ibm.mq.MQTopic](#). Constructing a topic object in this way enables you to write your own subclasses of `MQTopic`.

A topic must be opened either for publication or for subscription. The `MQQueueManager` class has eight `accessTopic` methods and the `Topic` class has eight constructors. In each case, four of these have a **destination** parameter and four have a **subscriptionName** parameter (including two that have both). These can only be used to open the topic for subscriptions. The two remaining methods have an **openAs** parameter, and the topic can be opened for either publication or subscription depending on the value of the **openAs** parameter.

To create a topic as a durable subscriber use either an `accessTopic` method of the `MQQueueManager` class or an `MQTopic` constructor that accepts a subscription name and, in either case, set the `CMQC.MQSO_DURABLE` option.

Processes

To access a process, use the `accessProcess` method of the `MQQueueManager`. For example, on a queue manager called `queueManager`, use the following code to create an `MQProcess` object:

```
MQProcess process =
    queueManager.accessProcess("PROCESSNAME",
                               CMQC.MQOO_FAIL_IF QUIESCING);
```

To access a process, use the `accessProcess` method of the `MQQueueManager`.

The `accessProcess` method returns a new object of class `MQProcess`.

When you have finished using the process object, use the `close()` method to close it, as in the following example:

```
process.close();
```

You can also create a process by using the `MQProcess` constructor. The parameters are exactly the same as for the `accessProcess` method, with the addition of a queue manager parameter. For example:

```
MQProcess process =
    new MQProcess(queueManager, "PROCESSNAME",
        CMQC.MQ00_FAIL_IF QUIESCING);
```

Constructing a process object in this way enables you to write your own subclasses of `MQProcess`.

Handling messages in WebSphere MQ classes for Java

Messages are represented by the `MQMessage` class. You put and get messages using methods of the `MQDestination` class, which has subclasses of `MQQueue` and `MQTopic`.

Put messages onto queues or topics using the `put()` method of the `MQDestination` class. You get messages from queues or topics using the `get()` method of the `MQDestination` class. Unlike the procedural interface, where `MQPUT` and `MQGET` put and get arrays of bytes, the Java programming language puts and gets instances of the `MQMessage` class. The `MQMessage` class encapsulates the data buffer that contains the actual message data, together with all the `MQMD` (message descriptor) parameters and message properties that describe that message.

To build a new message, create a new instance of the `MQMessage` class, and use the `writeXXX` methods to put data into the message buffer.

When the new message instance is created, all the `MQMD` parameters are automatically set to their default values, as defined in [Initial values and language declarations for MQMD](#). The `put()` method of `MQDestination` also takes an instance of the `MQPutMessageOptions` class as a parameter. This class represents the `MQPMO` structure. The following example creates a message and puts it onto a queue:

```
// Build a new message containing my age followed by my name
MQMessage myMessage = new MQMessage();
myMessage.writeInt(25);

String name = "Charlie Jordan";
myMessage.writeInt(name.length());
myMessage.writeBytes(name);

// Use the default put message options...
MQPutMessageOptions pmo = new MQPutMessageOptions();

// put the message!
queue.put(myMessage, pmo);
```

The `get()` method of `MQDestination` returns a new instance of `MQMessage`, which represents the message just taken from the queue. It also takes an instance of the `MQGetMessageOptions` class as a parameter. This class represents the `MQGMO` structure.

You do not need to specify a maximum message size, because the `get()` method automatically adjusts the size of its internal buffer to fit the incoming message. Use the `readXXX` methods of the `MQMessage` class to access the data in the returned message.

The following example shows how to get a message from a queue:

```
// Get a message from the queue
MQMessage theMessage = new MQMessage();
MQGetMessageOptions gmo = new MQGetMessageOptions();
queue.get(theMessage, gmo); // has default values

// Extract the message data
int age = theMessage.readInt();
```



```
int strLen = theMessage.readInt();
byte[] strData = new byte[strLen];
theMessage.readFully(strData,0,strLen);
String name = new String(strData,0);
```

You can alter the number format that the read and write methods use by setting the *encoding* member variable.

You can alter the character set to use for reading and writing strings by setting the *characterSet* member variable.

See “[MQMessage class](#)” on [page 1023](#) for more information.

Note: The writeUTF() method of MQMessage automatically encodes the length of the string as well as the Unicode bytes it contains. When your message will be read by another Java program (using readUTF()), this is the simplest way to send string information.

Improving the performance of nonpersistent messages in WebSphere MQ classes for Java

To improve performance when browsing messages or consuming nonpersistent messages from a client application, you can use *read ahead*. Client applications using MQGET or asynchronous consumption will benefit from the performance improvements when browsing messages or consuming nonpersistent messages.

For general information about the read ahead facility, see the related topic.

In WebSphere MQ classes for Java, you use the CMQC.MQSO_READ_AHEAD and CMQC.MQSO_NO_READ_AHEAD properties of an MQQueue or MQTopic object to determine whether message consumers and queue browsers are allowed to use read ahead on that object.

Putting messages asynchronously using WebSphere MQ classes for Java

To put a message asynchronously, set MQPMO_ASYNC_RESPONSE.

You put messages onto queues or topics using the put() method of the MQDestination class. To put a message asynchronously, that is, allowing the operation to complete without waiting for a response from the queue manager, you can set MQPMO_ASYNC_RESPONSE in the options field of MQPutMessageOptions. To determine the success or failure of asynchronous puts, use the MQQueueManager.getAsyncStatus call.

Publish/subscribe in WebSphere MQ classes for Java

In WebSphere MQ classes for Java, the topic is represented by the MQTopic class, and you publish to it using the MQTopic.put() methods.

For general information about WebSphere MQ publish/subscribe, see [Introduction to WebSphere MQ publish/subscribe messaging](#).

Handling WebSphere MQ message headers with WebSphere MQ classes for Java

Java classes are provided representing different types of message header. Two helper classes are also provided.

Header objects are described by the MQHeader interface, which provides general-purpose methods for accessing header fields and for reading and writing message content. Each header type has its own class that implements the MQHeader interface and adds getter and setter methods for individual fields. For example, the MQRFH2 header type is represented by the MQRFH2 class; the MQDLH header type by the MQDLH class, and so on. The header classes perform any necessary data conversion automatically, and can read or write data in any specified numeric encoding or character set (CCSID).

Two helper classes, MQHeaderIterator and MQHeaderList, assist with reading and decoding (parsing) the header content in messages:

- The MQHeaderIterator class works like a java.util.Iterator. For as long as there are more headers in the message, the next() method returns true, and the nextHeader() or next() method returns the next header object.
- The MQHeaderList works like a java.util.List. Like the MQHeaderIterator, it parses header content, but it also allows you to search for particular headers, add new headers, remove existing headers, update header fields and then write the header content back to a message. Alternatively, you can create an empty MQHeaderList, then populate it with header instances and write it to a message once or repeatedly.

The MQHeaderIterator and MQHeaderList classes use the information in the MQHeaderRegistry to know which WebSphere MQ header classes are associated with particular message types and formats. The MQHeaderRegistry is configured with knowledge of all current WebSphere MQ formats and header types and their implementation classes, and you can also register your own header types.

Support is provided for the following commonly used Websphere MQ headers

- MQRFH - Rules and formatting header
- MQRFH2 - Like MQRFH, used to pass messages to and from a message broker belonging to WebSphere Message Broker. Also used to contain message properties
- MQCIH - CICS Bridge
- MQDLH - Dead letter header
- MQIIH - IMS information header
- MQRMH - reference message header
- MQSAPH - SAP header
- MQWIH - Work information header
- MQXQH - Transmission Queue header
- MQDH - Distribution header
- MQEPH - Encapsulated PCF header

You can also define classes representing your own headers.

To use an MQHeaderIterator to get an RFH2 header, either set MQGMO_PROPERTIES_FORCE_MQRFH2 in the GetMessageOptions, or set the queue property PROPCTL to FORCE.

Printing all the headers in a message using WebSphere MQ classes for Java

In this example, an instance of MQHeaderIterator parses the headers in an MQMessage that has been received from a queue. The MQHeader objects returned from the nextHeader() method display their structure and contents when their toString method is invoked.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeader;
import com.ibm.mq.headers.MQHeaderIterator;

...
MQMessage message = ... // Message received from a queue.
MQHeaderIterator it = new MQHeaderIterator (message);

while (it.hasNext ())
{
    MQHeader header = it.nextHeader ();

    System.out.println ("Header type " + header.type () + ": " + header);
}
```

Skipping over the headers in a message using WebSphere MQ classes for Java

In this example, the skipHeaders() method of MQHeaderIterator positions the message read cursor immediately after the last header.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderIterator;

...
MQMessage message = ... // Message received from a queue.
MQHeaderIterator it = new MQHeaderIterator (message);

it.skipHeaders ();
```

Finding the reason code in a dead-letter message using WebSphere MQ classes for Java

In this example, the read method populates the MQDLH object by reading from the message. After the read operation, the message read cursor is positioned immediately after the MQDLH header content.

Messages on the queue manager's dead-letter queue are prefixed with a dead-letter header (MQDLH). To decide how to handle these messages - for example, to determine whether to retry or discard them - a dead-letter handling application must look at the reason code contained in the MQDLH.

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;

...
MQMessage message = ... // Message received from the dead-letter queue.
MQDLH dlh = new MQDLH ();

dlh.read (message);

System.out.println ("Reason: " + dlh.getReason ());
```

All header classes also provide a convenience constructor to initialize themselves directly from the message in a single step. So the code in this example could be simplified as follows:

```
import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;

...
MQMessage message = ... // Message received from the dead-letter queue.
MQDLH dlh = new MQDLH (message);

System.out.println ("Reason: " + dlh.getReason ());
```

Reading and removing the MQDLH from a dead-letter message using WebSphere MQ classes for Java

In this example, MQDLH is used to remove the header from a dead-letter message.

A dead-letter handling application will typically resubmit messages that have been rejected if their reason code indicates a transient error. Before resubmitting the message, it must remove the MQDLH header.

This example performs the following steps (see the comments in the example code):

1. The MQHeaderList reads the entire message, and each header encountered in the message becomes an item in the list.
2. Dead-letter messages contain an MQDLH as their first header, so this can be found in the first item of the header list. The MQDLH has already been populated from the message when the MQHeaderList is built, so there is no need to invoke its read method.
3. The reason code is extracted using the getReason() method provided by the MQDLH class.
4. The reason code has been inspected, and indicates that it is appropriate to resubmit the message. The MQDLH is removed using the MQHeaderList remove() method.
5. The MQHeaderList writes its remaining content to a new message object. The new message now contains everything in the original message except the MQDLH and can be written to a queue. The **true** argument to the constructor and to the write method indicates that the message body is to be held within the MQHeaderList, and written out again.
6. The format field in the message descriptor of the new message now contains the value that was previously in the MQDLH format field. The message data matches the numeric encoding and CCSID set in the message descriptor.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQDLH;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ... // Message received from the dead-letter queue.
MQHeaderList list = new MQHeaderList (message, true); // Step 1.
MQDLH dlh = (MQDLH) list.get (0); // Step 2.
int reason = dlh.getReason (); // Step 3.
...
list.remove (dlh); // Step 4.

MQMessage newMessage = new MQMessage ();

list.write (newMessage, true); // Step 5.
newMessage.format = list.getFormat (); // Step 6.

```

Printing the content of a message using WebSphere MQ classes for Java

This example uses MQHeaderList to print out the content of a message, including its headers.

The output contains a view of all the header contents as well as the body of the message. The MQHeaderList class decodes all the headers in one go, whereas the MQHeaderIterator steps through them one at a time under application control. You might use this technique to provide a simple debugging tool when writing Websphere MQ applications.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ... // Message received from a queue.

System.out.println (new MQHeaderList (message, true));

```

This example also prints out the message descriptor fields, using the MQMD class. The copyFrom() method of the com.ibm.mq.headers.MQMD class populates the header object from the message descriptor fields of the MQMessage rather than by reading the message body.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQMD;
import com.ibm.mq.headers.MQHeaderList;
...
MQMessage message = ...
MQMD md = new MQMD ();
...
md.copyFrom (message);
System.out.println (md + "\n" + new MQHeaderList (message, true));

```

Finding a specific type of header in a message using WebSphere MQ classes for Java

This example uses the indexOf(String) method of MQHeaderList to find an MQRFH2 header in a message, if one is present.

```

import com.ibm.mq.MQMessage;
import com.ibm.mq.headers.MQHeaderList;
import com.ibm.mq.headers.MQRFH2;
...
MQMessage message = ...
MQHeaderList list = new MQHeaderList (message);
int index = list.indexOf ("MQRFH2");

if (index >= 0)
{
    MQRFH2 rfh = (MQRFH2) list.get (index);
    ...
}

```

Analyzing an MQRFH2 header using WebSphere MQ classes for Java

This example shows how to access a known field value in a named folder, using the MQRFH2 class.

The MQRFH2 class provides a number of ways to access not only the fields in the fixed part of the structure, but also the XML-encoded folder contents that are carried within the NameValueData field. This

example shows how to access a known field value in a named folder - in this instance, the Rto field in the jms folder, which represents the reply queue name in an MQ JMS message.

```
MQRFH2 rfh = ...  
String value = rfh.getStringFieldValue ("jms", "Rto");
```

To discover the contents of an MQRFH2 (as opposed to requesting specific fields directly), you can use the getFolders method to return a list of MQRFH2.Element, which represents the structure of a folder that could contain fields and other folders. Setting a field or folder to null removes it from the MQRFH2. When you manipulate the NameValueData folder contents in this way, the StrucLength field is automatically updated accordingly.

Reading and writing byte streams other than MQMessage objects using WebSphere MQ classes for Java

These examples use the header classes to parse and manipulate WebSphere MQ header content when the data source is not an MQMessage object.

You can use the header classes to parse and manipulate WebSphere MQ header content even when the data source is something other than an MQMessage object. The MQHeader interface implemented by every header class provides the methods `int read (java.io.DataInput message, int encoding, int characterSet)` and `int write (java.io.DataOutput message, int encoding, int characterSet)`. The `com.ibm.mq.MQMessage` class implements the `java.io.DataInput` and `java.io.DataOutput` interfaces. This means that you can use the two MQHeader methods to read and write MQMessage content, overriding the encoding and CCSID specified in the message descriptor. This is useful for messages that contain a chain of headers in different encodings.

You can also obtain `DataInput` and `DataOutput` objects from other data streams, for example file or socket streams, or byte arrays carried in JMS messages. The `java.io.DataInputStream` classes implement `DataInput` and the `java.io.DataOutputStream` classes implement `DataOutput`. This example reads WebSphere MQ header content from a byte array:

```
import java.io.*;  
import com.ibm.mq.headers.*;  
...  
byte [] bytes = ...  
DataInput in = new DataInputStream (new ByteArrayInputStream (bytes));  
MQHeaderIterator it = new MQHeaderIterator (in, CMQC.MQENC_NATIVE,  
    CMQC.MQCCSI_DEFAULT);
```

The line starting `MQHeaderIterator` could be replaced with

```
MQDLH dlh = new MQDLH (in, CMQC.MQENC_NATIVE, CMQC.MQCCSI_DEFAULT);  
// or any other header type
```

This example writes to a byte array using a `DataOutputStream`:

```
MQHeader header = ... // Could be any header type  
ByteArrayOutputStream out = new ByteArrayOutputStream ();  
  
header.write (new DataOutputStream (out), CMQC.MQENC_NATIVE, CMQC.MQCCSI_DEFAULT);  
byte [] bytes = out.toByteArray ();
```

When you work with streams in this way, be careful to use the correct values for the encoding and characterSet arguments. When reading headers, specify the encoding and CCSID with which the byte content was originally written. When writing headers, specify the encoding and CCSID that you want to produce. The data conversion is performed automatically by the header classes.

Creating classes for new header types using WebSphere MQ classes for Java

You can create Java classes for header types not supplied with WebSphere MQ classes for Java.

To add a Java class representing a new header type that you can use in the same way as any header class supplied with WebSphere MQ classes for Java, you create a class that implements the `MQHeader` interface. The simplest approach is to extend the `com.ibm.mq.headers.impl.Header` class. This example

produces a fully-functional class representing the MQTM header structure. You do not have to add individual getter and setter methods for each field, but it is a useful convenience for users of the header class. The generic `getValue` and `setValue` methods that take a string for the field name will work for all fields defined in the header type. The inherited `read`, `write` and `size` methods will enable instances of the new header type to be read and written and will calculate the header size correctly based upon its field definition. The type definition is created just once, however many instances of this header class are created. To make the new header definition available for decoding using the `MQHeaderIterator` or `MQHeaderList` classes, you would register it using the `MQHeaderRegistry`. Note however that the MQTM header class is in fact already provided in this package and registered in the default registry.

```
import com.ibm.mq.headers.impl.Header;
import com.ibm.mq.headers.impl.HeaderField;
import com.ibm.mq.headers.CMQC;

public class MQTM extends Header {
    final static HeaderType TYPE = new HeaderType ("MQTM");
    final static HeaderField StrucId = TYPE.addMQChar ("StrucId", CMQC.MQTM_STRUC_ID);
    final static HeaderField Version = TYPE.addMQLong ("Version", CMQC.MQTM_VERSION_1);
    final static HeaderField QName = TYPE.addMQChar ("QName", CMQC.MQ_Q_NAME_LENGTH);
    final static HeaderField ProcessName = TYPE.addMQChar ("ProcessName",
        CMQC.MQ_PROCESS_NAME_LENGTH);
    final static HeaderField TriggerData = TYPE.addMQChar ("TriggerData",
        CMQC.MQ_TRIGGER_DATA_LENGTH);
    final static HeaderField ApplType = TYPE.addMQLong ("ApplType");
    final static HeaderField ApplId = TYPE.addMQChar ("ApplId", 256);
    final static HeaderField EnvData = TYPE.addMQChar ("EnvData", 128);
    final static HeaderField UserData = TYPE.addMQChar ("UserData", 128);

    protected MQTM (HeaderType type){
        super (type);
    }
    public String getStrucId () {
        return getStringValue (StrucId);
    }
    public int getVersion () {
        return getIntValue (Version);
    }
    public String getQName () {
        return getStringValue (QName);
    }
    public void setQName (String value) {
        setStringValue (QName, value);
    }
    // ...Add convenience getters and setters for remaining fields in the same way.
}
```

Handling PCF messages with WebSphere MQ classes for Java

Java classes are provided to create and parse PCF-structured messages, and to facilitate sending PCF requests and collecting PCF responses.

Classes `PCFMessage` & `MQCFGR` represent arrays of PCF parameter structures. They provide convenience methods for adding and retrieving PCF parameters.

PCF parameter structures are represented by the classes `MQCFH`, `MQCFIN`, `MQCFIN64`, `MQCFST`, `MQCFBS`, `MQCFIL`, `MQCFIL64`, `MQCFSL`, and `MQCFGR`. These share basic operational interfaces:

- Methods to read and write message content: `read ()`, `write ()`, and `size ()`
- Methods to manipulate parameters: `getValue ()`, `setValue ()`, `getParameter ()` and others
- The enumerator method `.nextParameter ()`, which parses PCF content in an `MQMessage`

The PCF filter parameter is used in inquire commands to provide a filter function. It is encapsulated in the following classes:

- `MQCFIF` - integer filter
- `MQCFSF` - string filter
- `MQCFBF` - byte filter

Two agent classes, PCFAgent and PCFMessageAgent are provided to manage the connection to a Queue Manager, the command server queue, and an associated response queue. PCFMessageAgent extends PCFAgent and should normally be used in preference to it. The PCFMessageAgent class converts the received MQMessages and passes them back to the caller as a PCFMessage array. PCFAgent returns an array of MQMessages, which you have to parse before use.

Handling message properties in WebSphere MQ classes for Java

Function calls to process message handles have no equivalent in WebSphere MQ classes for Java. To set, return, or delete message handle properties, use methods of the MQMessage class.

For general information about message properties, see [“Property names” on page 19](#).

In WebSphere MQ classes for Java access to messages is through the MQMessage class. Message handles are therefore not provided in the Java environment and there is no equivalent to the WebSphere MQ function calls MQCRTMH, MQDLTMH, MQMHBUF, and MQBUFMH

To set message handle properties in the procedural interface, you use the call MQSETMP. In WebSphere MQ classes for Java, use the appropriate method of the MQMessage class:

- setBooleanProperty
- setByteProperty
- setBytesProperty
- setShortProperty
- setIntProperty
- setInt2Property
- setInt4Property
- setInt8Property
- setLongProperty
- setFloatProperty
- setDoubleProperty
- setStringProperty
- setObjectProperty

These are sometimes referred to collectively as the *set*property* methods.

To return the value of message handle properties in the procedural interface, you use the call MQINQMP. In WebSphere MQ classes for Java, use the appropriate method of the MQMessage class:

- getBooleanProperty
- getByteProperty
- getBytesProperty
- getShortProperty
- getIntProperty
- getInt2Property
- getInt4Property
- getInt8Property
- getLongProperty
- getFloatProperty
- getDoubleProperty
- getStringProperty
- getObjectProperty

These are sometimes referred to collectively as the *get*property* methods.

To delete the value of message handle properties in the procedural interface, you use the call MQDLTMP. In WebSphere MQ classes for Java, use the deleteProperty method of the MQMessage class.

Handling errors in WebSphere MQ classes for Java

Handle errors arising from WebSphere MQ classes for Java using Java try and catch blocks.

Methods in the Java interface do not return a completion code and reason code. Instead, they throw an exception whenever the completion code and reason code resulting from a WebSphere MQ call are not both zero. This simplifies the program logic so that you do not have to check the return codes after each call to WebSphere MQ. You can decide at which points in your program you want to deal with the possibility of failure. At these points, you can surround your code with try and catch blocks, as in the following example:

```
try {
    myQueue.put(messageA,putMessageOptionsA);
    myQueue.put(messageB,putMessageOptionsB);
}
catch (MQException ex) {
    // This block of code is only executed if one of
    // the two put methods gave rise to a non-zero
    // completion code or reason code.
    System.out.println("An error occurred during the put operation:" +
        "CC = " + ex.completionCode +
        "RC = " + ex.reasonCode);
    System.out.println("Cause exception:" + ex.getCause() );
}
```

The WebSphere MQ call reason codes reported back in Java exceptions for z/OS are documented in [Reason codes for z/OS](#) and [Reason codes for all other platforms](#).

Exceptions that are thrown while a WebSphere MQ classes for Java application is running are also written to the log. However, an application can call the MQException.logExclude() method to prevent exceptions associated with a specific reason code from being logged. You might want to do this in situations where you expect many exceptions associated with a specific reason code to be thrown, and you do not want the log to be filled with these exceptions. For example, if your application attempts to get a message from a queue each time it iterates around a loop and, for most of these attempts, you expect no suitable message to be on the queue, you might want to prevent exceptions associated with the reason code MQRC_NO_MSG_AVAILABLE from being logged. If an application has previously prevented exceptions associated with a specific reason code from being logged, it can allow these exceptions to be logged again by calling the method MQException.logInclude().

Sometimes the reason code does not convey all details associated with the error. For each exception that is thrown, an application should check the linked exception. The linked exception itself might have another linked exception, and so the linked exceptions form a chain leading back to the original underlying problem. A linked exception is implemented by using the chained exception mechanism of the java.lang.Throwable class, and an application obtains a linked exception by calling the Throwable.getCause() method. From an exception that is an instance of MQException, MQException.getCause() retrieves the underlying instance of com.ibm.mq.jmqi.JmqiException, and getCause from this exception retrieves the underlying java.lang.Exception that caused the error.

By default, the MQException class automatically streams exceptions to System.err, which typically is directed to the console. If you want to stop exceptions appearing on the console, include a line in your application to set MQException.log=null.

Getting and setting attribute values in WebSphere MQ classes for Java

getXXX() and setXXX() methods are provided for many common attributes. Others can be accessed using the generic inquire() and set() methods.

For many of the common attributes, the classes MQManagedObject, MQDestination, MQQueue, MQTopic, MQProcess, and MQQueueManager contain getXXX() and setXXX() methods. These methods allow you to get and set their attribute values. Note that for MQDestination, MQQueue, and MQTopic, the methods work only if you specify the appropriate inquire and set flags when you open the object.

For less common attributes, the MQQueueManager, MQDestination, MQQueue, MQTopic, and MQProcess classes all inherit from a class called MQManagedObject. This class defines the inquire() and set() interfaces.

When you create a new queue manager object by using the *new* operator, it is automatically opened for inquire. When you use the accessProcess() method to access a process object, that object is automatically opened for inquire. When you use the accessQueue() method to access a queue object, that object is *not* automatically opened for either inquire or set operations. This is because adding these options automatically can cause problems with some types of remote queues. To use the inquire, set, getXXX, and setXXX methods on a queue, you must specify the appropriate inquire and set flags in the openOptions parameter of the accessQueue() method. The same is true for destination and topic objects.

The inquire and set methods take three parameters:

- selectors array
- intAttrs array
- charAttrs array

You do not need the SelectorCount, IntAttrCount, and CharAttrLength parameters that are found in MQINQ, because the length of an array in Java is always known. The following example shows how to make an inquiry on a queue:

```
// inquire on a queue
final static int MQIA_DEF_PRIORITY = 6;
final static int MQCA_Q_DESC = 2013;
final static int MQ_Q_DESC_LENGTH = 64;

int[] selectors = new int[2];
int[] intAttrs = new int[1];
byte[] charAttrs = new byte[MQ_Q_DESC_LENGTH]

selectors[0] = MQIA_DEF_PRIORITY;
selectors[1] = MQCA_Q_DESC;

queue.inquire(selectors,intAttrs,charAttrs);

System.out.println("Default Priority = " + intAttrs[0]);
System.out.println("Description : " + new String(charAttrs,0));
```

Multithreaded programs in Java

The Java runtime environment is inherently multithreaded. WebSphere MQ classes for Java allows a queue manager object to be shared by multiple threads but ensures that all access to the target queue manager is synchronized.

Multithreaded programs are hard to avoid in Java . Consider a simple program that connects to a queue manager and opens a queue at startup. The program displays a single button on the screen. When a user clicks that button, the program fetches a message from the queue.

The Java runtime environment is inherently multithreaded. Therefore, your application initialization occurs in one thread, and the code that executes in response to the button press executes in a separate thread (the user interface thread).

With the C based WebSphere MQ MQI client, this would cause a problem, because there are limitations to the sharing of handles by multiple threads. WebSphere MQ classes for Java relaxes this constraint, allowing a queue manager object (and its associated queue, topic and process objects) to be shared by multiple threads.

The implementation of WebSphere MQ classes for Java ensures that, for a particular connection (MQQueueManager object instance), all access to the target WebSphere MQ queue manager is synchronized. A thread that wants to issue a call to a queue manager is blocked until all other calls in progress for that connection are complete. If you require simultaneous access to the same queue manager from multiple threads within your program, create a new MQQueueManager object for each thread that requires concurrent access. (This is equivalent to issuing a separate MQCONN call for each thread.)

Note: Instances of the class `com.ibm.mq.MQGetMessageOptions` must not be shared between threads which are requesting messages concurrently. Instances of this class are updated with data during the corresponding MQGET request, which can result in unexpected consequences when multiple threads are operating concurrently on the same instance of the object.

Using channel exits in WebSphere MQ classes for Java

An overview of how to use channel exits in an application using the WebSphere MQ classes for Java.

The following topics describe how to write a channel exit in Java, how to assign it, and how to pass data to it. They then describe how to use channel exits written in C and how to use a sequence of channel exits.

Your application must have the correct security permission to load the channel exit class.

Creating a channel exit in WebSphere MQ classes for Java

You can provide your own channel exits by defining a Java class that implements an appropriate interface.

To implement an exit, you define a new Java class that implements the appropriate interface. Three exit interfaces are defined in the `com.ibm.mq.exits` package:

- `WMQSendExit`
- `WMQReceiveExit`
- `WMQSecurityExit`

Note: Channel exits are supported for client connections only; they are not supported for bindings connections. You cannot use a Java channel exit outside WebSphere MQ classes for Java, for example if you are using a client application written in C.

Any SSL encryption defined for a connection is performed *after* send and security exits have been invoked. Similarly, decryption is performed *before* receive and security exits are invoked.

The following sample defines a class that implements all three interfaces:

```
public class MyMQExits implements
    WMQSendExit, WMQReceiveExit, WMQSecurityExit {
    // Default constructor
    public MyMQExits(){
    }
    // This method comes from the send exit interface
    public ByteBuffer channelSendExit(
        MQCXP channelExitParms,
        MQCD channelDefinition,
        ByteBuffer agentBuffer)
    {
        // Fill in the body of the send exit here
    }
    // This method comes from the receive exit interface
    public ByteBuffer channelReceiveExit(
        MQCXP channelExitParms,
        MQCD channelDefinition,
        ByteBuffer agentBuffer)
    {
        // Fill in the body of the receive exit here
    }
    // This method comes from the security exit interface
    public ByteBuffer channelSecurityExit(
        MQCXP channelExitParms,
        MQCD channelDefinition,
        ByteBuffer agentBuffer)
    {
        // Fill in the body of the security exit here
    }
}
```

Each exit is passed an MQCXP object and an MQCD object. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

Any exit class you write must have a constructor. This can be either the default constructor or one that takes a string argument. If it takes a string then the user data will be passed into the exit class when it is

created. If the exit class contains both a default constructor and a single argument constructor, the single argument constructor has priority.

For the send and security exits, your exit code must return the data that you want to send to the server. For a receive exit, your exit code must return the modified data that you want WebSphere MQ to interpret.

The simplest possible exit body is:

```
{ return agentBuffer; }
```

Do not close the queue manager from within a channel exit.

Using existing channel exit classes

In versions of WebSphere MQ earlier than 7.0, you would implement these exits using the interfaces MQSendExit, MQReceiveExit, and MQSecurityExit, as in the following example. This method remains valid, but the new method is preferred for improved functionality and performance.

```
public class MyMQExits implements MQSendExit, MQReceiveExit, MQSecurityExit {
    // Default constructor
    public MyMQExits(){
    }
    // This method comes from the send exit
    public byte[] sendExit(MQChannelExit channelExitParms,
                          MQChannelDefinition channelDefParms,
                          byte agentBuffer[])
    {
        // Fill in the body of the send exit here
    }
    // This method comes from the receive exit
    public byte[] receiveExit(MQChannelExit channelExitParms,
                             MQChannelDefinition channelDefParms,
                             byte agentBuffer[])
    {
        // Fill in the body of the receive exit here
    }
    // This method comes from the security exit
    public byte[] securityExit(MQChannelExit channelExitParms,
                              MQChannelDefinition channelDefParms,
                              byte agentBuffer[])
    {
        // Fill in the body of the security exit here
    }
}
```

Assigning a channel exit in IBM WebSphere MQ classes for Java

You can assign a channel exit using IBM WebSphere MQ classes for Java.

There is no direct equivalent to the IBM WebSphere MQ channel in IBM WebSphere MQ classes for Java. Channel exits are assigned to an MQQueueManager. For example, having defined a class that implements the WMQSecurityExit interface, an application can use the security exit in one of four ways:

- By assigning an instance of the class to the MQEnvironment.channelSecurityExit field before creating an MQQueueManager object
- By setting the MQEnvironment.channelSecurityExit field to a string representing the security exit class before creating an MQQueueManager object
- By creating a key/value pair in the properties hashtable passed to MQQueueManager with a key of CMQC.SECURITY_EXIT_PROPERTY
- Using a client channel definition table (CCDT)

Any exit assigned by setting the MQEnvironment.channelSecurityExit field to a string, creating a key/value pair in the properties hashtable, or using a CCDT, must be written with a default constructor. An exit assigned as an instance of a class does not need a default constructor, depending on the application.

An application can use a send or a receive exit in a similar way. For example, the following code fragment shows you how to use the security, send, and receive exits that are implemented in the class `MyMQExits`, which was defined previously, using `MQEnvironment`:

```
MyMQExits myexits = new MyMQExits();
MQEnvironment.channelSecurityExit = myexits;
MQEnvironment.channelSendExit = myexits;
MQEnvironment.channelReceiveExit = myexits;
:
MQQueueManager jupiter = new MQQueueManager("JUPITER");
```

If more than one method is used to assign a channel exit, the order of precedence is as follows:

1. If the URL of a CCDT is passed to the `MQQueueManager`, the contents of the CCDT determine the channel exits to be used and any exit definitions in `MQEnvironment` or the properties hashtable are ignored.
2. If no CCDT URL is passed, exit definitions from `MQEnvironment` and the hashtable are merged
 - If the same exit type is defined in both `MQEnvironment` and the hashtable, the definition in the hashtable is used.
 - If equivalent old and new types of exit are specified (for example the `sendExit` field, which can only be used for the type of exit used in versions of IBM WebSphere MQ earlier than Version 7.0, and the `channelSendExit` field, which can be used for any send exit), the new exit (`channelSendExit`) is used rather than the old exit.

If you have declared a channel exit as a string, you must enable IBM WebSphere MQ to locate the channel exit program. You can do so in various ways, depending on the environment in which the application is running and on how the channel exit programs are packaged.

- For an application that is running in an application server, you must store the files in the directory shown in [Table 88 on page 659](#) or packaged in JAR files referenced by **`exitClasspath`**.
- For an application that is not running in an application server, the following rules apply:
 - If your channel exit classes are packaged in separate JAR files, these JAR files must be included in the **`exitClasspath`**.
 - If your channel exit classes are not packaged in JAR files, the class files can be stored in the directory shown in [Table 88 on page 659](#) or in any directory in the JVM system class path or **`exitClasspath`**.

The **`exitClasspath`** property can be specified in four ways. In order of priority, these ways are as follows:

1. The system property `com.ibm.mq.exitClasspath` (defined on the command line using the `-D` option)
2. The `exitPath` stanza of the `mqclient.ini` file
3. A hashtable entry with the key `CMQC.EXIT_CLASSPATH_PROPERTY`
4. The `MQEnvironment` variable **`exitClasspath`**

Separate multiple paths using the `java.io.File.pathSeparator` character.

Table 88. The directory for channel exit programs	
Platform	Directory
AIX, HP-UX, Linux, and Solaris	<code>/var/mqm/exits</code> (32-bit channel exit programs) <code>/var/mqm/exits64</code> (64-bit channel exit programs)
Windows	<code>install_data_dir\exits</code>
Note: <code>install_data_dir</code> is the directory that you chose for the IBM WebSphere MQ data files during installation. The default directory is <code>C:\Program Files\IBM\WebSphere MQ</code> .	

Passing data to channel exits in WebSphere MQ classes for Java

You can pass data to channel exits and return data from channel exits to your application.

The agentBuffer parameter

For a send exit, the *agentBuffer* parameter contains the data that is about to be sent. For a receive exit or a security exit, the *agentBuffer* parameter contains the data that has just been received. You do not need a length parameter, because the expression `agentBuffer.limit()` indicates the length of the array.

For the send and security exits, your exit code must return the data that you want to send to the server. For a receive exit, your exit code must return the modified data that you want WebSphere MQ to interpret.

The simplest possible exit body is:

```
{ return agentBuffer; }
```

Channel exits are called with a buffer that has a backing array. For best performance, the exit should return a buffer with a backing array.

User data

If an application connects to a queue manager by setting `channelSecurityExit`, `channelSendExit`, or `channelReceiveExit`, 32 bytes of user data can be passed to the appropriate channel exit class when it is called, using the `channelSecurityExitUserData`, `channelSendExitUserData`, or `channelReceiveExitUserData` fields. This user data is available to the channel exit class but is refreshed each time the exit is called. Any changes made to the user data in the channel exit will therefore be lost. If you want to make persistent changes to data in a channel exit, use the MQCXP `exitUserArea`. Data in this field is maintained between invocations of the exit.

If the application sets `securityExit`, `sendExit`, or `receiveExit`, no user data can be passed to these channel exit classes.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, any user data specified in a client connection channel definition is passed to channel exit classes when they are called. For more information about using a client channel definition table, see [“Using a client channel definition table with IBM WebSphere MQ classes for Java” on page 643](#).

Using channel exits not written in Java with WebSphere MQ classes for Java

How to use channel exit programs written in C from a Java application.

In WebSphere MQ Version 7.0, you can specify the name of a channel exit program written in C as a String passed to the `channelSecurityExit`, `channelSendExit`, or `channelReceiveExit` fields in the `MQEnvironment` object or properties Hashtable. However, you cannot use a channel exit written in Java in an application written in another language.

Specify the exit program name in the format `library(function)` and ensure that the location of the exit program is included in the path environment variable.

For information about how to write a channel exit in C, see [“Channel-exit programs for messaging channels” on page 378](#).

Using external exit classes

In versions of WebSphere MQ earlier than Version 7.0, three classes were provided to enable you to use channel exits written in languages other than Java:

- `MQExternalSecurityExit`, which implements the `MQSecurityExit` interface
- `MQExternalSendExit`, which implements the `MQSendExit` interface
- `MQExternalReceiveExit`, which implements the `MQReceiveExit` interface

The use of these classes remains valid but the new method is preferred.

To use a security exit that is not written in Java, an application first had to create an MQExternalSecurityExit object. The application specified, as parameters on the MQExternalSecurityExit constructor, the name of the library containing the security exit, the name of the entry point for the security exit, and the user data to be passed to the security exit when it was called. Channel exit programs that are not written in Java were stored in the directory shown in [Table 88 on page 659](#).

Using a sequence of channel send or receive exits in WebSphere MQ classes for Java

A WebSphere MQ classes for Java application can use a sequence of channel send or receive exits that are run in succession.

To use a sequence of send exits, an application can create either a List or a String containing the send exits. If a List is used, each element of the List can be any of the following:

- An instance of a user defined class that implements the WMQSendExit interface
- An instance of a user defined class that implements the MQSendExit interface (for a send exit written in Java)
- An instance of the MQExternalSendExit class (for a send exit not written in Java)
- An instance of the MQSendExitChain class
- An instance of the String class

A List cannot contain another List.

The application can use a sequence of receive exits in a similar manner.

If a String is used, it must consist of one or more comma-separated exit definitions, each of which can be the name of a Java class, or a C program in the format `library(function)`.

The application then assigns the List or String object to the MQEnvironment.channelSendExit field before creating an MQQueueManager object.

The context of information passed to exits is solely within the domain of the exits. For example, if a Java exit and a C exit are chained, the presence of the Java exit has no effect on the C exit.

Using exit chain classes

In versions of WebSphere MQ earlier than Version 7.0, two classes were provided to allow sequences of exits:

- MQSendExitChain, which implements the MQSendExit interface
- MQReceiveExitChain, which implements the MQReceiveExit interface

The use of these classes remains valid but the new method is preferred. Using the WebSphere MQ Classes for Java interfaces means that your application still has a dependency on `com.ibm.mq.jar`. If the new set of interfaces in the `com.ibm.mq.exits` package are used there is no dependency on `com.ibm.mq.jar`.

To use a sequence of send exits, an application created a list of objects, where each object was one of the following:

- An instance of a user defined class that implements the MQSendExit interface (for a send exit written in Java)
- An instance of the MQExternalSendExit class (for a send exit not written in Java)
- An instance of the MQSendExitChain class

The application created an MQSendExitChain object by passing this list of objects as a parameter on the constructor. The application would then have assigned the MQSendExitChain object to the MQEnvironment.sendExit field before creating an MQQueueManager object.

Channel compression in WebSphere MQ classes for Java

Compressing the data that flows on a channel can improve the performance of the channel and reduce network traffic. IBM WebSphere MQ classes for Java use the compression function built into IBM WebSphere MQ.

Using function supplied with IBM WebSphere MQ, you can compress the data that flows on message channels and MQI channels and, on either type of channel, you can compress header data and message data independently of each other. By default, no data is compressed on a channel. For a full description of channel compression, including how it is implemented in IBM WebSphere MQ, see [Data compression \(COMPMSG\)](#) and [Header compression \(COMPHDR\)](#).

A IBM WebSphere MQ classes for Java application specifies the techniques that can be used for compressing header or message data on a client connection by creating a `java.util.Collection` object. Each compression technique is an `Integer` object in the collection, and the order in which the application adds the compression techniques to the collection is the order in which the compression techniques are negotiated with the queue manager when the client connection starts. The application can then assign the collection to the `hdrCompList` field, for header data, or the `msgCompList` field, for message data, in the `MQEnvironment` class. When the application is ready, it can start the client connection by creating an `MQueueManager` object.

The following code fragments illustrate the approach described. The first code fragment shows you how to implement header data compression:

```
Collection headerComp = new Vector();
headerComp.add(new Integer(CMQXC.MQCOMPRESS_SYSTEM));
:
MQEnvironment.hdrCompList = headerComp;
:
MQueueManager qMgr = new MQueueManager(QM);
```

The second code fragment shows you how to implement message data compression:

```
Collection msgComp = new Vector();
msgComp.add(new Integer(CMQXC.MQCOMPRESS_RLE));
msgComp.add(new Integer(CMQXC.MQCOMPRESS_ZLIBHIGH));
:
MQEnvironment.msgCompList = msgComp;
:
MQueueManager qMgr = new MQueueManager(QM);
```

In the second example, the compression techniques are negotiated in the order RLE, then ZLIBHIGH, when the client connection starts. The compression technique that is selected cannot be changed during the lifetime of the `MQueueManager` object.

The compression techniques for header and message data that are supported by both the client and the queue manager on a client connection are passed to a channel exit as collections in the `hdrCompList` and `msgCompList` fields of an `MQChannelDefinition` object. The actual techniques that are currently being used for compressing header and message data on a client connection are passed to a channel exit in the `CurHdrCompression` and `CurMsgCompression` fields of an `MQChannelExit` object.

If compression is used on a client connection, the data is compressed before any channel send exits are processed and extracted after any channel receive exits are processed. The data passed to send and receive exits is therefore in a compressed state.

For more information about specifying compression techniques, and about which compression techniques are available, see [Class `com.ibm.mq.MQEnvironment`](#) and [Interface `com.ibm.mq.MQC`](#).

Sharing a TCP/IP connection in IBM WebSphere MQ classes for Java

Multiple instances of an MQI channel can be made to share a single TCP/IP connection.

In IBM WebSphere MQ classes for Java, you use the `MQEnvironment.sharingConversations` variable to control the number of conversations that can share a single TCP/IP connection.

The SHARECNV attribute is a best effort approach to connection sharing. Therefore when a SHARECNV value greater than 0 is used with the IBM WebSphere MQ classes for Java it is not guaranteed that a new connection request will always share an already established connection.

Connection pooling in WebSphere MQ classes for Java

WebSphere MQ classes for Java allows spare connections to be pooled for reuse.

WebSphere MQ classes for Java provides additional support for applications that deal with multiple connections to WebSphere MQ queue managers. When a connection is no longer required, instead of destroying it, it can be pooled and later reused. This can provide a substantial performance enhancement for applications and middleware that connect serially to arbitrary queue managers.

WebSphere MQ provides a default connection pool. Applications can activate or deactivate this connection pool by registering and deregistering tokens through the MQEnvironment class. If the pool is active when WebSphere MQ classes for Java constructs an MQQueueManager object, it searches this default pool and reuses any suitable connection. When an MQQueueManager.disconnect() call occurs, the underlying connection is returned to the pool.

Alternatively, applications can construct an MQSimpleConnectionManager connection pool for a particular use. Then, the application can either specify that pool during construction of an MQQueueManager object, or pass that pool to MQEnvironment for use as the default connection pool.

To prevent connections from using too much resource, you can limit the total number of connections that an MQSimpleConnectionManager object can handle, and you can limit the size of the connection pool. Setting limits is useful if there are conflicting demands for connections within a JVM.

By default, the getMaxConnections() method returns the value zero, which means that there is no limit to the number of connections that the MQSimpleConnectionManager object can handle. You can set a limit by using the setMaxConnections() method. If you set a limit and the limit is reached, a request for a further connection might cause an MQException to be thrown, with a reason code of MQRC_MAX_CONNS_LIMIT_REACHED.

Controlling the default connection pool in WebSphere MQ classes for Java

This example shows how to use the default connection pool.

Consider the following example application, MQApp1:

```
import com.ibm.mq.*;
public class MQApp1
{
    public static void main(String[] args) throws MQException
    {
        for (int i=0; i<args.length; i++) {
            MQQueueManager qmgr=new MQQueueManager(args[i]);
            :
            : (do something with qmgr)
            :
            qmgr.disconnect();
        }
    }
}
```

MQApp1 takes a list of local queue managers from the command line, connects to each in turn, and performs some operation. However, when the command line lists the same queue manager many times, it is more efficient to connect only once, and to reuse that connection many times.

WebSphere MQ classes for Java provides a default connection pool that you can use to do this. To enable the pool, use one of the MQEnvironment.addConnectionPoolToken() methods. To disable the pool, use MQEnvironment.removeConnectionPoolToken().

The following example application, MQApp2, is functionally identical to MQApp1, but connects only once to each queue manager.

```
import com.ibm.mq.*;
```

```

public class MQApp2
{
    public static void main(String[] args) throws MQException
    {
        MQPoolToken token=MQEnvironment.addConnectionPoolToken();

        for (int i=0; i<args.length; i++) {
            MQQueueManager qmgr=new MQQueueManager(args[i]);
            :
            : (do something with qmgr)
            :
            qmgr.disconnect();
        }

        MQEnvironment.removeConnectionPoolToken(token);
    }
}

```

The first bold line activates the default connection pool by registering an MQPoolToken object with MQEnvironment.

The MQQueueManager constructor now searches this pool for an appropriate connection and only creates a connection to the queue manager if it cannot find an existing one. The qmgr.disconnect() call returns the connection to the pool for later reuse. These API calls are the same as the sample application MQApp1.

The second highlighted line deactivates the default connection pool, which destroys any queue manager connections stored in the pool. This is important because otherwise the application would terminate with a number of live queue manager connections in the pool. This situation could cause errors that would appear in the queue manager logs.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, the MQQueueManager constructor first searches the table for a suitable client connection channel definition. If one is found, the constructor searches the default connection pool for a connection that can be used for the channel. If the constructor cannot find a suitable connection in the pool, it then searches the client channel definition table for the next suitable client connection channel definition, and proceeds as described previously. If the constructor completes its search of the client channel definition table and fails to find any suitable connection in the pool, the constructor starts a second search of the table. During this search, the constructor tries to create a new connection for each suitable client connection channel definition in turn, and uses the first connection that it manages to create.

The default connection pool stores a maximum of ten unused connections, and keeps unused connections active for a maximum of five minutes. The application can alter this (for details, see [“Supplying a different connection pool in WebSphere MQ classes for Java” on page 665](#)).

Instead of using MQEnvironment to supply an MQPoolToken, the application can construct its own:

```

MQPoolToken token=new MQPoolToken();
MQEnvironment.addConnectionPoolToken(token);

```

Some applications or middleware vendors provide subclasses of MQPoolToken in order to pass information to a custom connection pool. They can be constructed and passed to addConnectionPoolToken() in this way so that extra information can be passed to the connection pool.

The default connection pool and multiple components in WebSphere MQ classes for Java

This example shows how to add or remove MQPoolTokens from a static set of registered MQPoolToken objects.

MQEnvironment holds a static set of registered MQPoolToken objects. To add or remove MQPoolTokens from this set, use the following methods:

- MQEnvironment.addConnectionPoolToken()
- MQEnvironment.removeConnectionPoolToken()

An application might consist of many components that exist independently and perform work using a queue manager. In such an application, each component should add an MQPoolToken to the MQEnvironment set for its lifetime.

For example, the example application MQApp3 creates ten threads and starts each one. Each thread registers its own MQPoolToken, waits for a length of time, then connects to the queue manager. After the thread disconnects, it removes its own MQPoolToken.

The default connection pool remains active while there is at least one token in the set of MQPoolTokens, so it will remain active for the duration of this application. The application does not need to keep a master object in overall control of the threads.

```
import com.ibm.mq.*;
public class MQApp3
{
    public static void main(String[] args)
    {
        for (int i=0; i<10; i++) {
            MQApp3_Thread thread=new MQApp3_Thread(i*60000);
            thread.start();
        }
    }
}

class MQApp3_Thread extends Thread
{
    long time;

    public MQApp3_Thread(long time)
    {
        this.time=time;
    }

    public synchronized void run()
    {
        MQPoolToken token=MQEnvironment.addConnectionPoolToken();
        try {
            wait(time);
            MQQueueManager qmgr=new MQQueueManager("my.qmgr.1");
            :
            : (do something with qmgr)
            :
            qmgr.disconnect();
        }
        catch (MQException mqe) {System.err.println("Error occurred!");}
        catch (InterruptedException ie) {}

        MQEnvironment.removeConnectionPoolToken(token);
    }
}
```

Supplying a different connection pool in WebSphere MQ classes for Java

This example shows how to use the class **com.ibm.mq.MQSimpleConnectionManager** to supply a different connection pool.

This class provides basic facilities for connection pooling, and applications can use this class to customize the behavior of the pool.

Once it is instantiated, an MQSimpleConnectionManager can be specified on the MQQueueManager constructor. The MQSimpleConnectionManager then manages the connection that underlies the constructed MQQueueManager. If the MQSimpleConnectionManager contains a suitable pooled connection, that connection is reused and returned to the MQSimpleConnectionManager after an MQQueueManager.disconnect() call.

The following code fragment demonstrates this behavior:

```
MQSimpleConnectionManager myConnMan=new MQSimpleConnectionManager();
myConnMan.setActive(MQSimpleConnectionManager.MODE_ACTIVE);
MQQueueManager qmgr=new MQQueueManager("my.qmgr.1", myConnMan);
:
: (do something with qmgr)
```

```

:
qmgr.disconnect();

MQQueueManager qmgr2=new MQQueueManager("my.qmgr.1", myConnMan);
:
: (do something with qmgr2)
:
qmgr2.disconnect();
myConnMan.setActive(MQSimpleConnectionManager.MODE_INACTIVE);

```

The connection that is forged during the first `MQQueueManager` constructor is stored in `myConnMan` after the `qmgr.disconnect()` call. The connection is then reused during the second call to the `MQQueueManager` constructor.

The second line enables the `MQSimpleConnectionManager`. The last line disables `MQSimpleConnectionManager`, destroying any connections held in the pool. An `MQSimpleConnectionManager` is, by default, in `MODE_AUTO`, which is described later in this section.

An `MQSimpleConnectionManager` allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if it has not been used for five minutes, or if there are more than ten unused connections in the pool. You can alter these values by calling `MQSimpleConnectionManager.setTimeout()`.

You can also set up an `MQSimpleConnectionManager` for use as the default connection pool, to be used when no Connection Manager is supplied on the `MQQueueManager` constructor.

The following application demonstrates this:

```

import com.ibm.mq.*;
public class MQApp4
{
    public static void main(String []args)
    {
        MQSimpleConnectionManager myConnMan=new MQSimpleConnectionManager();
        myConnMan.setActive(MQSimpleConnectionManager.MODE_AUTO);
        myConnMan.setTimeout(3600000);
        myConnMan.setMaxConnections(75);
        myConnMan.setMaxUnusedConnections(50);
        MQEnvironment.setDefaultConnectionManager(myConnMan);
        MQApp3.main(args);
    }
}

```

The bold lines create and configure an `MQSimpleConnectionManager` object. The configuration does the following:

- Ends connections that are not used for an hour
- Limits the number of connections managed by `myConnMan` to 75
- Limits the number of unused connections in the pool to 50
- Sets `MODE_AUTO`, which is the default. This means that the pool is active only if it is the default connection manager, and there is at least one token in the set of `MQPoolTokens` held by `MQEnvironment`.

The new `MQSimpleConnectionManager` is then set as the default connection manager.

In the last line, the application calls `MQApp3.main()`. This runs a number of threads, where each thread uses WebSphere MQ independently. These threads use `myConnMan` when they forge connections.

Supplying your own ConnectionManager for WebSphere MQ classes for Java

WebSphere MQ classes for Java provides a partial implementation of the Java EE Connector Architecture, allowing implementations of `javax.resource.spi.ConnectionManager` to be used.

Applications and middleware providers can provide alternative implementations of connection pools. WebSphere MQ classes for Java provides a partial implementation of the Java EE Connector Architecture. Implementations of **`javax.resource.spi.ConnectionManager`** can either be used as the default Connection Manager or be specified on the `MQQueueManager` constructor.

WebSphere MQ classes for Java complies with the Connection Management contract of the Java EE Connector Architecture. Read this section in conjunction with the Connection Management contract of the Java EE Connector Architecture (see Sun's Java Web site at <https://java.sun.com>).

The ConnectionManager interface defines only one method:

```
package javax.resource.spi;
public interface ConnectionManager {
    Object allocateConnection(ManagedConnectionFactory mcf,
                             ConnectionRequestInfo cxRequestInfo);
}
```

The MQQueueManager constructor calls allocateConnection on the appropriate ConnectionManager. It passes appropriate implementations of ManagedConnectionFactory and ConnectionRequestInfo as parameters to describe the connection required.

The ConnectionManager searches its pool for a javax.resource.spi.ManagedConnection object that has been created with identical ManagedConnectionFactory and ConnectionRequestInfo objects. If the ConnectionManager finds any suitable ManagedConnection objects, it creates a java.util.Set that contains the candidate ManagedConnections. Then, the ConnectionManager calls the following:

```
ManagedConnection mc=mcf.matchManagedConnections(connectionSet, subject,
cxRequestInfo);
```

The WebSphere MQ implementation of ManagedConnectionFactory ignores the subject parameter. This method selects and returns a suitable ManagedConnection from the set, or returns null if it does not find a suitable ManagedConnection. If there is not a suitable ManagedConnection in the pool, the ConnectionManager can create one by using:

```
ManagedConnection mc=mcf.createManagedConnection(subject, cxRequestInfo);
```

Again, the subject parameter is ignored. This method connects to a WebSphere MQ queue manager and returns an implementation of javax.resource.spi.ManagedConnection that represents the newly-forged connection. Once the ConnectionManager has obtained a ManagedConnection (either from the pool or freshly created), it creates a connection handle using:

```
Object handle=mc.getConnection(subject, cxRequestInfo);
```

This connection handle can be returned from allocateConnection().

A ConnectionManager must register an interest in the ManagedConnection through:

```
mc.addConnectionEventListener()
```

The ConnectionEventListener is notified if a severe error occurs on the connection, or when MQQueueManager.disconnect() is called. When MQQueueManager.disconnect() is called, the ConnectionEventListener can do either of the following:

- Reset the ManagedConnection using the mc.cleanup() call, then return the ManagedConnection to the pool
- Destroy the ManagedConnection using the mc.destroy() call

If the ConnectionManager is the default ConnectionManager, it can also register an interest in the state of the MQEnvironment-managed set of MQPoolTokens. To do so, first construct an MQPoolServices object, then register an MQPoolServicesEventListener object with the MQPoolServices object:

```
MQPoolServices mqps=new MQPoolServices();
mqps.addMQPoolServicesEventListener(listener);
```

The listener is notified when an MQPoolToken is added or removed from the set, or when the default ConnectionManager changes. The MQPoolServices object also provides a way to query the current size of the set of MQPoolTokens.

JTA/JDBC coordination using WebSphere MQ classes for Java

WebSphere MQ classes for Java supports the MQQueueManager.begin() method, which allows WebSphere MQ to act as a coordinator for a database which provides a JDBC type 2 or JDBC type 4 compliant driver.

This support is not available on all platforms. To check which platforms support JDBC coordination, see <https://www.ibm.com/software/integration/wmq/requirements/>.

To use the XA-JTA support, you must use the special JTA switch library. The method for using this library varies depending on whether you are using Windows or one of the other platforms.

Configuring JTA/JDBC coordination on Windows

The XA library is supplied as a DLL with a name of the format `jdbcxxx.dll`.

V7.5.0.7 The supplied `jdbcora12.dll` provides compatibility with Oracle 12C, for an IBM WebSphere MQ Windows server installation.

On Windows systems, the XA library is supplied as a complete DLL. The name of this DLL is `jdbcxxx.dll` where `xxx` indicates the database for which the switch library has been compiled. This library is in the `java\lib\jdbc` or `java\lib64\jdbc` directory of your IBM WebSphere MQ classes for Java installation. You must declare the XA library, also described as the switch load file, to the queue manager. Use the IBM WebSphere MQ Explorer. Specify the details of the switch load file in the queue manager properties panel, under XA resource manager. You must only give the name of the library. For example:

For a Db2 database set the SwitchFile field to: `dbcd2`

For an Oracle database set the SwitchFile field to: `jdbcora`

Configuring JTA/JDBC coordination on platforms other than Windows

Object files are supplied. Link the appropriate one using the supplied makefile, and declare it to the queue manager using the configuration file.

For each database management system, WebSphere MQ provides two object files. You must link one object file to create a 32-bit switch library, and link the other object file to create a 64-bit switch library. For DB2, the name of each object file is `jdbcd2.o` and, for Oracle, the name of each object file is `jdbcora.o`.

You must link each object file using the appropriate makefile supplied with WebSphere MQ. A switch library requires other libraries, which might be stored in different locations on different systems. However, a switch library cannot use the library path environment variable to locate these libraries because the switch library is loaded by the queue manager, which runs in a setuid environment. The supplied makefile therefore ensures that a switch library contains the fully qualified path names of these libraries.

To create a switch library, enter a **make** command with the following format. To create a 32-bit switch library, enter the command in the `/java/lib/jdbc` directory of your WebSphere MQ installation. To create a 64-bit switch library, enter the command in the `/java/lib64/jdbc` directory.

```
make DBMS
```

where *DBMS* is the database management system for which you are creating the switch library. The valid values are `db2` for DB2 and `oracle` for Oracle.

Here is an example of a **make** command:

```
make db2
```

Note the following points:

- To run 32-bit applications, you must create both a 32-bit and a 64-bit switch library for each database management system that you are using. To run 64-bit applications, you need create only a 64-bit switch library. For DB2, the name of each switch library is jdbcdb2 and, for Oracle, the name of each switch library is jdbcora. The makefiles ensure that 32-bit and 64-bit switch libraries are stored in different WebSphere MQ directories. A 32-bit switch library is stored in the /java/lib/jdbc directory, and a 64-bit switch library is stored in the /java/lib64/jdbc directory.
- Because you can install Oracle anywhere on a system, the makefiles use the ORACLE_HOME environment variable to locate where Oracle is installed.

After you have created the switch libraries for DB2, Oracle, or both, you must declare them to your queue manager. If the queue manager configuration file (qm.ini) already contains XAResourceManager stanzas for DB2 or Oracle databases, you must replace the SwitchFile entry in each stanza by one of the following:

For a DB2 database

```
SwitchFile=jdbcdb2
```

For an Oracle database

```
SwitchFile=jdbcora
```

Do not specify the fully qualified path name of either the 32-bit or 64-bit switch library. Specify only the name of the library.

If the queue manager configuration file does not already contain XAResourceManager stanzas for DB2 or Oracle databases, or if you want to add additional XAResourceManager stanzas, see [Administering](#) for information about how to construct an XAResourceManager stanza. However, each SwitchFile entry in a new XAResourceManager stanza must be exactly as described previously for a DB2 or Oracle database. You must also include the entry ThreadOfControl=PROCESS.

After you have updated the queue manager configuration file, and made sure that all appropriate database environment variables have been set, you can restart the queue manager.

Using JTA/JDBC coordination

Code your API calls as in the supplied example.

The basic sequence of API calls for a user application is:

```
qMgr = new MQQueueManager("QM1")
Connection con = qMgr.getJDBCConnection( xads );
qMgr.begin()

< Perform MQ and DB operations to be grouped in a unit of work >

qMgr.commit() or qMgr.backout();
con.close()
qMgr.disconnect()
```

xads in the getJDBCConnection call is a database-specific implementation of the XADataSource interface, which defines the details of the database to connect to. See the documentation for your database to determine how to create an appropriate XADataSource object to pass into getJDBCConnection.

You must also update your class path with the appropriate database-specific jar files for performing JDBC work.

If you must connect to multiple databases, you must call getJDBCConnection several times to perform the transaction across several different connections.

There are two forms of the getJDBCConnection, reflecting the two forms of XADataSource.getXAConnection:

```
public java.sql.Connection getJDBCConnection(javax.sql.XADataSource xads)
    throws MQException, SQLException, Exception
```

```
public java.sql.Connection getJDBCConnection(XADataSource dataSource,
                                             String userid, String password)
    throws MQException, SQLException, Exception
```

These methods declare Exception in their throws clauses to avoid problems with the JVM verifier for customers who are not using the JTA functions. The actual exception thrown is javax.transaction.xa.XAException which requires the jta.jar file to be added to the class path for programs that did not previously require it.

To use the JTA/JDBC support, you must include the following statement in your application:

```
MQEnvironment.properties.put(CMQC.THREAD_AFFINITY_PROPERTY, new Boolean(true));
```

Known problems and limitations with JTA/JDBC coordination

There are certain problems and limitations of JTA/JDBC support, some depending on the database management system in use.

Because this support makes calls to JDBC drivers, the implementation of those JDBC drivers can have a significant effect on the system behavior. In particular, tested JDBC drivers behave differently when the database is shut down while an application is running. **Always** avoid abruptly shutting down a database while there are applications holding open connections to it.

Multiple XAResourceManager stanzas

The use of more than one XAResourceManager stanza in a queue manager configuration file, qm.ini, is not supported. Any XAResourceManager stanza other than the first is ignored.

DB2

Sometimes DB2 returns a SQL0805N error. This problem can be resolved with the following CLP command:

```
DB2 bind @db2cli.lst blocking all grant public
```

Refer to the DB2 documentation for more information.

The XAResourceManager stanza must be configured to use ThreadOfControl=PROCESS. For DB2 version 8.1 and higher this does not match the default thread of control setting for DB2, so toc=p must be specified in the XA Open String. An example XAResourceManager stanza for DB2 with JTA/JDBC coordination is as follows:

```
XAResourceManager:
  Name=jdbcdb2
  SwitchFile=jdbcdb2
  XAOpenString=uid=userid,db=dbalias,pwd=password,toc=p
  ThreadOfControl=PROCESS
```

This does not prevent the Java applications that use JTA/JDBC coordination from being multithreaded themselves.

Oracle

Calling the JDBC Connection.close() method after MQQueueManager.disconnect() generates an SQLException. Either call Connection.close() before MQQueueManager.disconnect(), or omit the call to Connection.close().

Secure Sockets Layer (SSL) support in WebSphere MQ classes for Java

WebSphere MQ classes for Java client applications support Secure Sockets Layer (SSL) encryption. You require a JSSE provider to use SSL encryption.

WebSphere MQ classes for Java client applications using TRANSPORT(CLIENT) support Secure Sockets Layer (SSL) encryption. SSL provides communication encryption, authentication, and message integrity. It is typically used to secure communications between any two peers on the Internet or within an intranet.

WebSphere MQ classes for Java uses Java Secure Socket Extension (JSSE) to handle SSL encryption, and so requires a JSSE provider. JSE v1.4 JVMs have a JSSE provider built in. Details of how to manage

and store certificates can vary from provider to provider. For information about this, refer to your JSSE provider's documentation.

This section assumes that your JSSE provider is correctly installed and configured, and that suitable certificates have been installed and made available to your JSSE provider.

If your WebSphere MQ classes for Java client application uses a client channel definition table (CCDT) to connect to a queue manager, see [“Using a client channel definition table with IBM WebSphere MQ classes for Java”](#) on page 643.

Enabling SSL in IBM WebSphere MQ classes for Java

To enable SSL, you specify a CipherSuite. There are two ways of specifying a CipherSuite.

SSL is supported only for client connections. To enable SSL, you must specify the CipherSuite to use when communicating with the queue manager, and this CipherSuite must match the CipherSpec set on the target channel. Additionally, the named CipherSuite must be supported by your JSSE provider. However, CipherSuites are distinct from CipherSpecs and so have different names. [“SSL CipherSpecs and CipherSuites in WebSphere MQ classes for Java”](#) on page 675 contains a table mapping the CipherSpecs supported by IBM WebSphere MQ to their equivalent CipherSuites as known to JSSE.

To enable SSL, specify the CipherSuite using the `sslCipherSuite` static member variable of `MQEnvironment`. The following example attaches to an `SVRCONN` channel named `SECURE.SVRCONN.CHANNEL`, which has been set up to require SSL with a CipherSpec of `RC4_MD5_EXPORT`:

```
MQEnvironment.hostname      = "your_hostname";
MQEnvironment.channel       = "SECURE.SVRCONN.CHANNEL";
MQEnvironment.sslCipherSuite = "SSL_RSA_EXPORT_WITH_RC4_40_MD5";
MQQueueManager qmgr = new MQQueueManager("your_q_manager");
```

Although the channel has a CipherSpec of `RC4_MD5_EXPORT`, the Java application must specify a CipherSuite of `SSL_RSA_EXPORT_WITH_RC4_40_MD5`. See [“SSL CipherSpecs and CipherSuites in WebSphere MQ classes for Java”](#) on page 675 for a list of mappings between CipherSpecs and CipherSuites.

An application can also specify a CipherSuite by setting the environment property `CMQC.SSL_CIPHER_SUITE_PROPERTY`.

Alternatively, use the Client Channel Definition Table (CCDT). For more information, see [“Using a client channel definition table with IBM WebSphere MQ classes for Java”](#) on page 643

If you require a client connection to use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS), an application can set the `sslFipsRequired` field in the `MQEnvironment` class to `true`. Alternatively, the application can set the environment property `CMQC.SSL_FIPS_REQUIRED_PROPERTY`. The default value is `false`, which means that a client connection can use any CipherSuite that is supported by IBM WebSphere MQ.

If an application uses more than one client connection, the value of the `sslFipsRequired` field that is used when the application creates the first client connection determines the value that is used when the application creates any subsequent client connection. Therefore when the application creates a subsequent client connection, the value of the `sslFipsRequired` field is ignored. You must restart the application if you want to use a different value for the `sslFipsRequired` field.

To connect successfully using SSL, the JSSE truststore must be set up with certificate authority root certificates from which the certificate presented by the queue manager can be authenticated. Similarly, if `SSLClientAuth` on the `SVRCONN` channel has been set to `MQSSL_CLIENT_AUTH_REQUIRED`, the JSSE keystore must contain an identifying certificate that is trusted by the queue manager.

Related reference

[Federal Information Processing Standards \(FIPS\) for UNIX, Linux and Windows](#)

Using the distinguished name of the queue manager in IBM WebSphere MQ classes for Java

The queue manager identifies itself using an SSL certificate, which contains a distinguished name (DN). A IBM WebSphere MQ classes for Java client application can use this DN to ensure that it is communicating with the correct queue manager.

A DN pattern is specified using the `sslPeerName` variable of `MQEnvironment`. For example, setting:

```
MQEnvironment.sslPeerName = "CN=QMGR.*, OU=IBM, OU=WEBSphere";
```

allows the connection to succeed only if the queue manager presents a certificate with a Common Name beginning QMGR., and at least two Organizational Unit names, the first of which must be IBM and the second WebSphere.

If `sslPeerName` is set, connections succeed only if it is set to a valid pattern and the queue manager presents a matching certificate.

An application can also specify the distinguished name of the queue manager by setting the environment property `CMQC.SSL_PEER_NAME_PROPERTY`. For more information about distinguished names, see [Distinguished names](#).

Using certificate revocation lists in IBM WebSphere MQ classes for Java

Specify the certificate revocation lists to use through the `java.security.cert.CertStore` class. IBM WebSphere MQ classes for Java then checks certificates against the specified CRL.

A certificate revocation list (CRL) is a set of certificates that have been revoked, either by the issuing certificate authority or by the local organization. CRLs are typically hosted on LDAP servers. With Java 2 v1.4, a CRL server can be specified at connect-time and the certificate presented by the queue manager is checked against the CRL before the connection is allowed. For more information about certificate revocation lists and IBM WebSphere MQ, see [Working with Certificate Revocation Lists and Authority Revocation Lists](#) and [Accessing CRLs and ARLs with WebSphere MQ classes for Java and WebSphere MQ classes for JMS](#).

Note: To use a `CertStore` successfully with a CRL hosted on an LDAP server, make sure that your Java Software Development Kit (SDK) is compatible with the CRL. Some SDKs require that the CRL conforms to RFC 2587, which defines a schema for LDAP v2. Most LDAP v3 servers use RFC 2256 instead.

The CRLs to use are specified through the `java.security.cert.CertStore` class. Refer to documentation on this class for full details of how to obtain instances of `CertStore`. To create a `CertStore` based on an LDAP server, first create an `LDAPCertStoreParameters` instance, initialized with the server and port settings to use. For example:

```
import java.security.cert.*;
CertStoreParameters csp = new LDAPCertStoreParameters("crl_server", 389);
```

Having created a `CertStoreParameters` instance, use the static constructor on `CertStore` to create a `CertStore` of type LDAP:

```
CertStore cs = CertStore.getInstance("LDAP", csp);
```

Other `CertStore` types (for example, `Collection`) are also supported. Commonly there are several CRL servers set up with identical CRL information to give redundancy. When you have a `CertStore` object for each of these CRL servers, place them all in a suitable `Collection`. The following example shows the `CertStore` objects placed in an `ArrayList`:

```
import java.util.ArrayList;
Collection crls = new ArrayList();
crls.add(cs);
```


This Collection can be set into the MQEnvironment static variable, sslCertStores, before connecting to enable CRL checking:

```
MQEnvironment.sslCertStores = crls;
```

The certificate presented by the queue manager when a connection is being set up is validated as follows:

1. The first CertStore object in the Collection identified by sslCertStores is used to identify a CRL server.
2. An attempt is made to contact the CRL server.
3. If the attempt is successful, the server is searched for a match for the certificate.
 - a. If the certificate is found to be revoked, the search process is over and the connection request fails with reason code MQRC_SSL_CERTIFICATE_REVOKED.
 - b. If the certificate is not found, the search process is over and the connection is allowed to proceed.
4. If the attempt to contact the server is unsuccessful, the next CertStore object is used to identify a CRL server and the process repeats from step 2.

If this was the last CertStore in the Collection, or if the Collection contains no CertStore objects, the search process failed, and the connection request fails with reason code MQRC_SSL_CERT_STORE_ERROR.

The Collection object determines the order in which CertStores are used.

The Collection of CertStores can also be set using the CMQC.SSL_CERT_STORE_PROPERTY. As a convenience, this property also allows a single CertStore to be specified without being a member of a Collection.

If sslCertStores is set to null, no CRL checking is performed. This property is ignored if sslCipherSuite is not set.

Renegotiating the secret key in WebSphere MQ classes for Java

A WebSphere MQ classes for Java client application can control when the secret key that is used for encryption on a client connection is renegotiated, in terms of the total number of bytes sent and received.

The application can do this in either of the following ways: If the application uses more than one of these ways, the usual precedence rules apply.

- By setting the sslResetCount field in the MQEnvironment class.
- By setting the environment property MQC.SSL_RESET_COUNT_PROPERTY in a Hashtable object. The application then assigns the hashtable to the properties field in the MQEnvironment class, or passes the hashtable to an MQQueueManager object on its constructor.

The value of the sslResetCount field or environment property MQC.SSL_RESET_COUNT_PROPERTY represents the total number of bytes sent and received by the WebSphere MQ classes for Java client code before the secret key is renegotiated. The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by the WebSphere MQ classes for Java client.

If the reset count is zero, which is the default value, the secret key is never renegotiated. The reset count is ignored if no CipherSuite is specified.

Supplying a customized SSLSocketFactory in IBM WebSphere MQ classes for Java

If you use a customized JSSE Socket Factory, set the MQEnvironment.sslSocketFactory to the customized factory object. Details vary between different JSSE implementations.

Different JSSE implementations can provide different features. For example, a specialized JSSE implementation might allow configuration of a particular model of encryption hardware. Additionally, some JSSE providers allow customization of keystores and truststores by program, or allow the choice of identity certificate from the keystore to be altered. In JSSE, all these customizations are abstracted into a factory class, javax.net.ssl.SSLSocketFactory.

See your JSSE documentation for details of how to create a customized SSLSocketFactory implementation. The details vary from provider to provider, but a typical sequence of steps might be:

1. Create an SSLContext object using a static method on SSLContext
2. Initialize this SSLContext with appropriate KeyManager and TrustManager implementations (created from their own factory classes)
3. Create an SSLSocketFactory from the SSLContext

When you have an SSLSocketFactory object, set the MQEnvironment.sslSocketFactory to the customized factory object. For example:

```
javax.net.ssl.SSLSocketFactory sf = sslContext.getSocketFactory();
MQEnvironment.sslSocketFactory = sf;
```

IBM WebSphere MQ classes for Java use this SSLSocketFactory to connect to the IBM WebSphere MQ queue manager. This property can also be set using the CMQC.SSL_SOCKET_FACTORY_PROPERTY. If sslSocketFactory is set to null, the default SSLSocketFactory of the JVM is used. This property is ignored if sslCipherSuite is not set.

When you use custom SSLSocketFactories, consider the effect of TCP/IP connection sharing. If connection sharing is possible then a new socket is not requested of the SSLSocketFactory supplied, even if the socket produced would be different in some way in the context of a subsequent connection request. For example, if a different client certificate is to be presented on a subsequent connection, then connection sharing must not be allowed.

Making changes to the JSSE keystore or truststore in WebSphere MQ classes for Java

If you change the JSSE keystore or truststore, you must perform certain actions for the changes to take effect.

If you change the contents of the JSSE keystore or truststore, or change the location of the keystore or truststore file, WebSphere MQ classes for Java applications that are running at the time do not automatically pick up the changes. For the changes to take effect, the following actions must be performed:

- The applications must close all their connections, and destroy any unused connections in connection pools.
- If your JSSE provider caches information from the keystore and truststore, this information must be refreshed.

After these actions have been performed, the applications can then re-create their connections.

Depending on how you design your applications, and on the function provided by your JSSE provider, it might be possible to perform these actions without stopping and restarting your applications. However, stopping and restarting the applications might be the simplest solution.

Error handling when using SSL with WebSphere MQ classes for Java

A number of reason codes can be issued by WebSphere MQ classes for Java when connecting to a queue manager using SSL.

These are explained in the following list:

MQRC_SSL_NOT_ALLOWED

The sslCipherSuite property was set, but bindings connect was used. Only client connect supports SSL.

MQRC_JSSE_ERROR

The JSSE provider reported an error that could not be handled by WebSphere MQ. This could be caused by a configuration problem with JSSE, or because the certificate presented by the queue manager could not be validated. The exception produced by JSSE can be retrieved using the getCause() method on MQException.

MQRC_SSL_INITIALIZATION_ERROR

An MQCONN or MQCONNX call was issued with SSL configuration options specified, but an error occurred during the initialization of the SSL environment.

MQRC_SSL_PEER_NAME_MISMATCH

The DN pattern specified in the sslPeerName property did not match the DN presented by the queue manager.

MQRC_SSL_PEER_NAME_ERROR

The DN pattern specified in the sslPeerName property was not valid.

MQRC_UNSUPPORTED_CIPHER_SUITE

The CipherSuite named in sslCipherSuite was not recognized by the JSSE provider. A full list of CipherSuites supported by the JSSE provider can be obtained by a program using the `SSLConnectionFactory.getSupportedCipherSuites()` method. A list of CipherSuites that can be used to communicate with WebSphere MQ can be found in [“SSL CipherSpecs and CipherSuites in WebSphere MQ classes for Java” on page 675](#).

MQRC_SSL_CERTIFICATE_REVOKED

The certificate presented by the queue manager was found in a CRL specified with the sslCertStores property. Update the queue manager to use trusted certificates.

MQRC_SSL_CERT_STORE_ERROR

None of the supplied CertStores could be searched for the certificate presented by the queue manager. The `MQException.getCause()` method returns the error that occurred while searching the first CertStore attempted. If the causal exception is `NoSuchElementException`, `ClassCastException`, or `NullPointerException`, check that the Collection specified on the sslCertStores property contains at least one valid CertStore object.

SSL CipherSpecs and CipherSuites in WebSphere MQ classes for Java

Whether a IBM WebSphere MQ classes for Java application can establish a connection to a queue manager depends on the CipherSpec specified at the server end of the MQI channel and the CipherSuite specified at the client end.

For each combination of CipherSpec and CipherSuite, whether a IBM WebSphere MQ classes for Java application can connect to a queue manager depends on the value of the `sslFipsRequired` field in the `MQEnvironment` class, or on the value of the environment property `CMQC.SSL_FIPS_REQUIRED_PROPERTY`.

At the server end of an MQI channel, the name of a CipherSpec can be specified as the value of the `SSLCIPH` parameter on a `DEFINE CHANNEL CHLTYPE(SVRCONN)` command. At the client end of an MQI channel, a IBM WebSphere MQ classes for Java application can set the `sslCipherSuite` field in the `MQEnvironment` class, or set the environment property `CMQC.SSL_CIPHER_SUITE_PROPERTY`.

Configuring your application to use IBM Java or Oracle Java CipherSuite mappings

From IBM WebSphere MQ Version 7.5.0, Fix Pack 5, you can configure whether your application uses the default IBM Java CipherSuite to WebSphere MQ CipherSpec mappings, or the Oracle CipherSuite to WebSphere MQ CipherSpec mappings. Therefore, you can use TLS CipherSuites whether your application uses an IBM JRE or an Oracle JRE. The Java System Property `com.ibm.mq.cfg.useIBMCipherMappings` controls which mappings are used. The property can be one of the following values:

true

Use the IBM Java CipherSuite to WebSphere MQ CipherSpec mappings.

This value is the default value.

false

Use the Oracle CipherSuite to WebSphere MQ CipherSpec mappings.

The following table lists the CipherSpecs supported by IBM WebSphere MQ and their equivalent CipherSuites. The table also indicates whether a IBM WebSphere MQ classes for Java application can

establish a connection to a queue manager if a CipherSpec is specified at the server end of the MQI channel and the equivalent CipherSuite is specified at the client end.

<i>Table 89. CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites</i>		
CipherSpec	Equivalent CipherSuite	Connection possible if SFIPS ¹ is set to YES?
NULL_MD5	SSL_RSA_WITH_NULL_MD5	No
NULL_SHA	SSL_RSA_WITH_NULL_SHA	No
RC4_MD5_EXPORT	SSL_RSA_EXPORT_WITH_RC4_40_MD5 (IBM JRE) No equivalent for Oracle JRE.	No
RC4_MD5_US	SSL_RSA_WITH_RC4_128_MD5	No
RC4_SHA_US	SSL_RSA_WITH_RC4_128_SHA (IBM JRE) No equivalent for Oracle JRE.	No
RC2_MD5_EXPORT	SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5 (IBM JRE) SSL_RSA_EXPORT_WITH_RC4_40_MD5 (Oracle JRE)	No
DES_SHA_EXPORT	SSL_RSA_WITH_DES_CBC_SHA (IBM JRE) No equivalent for Oracle JRE.	No
RC4_56_SHA_EXPORT1024	SSL_RSA_EXPORT1024_WITH_RC4_56_SHA (IBM JRE) No equivalent for Oracle JRE.	No
DES_SHA_EXPORT1024	SSL_RSA_EXPORT1024_WITH_DES_CBC_SHA (IBM JRE) No equivalent for Oracle JRE.	No
TRIPLE_DES_SHA_US	SSL_RSA_WITH_3DES_EDE_CBC_SHA (IBM JRE) No equivalent for Oracle JRE.	No
TLS_RSA_WITH_NULL_SHA256	SSL_RSA_WITH_NULL_SHA256 (IBM JRE) TLS_RSA_WITH_NULL_SHA256 (Oracle JRE)	No ⁷
TLS_RSA_WITH_AES_128_CBC_SHA	SSL_RSA_WITH_AES_128_CBC_SHA (IBM JRE) TLS_RSA_WITH_AES_128_CBC_SHA (Oracle JRE)	Yes ^{5 7}
TLS_RSA_WITH_AES_128_CBC_SHA256	SSL_RSA_WITH_AES_128_CBC_SHA256 (IBM JRE) TLS_RSA_WITH_AES_128_CBC_SHA256 (Oracle JRE)	Yes ^{5 7}
TLS_RSA_WITH_AES_256_CBC_SHA	SSL_RSA_WITH_AES_256_CBC_SHA (IBM JRE) TLS_RSA_WITH_AES_256_CBC_SHA (Oracle JRE)	Yes ^{5 7}

Table 89. CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites (continued)		
CipherSpec	Equivalent CipherSuite	Connection possible if SFIPS ¹ is set to YES ?
TLS_RSA_WITH_AES_256_CBC_SHA256	SSL_RSA_WITH_AES_256_CBC_SHA256 (IBM JRE) TLS_RSA_WITH_AES_256_CBC_SHA256 (Oracle JRE)	Yes ^{5 7}
AES_SHA_US ²		
TLS_RSA_WITH_DES_CBC_SHA ⁸	SSL_RSA_WITH_DES_CBC_SHA	No ³
TLS_RSA_WITH_3DES_EDE_CBC_SHA ^{8 9}	SSL_RSA_WITH_3DES_EDE_CBC_SHA	Yes
FIPS_WITH_DES_CBC_SHA	SSL_RSA_FIPS_WITH_DES_CBC_SHA (IBM JRE) No equivalent for Oracle JRE.	No ⁴
FIPS_WITH_3DES_EDE_CBC_SHA	SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA (IBM JRE) No equivalent for Oracle JRE.	No ⁶

Notes:

1. In a IBM WebSphere MQ classes for Java application, indicate that only FIPS-certified algorithms are to be used by setting the `sslFipsRequired` field in the `MQEnvironment` class to `true` and indicate that non-FIPS-certified algorithms can also be used by setting the `sslFipsRequired` field to `false`. Alternatively, set the environment property `CMQC.SSL_FIPS_REQUIRED_PROPERTY`.
2. This CipherSpec has no equivalent CipherSuite.
3. This CipherSpec was FIPS 140-2 certified before 19th May 2007.
4. This CipherSpec was FIPS 140-2 certified before 19th May 2007. The name `FIPS_WITH_DES_CBC_SHA` is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.
5. These CipherSpecs (`TLS_RSA_WITH_AES_128_CBC_SHA`, `TLS_RSA_WITH_AES_128_CBC_SHA256`, `TLS_RSA_WITH_AES_256_CBC_SHA`, `TLS_RSA_WITH_AES_256_CBC_SHA256`) cannot be used to secure a connection from the WebSphere MQ Explorer to a queue manager unless the appropriate unrestricted policy files are applied to the JRE used by the Explorer.

See [Security information](#) for further information on policy files.
6. The name `FIPS_WITH_3DES_EDE_CBC_SHA` is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.
7. These CipherSpecs (`TLS_RSA_WITH_NULL_SHA256`, `TLS_RSA_WITH_AES_128_CBC_SHA`, `TLS_RSA_WITH_AES_256_CBC_SHA256`, `TLS_RSA_WITH_AES_256_CBC_SHA`, `TLS_RSA_WITH_AES_256_CBC_SHA256`) require IBM JREs 6.0 SR13 FP2 , 7.0 SR4 FP2 or later.
8. These CipherSpecs (`TLS_RSA_WITH_3DES_EDE_CBC_SHA`, `TLS_RSA_WITH_DES_CBC_SHA`, `TLS_RSA_WITH_RC4_128_SHA256`) can use either SSLv3 or TLS. By default, when FIPS is not enabled, SSLv3 is used. To use TLS, set the Java System Property **`com.ibm.mq.cfg.preferTLS`** to `true`.
9. This CipherSpec `TLS_RSA_WITH_3DES_EDE_CBC_SHA` is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

Related information

[Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client](#)

[Federal Information Processing Standards \(FIPS\) for UNIX, Linux and Windows](#)

[MQdev blog: MQ Java, TLS Ciphers, Non-IBM JREs & APARs IT06775, IV66840, IT09423, IT10837](#)

[MQdev blog: The relationship between MQ CipherSpecs and Java Cipher Suites](#)

Running WebSphere MQ classes for Java applications

If you write an application (a class that contains a `main()` method), using either the client or the bindings mode, run your program using the Java interpreter.

Use the command:

```
java -Djava.library.path=library_path MyClass
```

where *library_path* is the path to the WebSphere MQ classes for Java libraries (see [The WebSphere MQ classes for Java libraries](#)).

WebSphere MQ classes for Java environment-dependent behavior

WebSphere MQ classes for Java allow you to create applications that can run against different versions of WebSphere MQ. This collection of topics describes the behavior of the Java classes dependent on these different versions.

WebSphere MQ classes for Java provides a core of classes, which provide consistent function and behavior in all the environments. Features outside this core depend on the capability of the queue manager to which the application is connected.

Except where noted here, the behavior exhibited is as described in the Application Programming Reference appropriate to the queue manager.

Core classes in WebSphere MQ classes for Java

WebSphere MQ classes for Java contains a core set of classes, which can be used in all environments.

The following set of classes are considered core classes, and can be used in all environments with only the minor variations listed in [“Restrictions and variations for core classes of WebSphere MQ classes for Java” on page 679](#).

- MQEnvironment
- MQException
- MQGetMessageOptions

Excluding:

- MatchOptions
- GroupStatus
- SegmentStatus
- Segmentation

- MQManagedObject

Excluding:

- inquire()
- set()

- MQMessage

Excluding:

- groupId

- messageFlags
- messageSequenceNumber
- offset
- originalLength
- MQPoolServices
- MQPoolServicesEvent
- MQPoolServicesEventListener
- MQPoolToken
- MQPutMessageOptions

Excluding:

- knownDestCount
- unknownDestCount
- invalidDestCount
- recordFields
- MQProcess
- MQQueue
- MQQueueManager

Excluding:

- begin()
- accessDistributionList()
- MQSimpleConnectionManager
- MQTopic
- MQC

Note:

1. Some constants are not included in the core (see [“Restrictions and variations for core classes of WebSphere MQ classes for Java”](#) on page 679 for details); do not use them in completely portable programs.
2. Some platforms do not support all connection modes. On these platforms, you can use only the core classes and options that relate to the supported modes. (See [“Connection options for WebSphere MQ classes for Java”](#) on page 626.)

Restrictions and variations for core classes of WebSphere MQ classes for Java

The core classes generally behave consistently across all environments, even if the equivalent MQI calls normally have environment differences. The behavior is as if a Windows, UNIX or Linux WebSphere MQ queue manager is used, except for the following minor restrictions and variations.

Restrictions for MQGMO_ values in WebSphere MQ classes for Java*

Certain MQGMO_* values are not supported by all queue managers.

Use of the following MQGMO_* values might result in an MQException being thrown from an MQQueue.get():

```
MQGMO_SYNCPOINT_IF_PERSISTENT
MQGMO_MARK_SKIP_BACKOUT
MQGMO_BROWSE_MSG_UNDER_CURSOR
MQGMO_LOCK
MQGMO_UNLOCK
MQGMO_LOGICAL_ORDER
MQGMO_COMPLETE_MESSAGE
```

MQGMO_ALL_MSGS_AVAILABLE
MQGMO_ALL_SEGMENTS_AVAILABLE
MQGMO_UNMARKED_BROWSE_MSG
MQGMO_MARK_BROWSE_HANDLE
MQGMO_MARK_BROWSE_CO_OP
MQGMO_UNMARK_BROWSE_HANDLE
MQGMO_UNMARK_BROWSE_CO_OP

Additionally, MQGMO_SET_SIGNAL is not supported when used from Java.

Restrictions for MQPMRF_ values in WebSphere MQ classes for Java*

These are used only when putting messages to a distribution list, and are supported only by queue managers supporting distribution lists. For example, z/OS queue managers do not support distribution lists.

Restrictions for MQPMO_ values in WebSphere MQ classes for Java*

Certain MQPMO_* values are not supported by all queue managers

Use of the following MQPMO_* values might result in an MQException being thrown from an MQQueue.put() or an MQQueueManager.put():

MQPMO_LOGICAL_ORDER
MQPMO_NEW_CORREL_ID
MQPMO_NEW_MESSAGE_ID
MQPMO_RESOLVE_LOCAL_Q

Restrictions and variations for MQCNO_ values in WebSphere MQ classes for Java*

Certain MQCNO_* values are not supported.

- Automatic client reconnect is not supported by the WebSphere MQ classes for Java. Whatever value MQCNO_RECONNECT_* you set, the connection continues to behave as if you set MQCNO_RECONNECT_DISABLED.
- MQCNO_FASTPATH is ignored on queue managers that do not support MQCNO_FASTPATH. It is also ignored by client connections.

Restrictions for MQRO_ values in WebSphere MQ classes for Java*

The following report options can be set.

MQRO_EXCEPTION_WITH_FULL_DATA
MQRO_EXPIRATION_WITH_FULL_DATA
MQRO_COA_WITH_FULL_DATA
MQRO_COD_WITH_FULL_DATA
MQRO_DISCARD_MSG
MQRO_PASS_DISCARD_AND_EXPIRY

For more information see, [Report](#).

Features outside the core classes of WebSphere MQ classes for Java

WebSphere MQ classes for Java contain certain functions that are specifically designed to use API extensions that are not supported by all queue managers. This collection of topics describes how they behave when using a queue manager that does *not* support them.

Variations in the MQQueueManager constructor option

Some of the MQQueueManager constructors include an optional integer argument. Some values of this argument are not accepted on all platforms.

Where an MQQueueManager constructor include an optional integer argument, it maps onto the MQCNO options field of the MQI, and is used to switch between normal and fast path connection. This extended form of the constructor is accepted in all environments, if the only options used are

MQCNO_STANDARD_BINDING or MQCNO_FASTPATH_BINDING. Any other options cause the constructor to fail with MQRC_OPTIONS_ERROR. The fast path option CMQC.MQCNO_FASTPATH_BINDING is honored only with a bindings connection to a queue manager that supports it. In other environments, it is ignored.

Restrictions on the MQQueueManager.begin() method

This method can be used only against a WebSphere MQ queue manager on UNIX, Linux, or Windows systems in bindings mode. Otherwise, it fails with MQRC_ENVIRONMENT_ERROR.

See [“JTA/JDBC coordination using WebSphere MQ classes for Java” on page 668](#) for more details.

Variations in the MQGetMessageOptions fields

Some queue managers do not support the Version 2 MQGMO structure, so you must set some fields to their default values.

When using a queue manager that does not support the Version 2 MQGMO structure, leave the following fields set to their default values:

- GroupStatus
- SegmentStatus
- Segmentation

Also, the MatchOptions field supports only MQMO_MATCH_MSG_ID and MQMO_MATCH_CORREL_ID. If you put unsupported values into these fields, the subsequent MQDestination.get() fails with MQRC_GMO_ERROR. If the queue manager does not support the Version 2 MQGMO structure, these fields are not updated after a successful MQDestination.get().

Restrictions in distribution lists in WebSphere MQ classes for Java

Not all queue managers allow you to open an MQDistributionList.

The following classes are used to create distribution lists:

- MQDistributionList
- MQDistributionListItem
- MQMessageTracker

You can create and populate MQDistributionLists and MQDistributionListItems in any environment, but not all queue managers allow you to open an MQDistributionList. In particular, z/OS queue managers do not support distribution lists. Attempting to open an MQDistributionList when using such a queue manager results in MQRC_OD_ERROR.

Variations in MQPutMessageOptions fields

If a queue manager does not support distribution lists, certain MQPMO fields are treated differently.

Four fields in the MQPMO are rendered as the following member variables in the MQPutMessageOptions class:

- knownDestCount
- unknownDestCount
- invalidDestCount
- recordFields

These fields are primarily intended for use with distribution lists. However, a queue manager that supports distribution lists also fills in the DestCount fields after an MQPUT to a single queue. For example, if the queue resolves to a local queue, knownDestCount is set to 1 and the other two count fields are set to 0.

If the queue manager does not support distribution lists, these values are simulated as follows:

- If the put() succeeds, unknownDestCount is set to 1, and the others are set to 0.
- If the put() fails, invalidDestCount is set to 1, and the others are set to 0.

The recordFields variable is used with distribution lists. A value can be written into recordFields at any time, regardless of the environment. It is ignored if the MQPutMessageOptions object is used on a subsequent MQDestination.put() or MQQueueManager.put(), rather than MQDistributionList.put().

Restrictions in MQMD fields with WebSphere MQ classes for Java

Certain MQMD fields concerned with message segmentation should be left at their default value when using a queue manager that does not support segmentation.

The following MQMD fields are largely concerned with message segmentation:

- GroupId
- MsgSeqNumber
- Offset
- MsgFlags
- OriginalLength

If an application sets any of these MQMD fields to values other than their defaults, and then does a put() or get() on a queue manager that does not support these, the put() or get() raises an MQException with MQRC_MD_ERROR. A successful put() or get() with such a queue manager always leaves the MQMD fields set to their default values. Do not send a grouped or segmented message to a Java application that runs against a queue manager that does not support message grouping and segmentation.

If a Java application attempts to get() a message from a queue manager that does not support these fields, and the physical message to be retrieved is part of a group of segmented messages (that is, it has non-default values for the MQMD fields), it is retrieved without error. However, the MQMD fields in the MQMessage are not updated, the MQMessage format property is set to MQFMT_MD_EXTENSION, and the true message data is prefixed with an MQMDE structure that contains the values for the new fields.

Restrictions for WebSphere MQ classes for Java under CICS Transaction Server

In the CICS Transaction Server for z/OS environment, only the main (first) thread is allowed to issue CICS or WebSphere MQ calls.

Note, that WebSphere MQ JMS classes are not supported for use in a CICS Java application.

It is therefore not possible to share MQQueueManager or MQQueue objects between threads in this environment, or to create a new MQQueueManager on a child thread.

Running IBM WebSphere MQ classes for Java applications within Java platform Enterprise Edition

There are certain restrictions and design considerations that must be taken into account before using IBM WebSphere MQ classes for Java in Java EE

IBM WebSphere MQ classes for Java has restrictions when used within a Java EE environment. There are also additional considerations that must be taken into account when designing, implementing, and managing an IBM WebSphere MQ classes for Java application that runs inside a Java EE environment. These restrictions and considerations are outlined in the following sections.

JTA transactions restrictions

The only supported transaction manager for applications using IBM WebSphere MQ classes for Java is IBM WebSphere MQ itself. Although an application under JTA control can make use IBM WebSphere MQ classes for Java, any work performed through these classes is not controlled by the JTA units of work. They instead form local units of work separate from those managed by the application server through the JTA interfaces. In particular, any rollback of the JTA transaction does not result in a rollback of any sent or received messages. This restriction applies to application or bean managed transactions and to container managed transactions, and all Java EE containers. To perform messaging work directly with IBM WebSphere MQ inside application server-coordinated transactions, IBM WebSphere MQ classes for JMS must be used instead.

Thread creation

IBM WebSphere MQ classes for Java creates threads internally for various operations. For example, when running in BINDINGS mode to call directly on a local queue manager, the calls are made on a 'worker' thread created internally by IBM WebSphere MQ classes for Java. Other threads can be created internally, for example to clear unused connections from a connection pool or to remove subscriptions for terminated publish/subscribe applications.

Some Java EE applications (for example those running in EJB and Web containers) must not create new threads. Instead, all work must be performed on the main application threads managed by the application server. When applications use IBM WebSphere MQ classes for Java, the application server might not be able to distinguish between application code and the IBM WebSphere MQ classes for Java code, so the threads previously described cause the application to be non-compliant with the container specification. IBM WebSphere MQ classes for JMS does not break these Java EE specifications and so can be used instead.

Security restrictions

Security policies implemented by an application server might prevent certain operations that are undertaken by the IBM WebSphere MQ classes for Java API, such as creating and operating new threads of control (as described in the preceding sections).

For example, application servers typically run with Java Security disabled by default, and allow it to be enabled through some application server-specific configuration (some application servers also allow more detailed configuration of the policies used within Java Security). When Java Security is enabled, IBM WebSphere MQ classes for Java might break the Java Security policy threading rules defined for the application server, and the API might not be able to create all the threads that it needs in order to function. To prevent problems with thread management, the use of IBM WebSphere MQ classes for Java is not supported in environments where Java Security is enabled.

Application isolation considerations

An intended benefit of running applications within a Java EE environment is application isolation. The design and implementation of IBM WebSphere MQ classes for Java predate the Java EE environment. IBM WebSphere MQ classes for Java can be used in a manner which does not support the concept of application isolation. Specific examples of considerations in this area include:

- The use of static (JVM process-wide) settings within the MQEnvironment class, such as:
 - the user ID and password to be used for connection identification and authentication
 - the host name, port, and channel used for client connections
 - SSL configuration for secured client connections

Modifying any of the MQEnvironment properties for the benefit of one application also affect other applications using the same properties. When running in a multi-application environment such as Java EE, each application must use its own distinct configuration through the creation of MQQueueManager objects with a specific set of properties, rather than defaulting to the properties configured in the process-wide MQEnvironment class.

- The MQEnvironment class introduces a number of static methods which act globally on all applications using IBM WebSphere MQ classes for Java within the same JVM process, and there is no way to override this behavior for particular applications. Examples include:
 - configuring SSL properties, such as the location of the keystore
 - configuring client channel exits
 - enabling or disabling diagnostic tracing
 - managing the default connection pool used to optimize the use of connections to queue managers

Invoking such methods affects all applications running in the same Java EE environment.

- Connection pooling is enabled to optimize the process of making multiple connections to the same queue manager. The default connection pool manager is process-wide, and shared by multiple applications. Changes to connection pool configuration, such as replacing the default connection manager for one application using the `MQEnvironment.setDefaultConnectionManager()` method therefore affects other applications running in the same Java EE application server.
- SSL is configured for applications using IBM WebSphere MQ classes for Java using the `MQEnvironment` class and `MQQueueManager` object properties. It is not integrated with the managed security configuration of the application server itself. You must ensure that you configure IBM WebSphere MQ classes for Java appropriately to provide your required level of security, and not use the application server configuration.

Bindings mode restrictions

IBM WebSphere MQ and WebSphere Application Server can be installed on the same machine such that the major versions of the queue manager and of the IBM WebSphere MQ resource adapter (RA) shipped in WebSphere Application Server are different. For instance WebSphere Application Server Version 7.0, which ships an IBM WebSphere MQ RA level of 7.0.1, can be installed on the same machine as a Version 6.0 queue manager.

If the queue manager and resource adapter major versions are different, bindings connections cannot be used. Any connections from WebSphere Application Server to the queue manager using the resource adapter must use client type connections. Bindings connections can be used if the versions are the same.

Using WebSphere MQ classes for JMS

WebSphere MQ classes for Java Message Service (WebSphere MQ classes for JMS) is the JMS provider that is supplied with WebSphere MQ. As well as implementing the interfaces defined in the `javax.jms` package, WebSphere MQ classes for JMS provides two sets of extensions to the JMS API.

The JMS specification defines a set of interfaces that applications can use to perform messaging operations. The `javax.jms` package defines the JMS interfaces, and a JMS provider implements these interfaces for a specific messaging product. WebSphere MQ version 7.5 currently uses the JMS 1.1 specification. WebSphere MQ classes for JMS is a JMS provider that implements the JMS interfaces for WebSphere MQ.

The JMS specification expects `ConnectionFactory` and `Destination` objects to be administered objects. An administrator creates and maintains administered objects in a central repository, and a JMS application retrieves these objects using the Java Naming and Directory Interface (JNDI). WebSphere MQ classes for JMS supports the use of administered objects, and an administrator can use either the WebSphere MQ JMS administration tool or WebSphere MQ Explorer to create and maintain administered objects.

WebSphere MQ classes for JMS also provides two sets of extensions to the JMS API. The main focus of these extensions concerns creating and configuring connection factories and destinations dynamically at run time, but the extensions also provide function that is not directly related to messaging, such as function for problem determination.

The WebSphere MQ JMS extensions

Previous releases of WebSphere MQ classes for JMS contain extensions that are implemented in objects such as `MQConnectionFactory`, `MQQueue`, and `MQTopic` objects. These objects have properties and methods that are specific to WebSphere MQ. The objects can be administered objects, or an application can create the objects dynamically at run time. This release of WebSphere MQ classes for JMS maintains these extensions, which are now known as the WebSphere MQ JMS extensions. You can continue to use, without change, any applications that use these extensions.

The IBM JMS extensions

This release of WebSphere MQ classes for JMS provides a more generic set of extensions to the JMS API, which are not specific to WebSphere MQ as the messaging system. These extensions are known as the IBM JMS extensions and have the following broad objectives:

- To provide a greater level of consistency across IBM JMS providers
- To make it easier to write a bridge application between two IBM messaging systems

- To make it easier to port an application from one IBM JMS provider to another

The extensions provide function that is similar to that provided in Message Service Client for C/C++ and Message Service Client for .NET.

Why should I use WebSphere MQ classes for JMS?

Using WebSphere MQ classes for JMS has the following advantages:

- You can reuse JMS skills.

WebSphere MQ classes for JMS is a JMS provider that implements the JMS interfaces for WebSphere MQ as the messaging system. If your organization is new to WebSphere MQ, but already has JMS application development skills, you might find it easier to use the familiar JMS API to access WebSphere MQ resources rather than one of the other APIs provided with WebSphere MQ.

- JMS is an integral part of Java Platform, Enterprise Edition (Java EE).

JMS is the natural API to use for messaging on the Java EE platform. Every application server that is Java EE compliant must include a JMS provider. You can use JMS in application clients, servlets, JavaServer pages (JSPs), enterprise Java beans (EJBs), and message driven beans (MDBs). Note in particular that Java EE applications use MDBs to process messages asynchronously, and all messages are delivered to MDBs as JMS messages.

- An administrator can create and maintain JMS administered objects in a central repository, and WebSphere MQ classes for JMS applications can retrieve these objects using the Java Naming and Directory Interface (JNDI).

JMS connection factories and destinations encapsulate WebSphere MQ specific information such as queue manager names, channel names, connection options, queue names, and topic names. If connection factories and destinations are stored as administered objects, this information is not hard coded into an application. This arrangement therefore provides the application with a degree of independence from the underlying WebSphere MQ configuration.

- JMS is an industry standard API that can provide application portability.

A JMS application can use JNDI to retrieve connection factories and destinations that are stored as administered objects, and use only the interfaces defined in the `javax.jms` package to perform messaging operations. The application is then entirely independent of any JMS provider, such as WebSphere MQ classes for JMS, and can be ported from one JMS provider to another without any change to the application.

If JNDI is not available in a particular application environment, a WebSphere MQ classes for JMS application can use extensions to the JMS API to create and configure connection factories and destinations dynamically at run time. The application is then completely self contained, but is tied to WebSphere MQ classes for JMS as the JMS provider.

- Bridge applications might be easier to write using JMS.

A bridge application is an application that receives messages from one messaging system and sends them to another messaging system. Writing a bridge application can be complicated using product specific APIs and message formats. Instead, you can write a bridge application using two JMS providers, one for each messaging system. The application then uses only one API, the JMS API, and processes only JMS messages.

Getting started with WebSphere MQ classes for JMS

This topic provides an overview of WebSphere MQ classes for JMS and tells you what you need to know before using WebSphere MQ classes for JMS.

Prerequisites for WebSphere MQ classes for JMS

To develop and run WebSphere MQ classes for JMS applications, you need certain software components as prerequisites.

For the latest information about the prerequisites for WebSphere MQ classes for JMS, see the [WebSphere MQ readme file](#).

To develop WebSphere MQ classes for JMS applications, you need a Java 2 Software Development Kit (SDK). Details of the JDKs supported with your operating system can be found on the [WebSphere MQ System requirements page](#). See [WebSphere MQ Requirements](#).

To run WebSphere MQ classes for JMS applications, you need the following software components:

- A WebSphere MQ queue manager
- A Java Runtime Environment (JRE), for each system on which you run applications

If you require SSL connections to use cryptographic modules that are FIPS 140-2 certified, you need the IBM Java JSSE FIPS provider (IBMJSEFIPS). Every IBM Java 2 SDK and JRE at Version 5 or later contains IBMJSEFIPS.

You can use Internet Protocol Version 6 (IPv6) addresses in your WebSphere MQ classes for JMS applications provided IPv6 addresses are supported by your Java virtual machine (JVM) and the TCP/IP implementation on your operating system. The WebSphere MQ JMS administration tool (see [“Using the WebSphere MQ JMS administration tool” on page 893](#)) also accepts IPv6 addresses.

The WebSphere MQ JMS administration tool and WebSphere MQ Explorer use the Java Naming and Directory Interface (JNDI) to access a directory service, which stores administered objects. WebSphere MQ classes for JMS applications can also use JNDI to retrieve administered objects from a directory service. A service provider is code that provides access to a directory service by mapping JNDI calls to calls to the directory service. The following service providers are supplied with WebSphere MQ classes for JMS:

- A Lightweight Directory Access Protocol (LDAP) service provider in the files `ldap.jar` and `providerutil.jar`. The LDAP service provider provides access to a directory service based on an LDAP server.
- A file system service provider in the files `fscontext.jar` and `providerutil.jar`. The file system service provider provides access to a directory service based on the local file system.

If you intend to use a directory service based on an LDAP server, you must install and configure an LDAP server, or have access to an existing LDAP server. In particular, you must configure the LDAP server to store Java objects. For information about how to install and configure your LDAP server, see the documentation that is supplied with the server.

Preparing JMS programs for the IBM WebSphere MQ client for HP Integrity NonStop Server

This topic explains what you need to know before you develop and run JMS programs for the IBM WebSphere MQ client for HP Integrity NonStop Server.

The IBM WebSphere MQ classes for JMS are installed as part of the IBM WebSphere MQ client for HP Integrity NonStop Server installation. For details of a summary of the contents of the installation, see [File system](#).

Some aspects of client functionality are specific to the host operating system. For more information about the supported features for the IBM WebSphere MQ client for HP Integrity NonStop Server, see [IBM WebSphere MQ client for HP Integrity NonStop Server supported environments and features](#).

Prerequisites

To build and run JMS applications, the *HP Integrity NonStop Server for Java* component must be installed and available.

Setup

For information about setting up the environment to run and build applications in which you can use the IBM WebSphere MQ classes for JMS, see [“Environment variables used by IBM WebSphere MQ classes for JMS” on page 691](#).

For information about the steps required to configure a queue manager to accept connections from client applications, see [“Post installation setup for WebSphere MQ classes for JMS applications” on page 737](#).

For information about validating your IBM WebSphere MQ classes for JMS environment, see [“The point-to-point installation verification test for WebSphere MQ classes for JMS” on page 740](#).

Writing applications

For more information about writing JMS applications, see [“Writing WebSphere MQ classes for JMS applications” on page 770](#).

For more information about using the IBM WebSphere MQ JMS administration tool, see [“Using the WebSphere MQ JMS administration tool” on page 893](#).

Samples

Sample applications are provided in the following subdirectory of the installation: `opt/mqm/samp/jms`.

For more information about the configuration steps required to run the samples, see [“Preparing and running the sample programs” on page 104](#).

Problem solving

For information about solving problems, see [“Solving problems with IBM WebSphere MQ classes for JMS” on page 761](#).

Installation and configuration of WebSphere MQ classes for JMS

This section describes the directories and files that are created when you install WebSphere MQ classes for JMS and tells you how to configure WebSphere MQ classes for JMS after installation.

Related concepts

[“What is installed for IBM WebSphere MQ classes for JMS” on page 688](#)

A number of files and directories are created when you install IBM WebSphere MQ classes for JMS. On Windows some configuration is performed during installation by automatically setting environment variables. On other platforms, and in certain Windows environments, you must set environment variables before you can run IBM WebSphere MQ classes for JMS applications.

[“Running WebSphere MQ classes for JMS applications under the Java security manager” on page 697](#)
WebSphere MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the security manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

[“The IBM WebSphere MQ resource adapter” on page 701](#)

The resource adapter allows applications running in an application server to access IBM WebSphere MQ resources. It supports inbound and outbound communication.

[“Post installation setup for WebSphere MQ classes for JMS applications” on page 737](#)

This topic tells you what authorities WebSphere MQ classes for JMS applications need in order to access the resources of a queue manager. It also introduces connection modes and describes how to configure a queue manager so that applications can connect in client mode.

[“The point-to-point installation verification test for WebSphere MQ classes for JMS” on page 740](#)

A point-to-point installation verification test (IVT) program is supplied with WebSphere MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, sends a message to the queue called `SYSTEM.DEFAULT.LOCAL.QUEUE`, and then receives the message from the queue. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

[“The publish/subscribe installation verification test for WebSphere MQ classes for JMS” on page 744](#)

A publish/subscribe installation verification test (IVT) program is supplied with WebSphere MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, subscribes to a topic, publishes a message on the topic, and then receives the message that it has just published. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

[“The installation verification test program for the WebSphere MQ resource adapter” on page 747](#)

The IVT program is supplied as an EAR file. To use the program, you must deploy it and define some objects as JCA resources.

[“Configuring the resource adapter for outbound communication” on page 718](#)

To configure outbound communication, define the properties of a `ConnectionFactory` object and an administered destination object.

[“Support for OSGi” on page 760](#)

OSGi provides a framework that supports the deployment of applications as bundles. Nine OSGi bundles are supplied as part of the IBM WebSphere MQ classes for JMS.

[“Solving problems with IBM WebSphere MQ classes for JMS” on page 761](#)

You can investigate problems by running the installation verification programs, and by using the trace and log facilities.

Related tasks

[“Installing and testing the MQ resource adapter in WAS CE” on page 750](#)

Installing the IBM WebSphere MQ resource adapter and running the installation verification test (IVT) application in WebSphere Application Server CE.

[“Deploying the IVT application on WAS CE with a custom MQ environment” on page 752](#)

If you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, you must modify the IVT application and associated scripts in WebSphere Application Server CE before deploying the resource adapter or IVT application.

[“Deploying the IVT application in JBoss with a custom IBM WebSphere MQ environment” on page 755](#)

When installing the IBM WebSphere MQ resource adapter in JBoss, if you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, you must first modify the IVT application and associated scripts in JBoss before deploying the resource adapter or IVT application.

[Problem determination for the IBM WebSphere MQ resource adapter](#)

Related reference

[“Scripts provided with WebSphere MQ classes for JMS” on page 759](#)

A number of scripts are provided to assist with common tasks that need to be performed when using WebSphere MQ classes for JMS.

What is installed for IBM WebSphere MQ classes for JMS

A number of files and directories are created when you install IBM WebSphere MQ classes for JMS.

On Windows some configuration is performed during installation by automatically setting environment variables. On other platforms, and in certain Windows environments, you must set environment variables before you can run IBM WebSphere MQ classes for JMS applications.

For most operating systems, the IBM WebSphere MQ classes for JMS are installed as an optional component when you install IBM WebSphere MQ. For the IBM WebSphere MQ client for HP Integrity NonStop Server, the IBM WebSphere MQ classes for JMS are installed by default. For more information about installing IBM WebSphere MQ, see:

[Installing a WebSphere MQ server](#)

[Installing an IBM WebSphere MQ client](#)

Table 90 on page 689 shows where the IBM WebSphere MQ classes for JMS files are installed on each platform.

<i>Table 90. IBM WebSphere MQ classes for JMS installation directories</i>	
Platform	Directory
AIX	<code>MQ_INSTALLATION_PATH/java</code>
HP Integrity NonStop Server	<code>MQ_INSTALLATION_PATH/opt/mqm/java</code>
HP-UX, Linux, and Solaris	<code>MQ_INSTALLATION_PATH/java</code>
Windows	<code>MQ_INSTALLATION_PATH\java</code>
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which IBM WebSphere MQ is installed.	

The installation directory includes:

- The IBM WebSphere MQ classes for JMS JAR files, which are located in the `MQ_INSTALLATION_PATH\java\lib` directory.
- The IBM WebSphere MQ native libraries, which are used by applications that use the Java Native Interface.

The 32-bit native libraries are installed into the `MQ_INSTALLATION_PATH\java\lib` directory and the 64-bit native libraries can be found in the `MQ_INSTALLATION_PATH\java\lib64` directory.

For more information about the IBM WebSphere MQ native libraries, see [“Configuring the Java Native Interface \(JNI\) libraries”](#) on page 693.

- Additional scripts described in [“Scripts provided with WebSphere MQ classes for JMS”](#) on page 759. These scripts are located in the `MQ_INSTALLATION_PATH\java\bin` directory.
- The specifications of the IBM WebSphere MQ classes for JMS API. The Javadoc tool has been used to generate the HTML pages containing the specifications of the API.

The HTML pages are in the `MQ_INSTALLATION_PATH\java\doc\WMQJMSClasses` directory.

On UNIX, Linux, and Windows systems, this subdirectory contains the individual HTML pages.

- Support for OSGi. OSGi bundles are installed in the `java\lib\OSGi` directory and described in [“Support for OSGi”](#) on page 760.
- The IBM WebSphere MQ Resource Adapter, which can be deployed into any JCA 1.5 (or later) compliant application server.

The IBM WebSphere MQ Resource Adapter is located in the `MQ_INSTALLATION_PATH\java\lib\jca` directory; for further information, see [“The IBM WebSphere MQ resource adapter”](#) on page 701

- On Windows, symbols that can be used for debugging are installed in the `MQ_INSTALLATION_PATH\java\lib\symbols` directory.

The installation directory also includes some files that belong to other IBM WebSphere MQ components. These directories are as follows:

- The IBM WebSphere MQ transport for SOAP, which provides a JMS transport for SOAP, is installed into the `MQ_INSTALLATION_PATH\java\lib\soap` directory. For further information on IBM WebSphere MQ transport for SOAP, see the section of the infocenter describing [“WebSphere MQ transport for SOAP”](#) on page 905.

- On distributed platforms, the IBM WebSphere MQ Bridge for HTTP is installed in the `MQ_INSTALLATION_PATH\java\lib\http` directory. For further information on IBM WebSphere MQ bridge for HTTP, see the section of the infocenter describing [“WebSphere MQ bridge for HTTP”](#) on page 977

Some sample applications are supplied with IBM WebSphere MQ classes for JMS. Table 91 on page 690 shows where the sample applications are installed on each platform.

Table 91. Samples directories	
Platform	Directory
AIX	<code>MQ_INSTALLATION_PATH/samp/jms</code>
HP Integrity NonStop Server	<code>MQ_INSTALLATION_PATH/opt/mqm/samp/jms</code>
HP-UX, Linux, and Solaris	<code>MQ_INSTALLATION_PATH/samp/jms</code>
Windows	<code>MQ_INSTALLATION_PATH\tools\jms</code>
<i>MQ_INSTALLATION_PATH</i> represents the high-level directory in which IBM WebSphere MQ is installed.	

After installation, you might need to perform some configuration tasks to compile and run applications.

[“Environment variables used by IBM WebSphere MQ classes for JMS”](#) on page 691 describes the classpath required to run simple IBM WebSphere MQ classes for JMS applications. This topic also describes additional JAR files that need to be referenced in special circumstances and the environment variables that you must set to run the scripts provided with IBM WebSphere MQ classes for JMS.

Should you need your IBM WebSphere MQ classes for JMS application to link to code written in languages other than Java (for example, to use the bindings transport when connecting to a Queue Manager), [“Configuring the Java Native Interface \(JNI\) libraries”](#) on page 693 explains where to find the location of the Java Native Interface (JNI) libraries to specify as a parameter of the Java command.

To control properties, such as tracing and logging of an application, you need to provide a configuration properties file. The IBM WebSphere MQ classes for JMS configuration properties file is described in [“The IBM WebSphere MQ classes for JMS configuration file”](#) on page 695.

Installing and upgrading the WebSphere MQ classes for JMS JAR files

The only supported way to get the IBM WebSphere MQ classes for JMS JAR files onto a system, is to install either the IBM WebSphere MQ product, or the [WebSphere MQ V7.5 Clients SupportPac-MQC75](#), or by using a software management tool such as Apache Maven, for more information see [“IBM WebSphere MQ classes for JMS and software management tools”](#) on page 697.

Do not move, or copy, the IBM WebSphere MQ classes for JMS JAR files or native libraries, to other machines, or to a different location on a machine where the IBM WebSphere MQ classes for JMS have been installed, unless you are using a software management tool.

- Fix packs cannot be applied to an "installation" where JAR files have been copied from another machine, because this makes it much more difficult to ensure that all the JAR files are kept in step with each other, and are at compatible levels.
- Copying the IBM WebSphere MQ classes for JMS JAR files between machines can also result in multiple copies of the files residing on the same machine, which can cause problems servicing the code, and debugging problems.
- The `dspmqver` command, used to display version information from an IBM WebSphere MQ installation, only displays version information for the IBM WebSphere MQ classes for JMS installed into the `\java\lib` directory.

If multiple copies of the files reside on the same machine, running **dspmqver** might not give accurate information about the version of the IBM WebSphere MQ classes for JMS being used by an application.

Do not include the IBM WebSphere MQ classes for JMS JAR files within application archives (such as enterprise application archives, or EAR files).

- Updates to the IBM WebSphere MQ classes for JMS cannot be applied using an IBM WebSphere MQ Fix Pack.
- It is not possible for IBM Support to easily determine the version of the IBM WebSphere MQ classes for JMS that are being used by the application.
- Problems can arise if multiple applications running inside the same Java Runtime Environment include different versions of the IBM WebSphere MQ classes for JMS, as multiple versions of the IBM WebSphere MQ classes for JMS are loaded into the Java Runtime Environment at the same time.

Examples of these problems include the following exceptions:

```
java.lang.ClassCastException :
com.ibm.mq.jmqi.system.JmqiSystemEnvironment incompatible with
com.ibm.mq.jmqi.system.JmqiSystemEnvironment

java.lang.ClassCastException :
com.ibm.mq.jms.MQQueue incompatible with com.ibm.mq.jms.MQQueue
```

- If an application uses the BINDINGS transport to connect to a queue manager, any major upgrades to the queue manager also require the application to be updated to include the corresponding level of the IBM WebSphere MQ classes for JMS.

For example, if a queue manager is upgraded to the IBM WebSphere MQ Version 7.5 level, any applications that connect to the queue manager using the BINDINGS transport also need to be updated to include the IBM WebSphere MQ Version 7.5 classes for JMS.

Environment variables used by IBM WebSphere MQ classes for JMS

Before you can compile and run IBM WebSphere MQ classes for JMS applications, the setting for your CLASSPATH environment variable must include the IBM WebSphere MQ classes for JMS Java archive (JAR) file. Depending on your requirements, you might need to add other JAR files to your class path. To run the scripts provided with IBM WebSphere MQ classes for JMS, other environment variables must be set.

To compile and run IBM WebSphere MQ classes for JMS applications, use the CLASSPATH setting for your platform as shown in Table 92 on page 691. The setting includes the samples directory, so that you can compile and run the IBM WebSphere MQ classes for JMS sample applications. Alternatively, you can specify the class path on the **java** command instead of using the environment variable.

Table 92. CLASSPATH setting to compile and run IBM WebSphere MQ classes for JMS applications, including the sample applications

Platform	CLASSPATH setting
AIX	CLASSPATH=MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms:
HP Integrity NonStop Server	CLASSPATH=MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms:
HP-UX, Linux, and Solaris	CLASSPATH=MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/samp/jms:
IBM i	CLASSPATH=/QIBM/ProdData/mqm/java/lib/com.ibm.mqjms.jar: /QIBM/ProdData/mqm/java/samples/jms:

Table 92. CLASSPATH setting to compile and run IBM WebSphere MQ classes for JMS applications, including the sample applications (continued)

Platform	CLASSPATH setting
Windows	CLASSPATH=MQ_INSTALLATION_PATH\java\lib\com.ibm.mqjms.jar; MQ_INSTALLATION_PATH\tools\jms;
z/OS	CLASSPATH=MQ_INSTALLATION_PATH/mqm/V7R0M0/java/lib/ com.ibm.mqjms.jar: MQ_INSTALLATION_PATH/mqm/V6R0M0/java/samples/jms:
MQ_INSTALLATION_PATH represents the high-level directory in which IBM WebSphere MQ is installed.	

The manifest of the JAR file com.ibm.mqjms.jar contains references to most of the other JAR files required by IBM WebSphere MQ classes for JMS applications, and so you do not need to add these JAR files to your class path. These JAR files include those required by applications that use the Java Naming and Directory Interface (JNDI) to retrieve administered objects from a directory service and by applications that use the Java Transaction API (JTA).

However, you must include additional JAR files in your class path in the following circumstances:

- If you are using channel exit classes that implement the channel exit interfaces defined in the com.ibm.mq package, instead of those defined in the com.ibm.mq.exits package, you must add the IBM WebSphere MQ classes for Java JAR file, com.ibm.mq.jar, to your class path.
- If you compile your Java code using a Java 2 Software Development Kit (SDK) at Version 1.4.2, you must add the following JAR files to your class path:

- jms.jar
- com.ibm.mq.jmqi.jar

Additionally, if your application uses JNDI to retrieve administered objects from a directory service, you must also add the following JAR files to your class path:

- fscontext.jar
- jndi.jar
- ldap.jar
- providerutil.jar

And if your application uses the JTA, you must also add jta.jar to your class path.

Note that these additional JAR files are required only for compiling your applications, not for running them.

If you compile using the -Xlint option, you might see a message warning you that com.ibm.mq.es.jar is not present. You can ignore the warning. This file is only present if you have installed the Extended Security Edition.

The scripts provided with IBM WebSphere MQ classes for JMS use the following environment variables:

MQ_JAVA_DATA_PATH

This environment variable specifies the directory for log and trace output.

MQ_JAVA_INSTALL_PATH

This environment variable specifies the directory where WebSphere MQ classes for JMS is installed.

MQ_JAVA_LIB_PATH

This environment variable specifies the directory where the WebSphere MQ classes for JMS libraries are stored, as shown in [Table 93 on page 693](#).

On Windows, all the environment variables are set automatically during installation. On any other platform, you must set them yourself.

To set the environment variables if you are using a 32-bit JVM on UNIX, HP Integrity NonStop Server, or Linux systems, you can use the script `setjmsenv`. To set the environment variables if you are using a 64-bit JVM on a UNIX or Linux system, you can use the script `setjmsenv64`. These scripts are in the `MQ_INSTALLATION_PATH/java/bin` directory, where `MQ_INSTALLATION_PATH` represents the high-level directory in which IBM WebSphere MQ is installed.

You can use the `setjmsenv` or `setjmsenv64` script in a variety of ways: You can use it as a basis for setting the required environment variables, as shown in the table, or add them to `.profile` using a text editor. If you have a non-typical setup, edit the script contents as necessary. Alternatively, you can run the script in every session from which JMS startup scripts are to be run. If you choose this option you need to run the script in every shell window you start, during the JMS verification process by typing `./setjmsenv` or `./setjmsenv64`.

Configuring the Java Native Interface (JNI) libraries

IBM WebSphere MQ classes for JMS applications, that either connect to a queue manager using the bindings transport, or that connect to a queue manager using the client transport, and use the channel exit programs written in languages other than Java, need to be run in an environment that accesses the Java Native Interface (JNI) libraries.

About this task

To set up this environment, you must configure the environment's library path so that the Java virtual machine (JVM) can load the `mqjbnd` library before you start the IBM WebSphere MQ classes for JMS application.

IBM WebSphere MQ provides two Java Native Interface (JNI) libraries:

mqjbnd

This library is used by applications that connect to a queue manager using the bindings transport. It provides the interface between the IBM WebSphere MQ classes for JMS and the queue manager. The `mqjbnd` library installed with IBM WebSphere MQ Version 7.5 can be used to connect to any IBM WebSphere MQ Version 7.5 (or earlier) queue manager.

mqjexitstub02

The `mqjexitstub02` library is loaded by the IBM WebSphere MQ classes for JMS when an application connects to a queue manager using the client transport and uses a channel exit program written in a language other than Java.

On certain platforms, IBM WebSphere MQ installs 32-bit and 64-bit versions of these JNI libraries. The location of the libraries for each platform is shown in [Table 1](#)

Table 93. The location of the IBM WebSphere MQ classes for JMS libraries for each platform	
Platform	Directory containing the IBM WebSphere MQ classes for JMS libraries
AIX	<code>MQ_INSTALLATION_PATH/java/lib</code> (32 bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64 bit libraries)
HP-UX	<code>MQ_INSTALLATION_PATH/java/lib</code> (32 bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64 bit libraries)
Linux (POWER, x86-64 and zSeries s390x platforms)	<code>MQ_INSTALLATION_PATH/java/lib</code> (32 bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64 bit libraries)

Table 93. The location of the IBM WebSphere MQ classes for JMS libraries for each platform (continued)	
Platform	Directory containing the IBM WebSphere MQ classes for JMS libraries
Linux (x86 platform) Linux (zSeries platform)	<code>MQ_INSTALLATION_PATH /java/lib</code>
Solaris (x86-64 and SPARC platforms)	<code>MQ_INSTALLATION_PATH/java/lib</code> (32 bit libraries) <code>MQ_INSTALLATION_PATH/java/lib64</code> (64 bit libraries)
Windows	<code>MQ_INSTALLATION_PATH\java\lib</code> (32 bit libraries) <code>MQ_INSTALLATION_PATH\java\lib64</code> (64 bit libraries)
<code>MQ_INSTALLATION_PATH</code> represents the high-level directory in which IBM WebSphere MQ is installed.	

Procedure

1. Configure the JVM's **java.library.path** property, which can be done in two ways:

- By specifying the JVM argument as shown in the following example:

```
-Djava.library.path=<path_to_library_directory>
```

Linux

For example, for a 64-bit JVM on Linux for a default location installation, specify:

```
-Djava.library.path=/opt/mqm/java/lib64
```

- By configuring the shell's environment such that the JVM will set up its own `java.library.path`. This path varies by platform and by the location in which you installed IBM WebSphere MQ. For example, for a 64-bit JVM and a default IBM WebSphere MQ installation location, you can use the following settings:

AIX

```
export LIBPATH=/usr/mqm/java/lib64:$LIBPATH
```

Solaris

Linux

HP-UX

```
export LD_LIBRARY_PATH=/opt/mqm/java/lib64:$LD_LIBRARY_PATH
```

Windows

```
set PATH=C:\Program Files\IBM\MQ\java\lib64;%PATH%
```

An example of the exception stack that you see when the environment has not been configured correctly is as follows:

```
Caused by: com.ibm.mq.jmqi.local.LocalMQ$4: CC=2;RC=2495;
AMQ8598: Failed to load the WebSphere MQ native JNI library: 'mqjbnd'.
    at com.ibm.mq.jmqi.local.LocalMQ.loadLib(LocalMQ.java:1268)
    at com.ibm.mq.jmqi.local.LocalMQ$1.run(LocalMQ.java:309)
    at java.security.AccessController.doPrivileged(AccessController.java:400)
    at com.ibm.mq.jmqi.local.LocalMQ.initialise_inner(LocalMQ.java:259)
    at com.ibm.mq.jmqi.local.LocalMQ.initialise(LocalMQ.java:221)
    at com.ibm.mq.jmqi.local.LocalMQ.<init>(LocalMQ.java:1350)
    at com.ibm.mq.jmqi.local.LocalServer.<init>(LocalServer.java:230)
    at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method)
    at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:86)
    at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:58)
    at java.lang.reflect.Constructor.newInstance(Constructor.java:542)
    at com.ibm.mq.jmqi.JmqiEnvironment.getInstance(JmqiEnvironment.java:706)
```

```

        at com.ibm.mq.jmqi.JmqiEnvironment.getMqi(JmqiEnvironment.java:640)
        at
com.ibm.msg.client.wmq.factories.WMQConnectionFactory.createV7ProviderConnection(WMQConnectionFactory.java:8437)
        ... 7 more
Caused by: java.lang.UnsatisfiedLinkError: mqjbn (Not found in java.library.path)
        at java.lang.ClassLoader.loadLibraryWithPath(ClassLoader.java:1235)
        at java.lang.ClassLoader.loadLibraryWithClassLoader(ClassLoader.java:1205)
        at java.lang.System.loadLibrary(System.java:534)
        at com.ibm.mq.jmqi.local.LocalMQ.loadLib(LocalMQ.java:1240)
        ... 20 more

```

2. After either the 32-bit or 64-bit environment has been set up, start the IBM WebSphere MQ classes for JMS application using the command:

```
java application-name
```

where *application-name* is the name of the IBM WebSphere MQ classes for JMS application to be run.

An exception containing IBM WebSphere MQ Reason Code 2495 (MQRC_MODULE_NOT_FOUND) is thrown by the IBM WebSphere MQ classes for JMS if:

- The IBM WebSphere MQ classes for JMS application is run in a 32-bit Java runtime environment, and a 64-bit environment has been set up for the IBM WebSphere MQ classes for JMS, as the 32-bit Java runtime environment is unable to load the 64-bit Java Native Library.
- The IBM WebSphere MQ classes for JMS application is run in a 64-bit Java runtime environment, and a 32-bit environment has been set up for the IBM WebSphere MQ classes for JMS, as the 64-bit Java runtime environment is unable to load the 32-bit Java Native Library.

The IBM WebSphere MQ classes for JMS configuration file

A WebSphere MQ classes for JMS configuration file specifies properties that are used to configure WebSphere MQ classes for JMS.

The format of a WebSphere MQ classes for JMS configuration file is that of a standard Java properties file. A sample configuration file called `jms.config` is supplied in the `bin` subdirectory of the WebSphere MQ classes for JMS installation directory. This file documents all the supported properties and their default values.

You can choose the name and location of a WebSphere MQ classes for JMS configuration file. When you start your application, use a **java** command with the following format:

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

In the command, *config_file_url* is a uniform resource locator (URL) that specifies the name and location of the WebSphere MQ classes for JMS configuration file. URLs of the following types are supported: `http`, `file`, 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 WebSphere MQ classes for JMS configuration file as the file `D:\mydir\mjms.config` on the local Windows system.

When an application starts, WebSphere MQ classes for JMS reads the contents of the configuration file and stores the specified properties in an internal property store. If the **java** command does not identify a configuration file, or if the configuration file cannot be found, WebSphere MQ classes for JMS uses the default values for all the properties. If required, you can override any property in the configuration file by specifying it as a system property on the **java** command.

A WebSphere MQ classes for JMS configuration file can be used with any of the supported transports between an application and a queue manager or broker.

Note that you cannot specify startup trace by setting a property in the WebSphere MQ classes for JMS configuration file. You can specify startup trace only by setting a system property on the **java** command, as shown in the following example:

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

Overriding properties specified in a WebSphere MQ MQI client configuration file

A WebSphere MQ MQI client configuration file can also specify properties that are used to configure WebSphere MQ classes for JMS. However, properties specified in a WebSphere MQ MQI client configuration file apply only when an application connects to a queue manager in client mode.

If required, you can override any attribute in a WebSphere MQ MQI client configuration file by specifying it as a property in a WebSphere MQ classes for JMS configuration file. To override an attribute in a WebSphere MQ MQI client configuration file, use an entry with the following format in the WebSphere MQ classes for JMS configuration file:

```
com.ibm.mq.cfg.stanza.propName=propValue
```

The variables in the entry have the following meanings:

stanza

The name of the stanza in the WebSphere MQ MQI client configuration file that contains the attribute

propName

The name of the attribute as specified in the WebSphere MQ MQI client configuration file

propValue

The value of the property that overrides the value of the attribute specified in the WebSphere MQ MQI client configuration file

Alternatively, you can override an attribute in a WebSphere MQ MQI client configuration file by specifying the property as a system property on the **java** command. Use the preceding format to specify the property as a system property.

Only the following attributes in a WebSphere MQ MQI client configuration file are relevant to WebSphere MQ classes for JMS. If you specify or override other attributes, it has no effect.

Stanza	Attribute
ClientExitPath stanza of the client configuration file	ExitsDefaultPath
ClientExitPath stanza of the client configuration file	ExitsDefaultPath64
ClientExitPath stanza of the client configuration file	JavaExitsClasspath
MessageBuffer stanza of the client configuration file	MaximumSize
MessageBuffer stanza of the client configuration file	PurgeTime
MessageBuffer stanza of the client configuration file	UpdatePercentage
TCP stanza of the client configuration file	CIntRcvBufSize
TCP stanza of the client configuration file	CIntSndBufSize
TCP stanza of the client configuration file	Connect_Timeout
TCP stanza of the client configuration file	KeepAlive

IBM WebSphere MQ classes for JMS and software management tools

Software management tools such as Apache Maven can be used with the IBM WebSphere MQ classes for JMS.

Many large development organizations use these tools to centrally manage repositories of third-party libraries.

The IBM WebSphere MQ classes for JMS are composed of a number of JAR files. When you are developing Java language applications by using this API, an installation of an IBM WebSphere MQ Server, Client, or Client SupportPac is required on the machine where the application is being developed.

If you want to use such a tool and add the JAR files that make up the IBM WebSphere MQ classes for JMS to a centrally managed repository, the following points must be observed:

- A repository or container must be made available only to developers within your organization. Any distribution outside of the organization is not permitted.
- The repository needs to contain a complete and consistent set of JAR files from a single IBM WebSphere MQ release or Fix Pack.
- You are responsible for updating the repository with any maintenance provided by IBM Support.


For IBM WebSphere MQ Version 7.5, the following JAR files need to be installed into the repository:

- com.ibm.mqjms.jar.
- com.ibm.mq.jar
- com.ibm.mq.jmqi.jar
- com.ibm.mq.pcf.jar
- com.ibm.mq.headers.jar
- CL3Export.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- CL3Nonexport.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- jndi.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- ldap.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- rmm.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- dnhbcore.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- jms.jar is required if you are using the IBM WebSphere MQ classes for JMS.
- fscontext.jar is required if you are using the IBM WebSphere MQ classes for JMS and accessing JMS administered objects that are stored in a file system JNDI context.
- providerutil.jar if you are using the IBM WebSphere MQ classes for JMS and accessing JMS administered objects that are stored in a file system JNDI context.

Running WebSphere MQ classes for JMS applications under the Java security manager

WebSphere MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the security manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

The simplest way to do this is to change the policy configuration file supplied with your JRE. On most systems, this file is stored in the path `lib/security/java.policy`, relative to your JRE directory. You can edit the policy configuration file using your preferred editor or the `policytool` program supplied with your JRE.

Important:  Wherever possible, the term *allowlist* has replaced the term *whitelist*. One exception is the following Java system property names.

If you use the Java security manager mechanism with your application, you must grant the following permissions:

- FilePermission on any allowlist file that you use, with read permission for ENFORCEMENT mode, write permission for DISCOVER mode.
- PropertyPermission (read) on the com.ibm.mq.jms.whitelist, com.ibm.mq.jms.whitelist.discover, and com.ibm.mq.jms.whitelist.mode properties.

ClassName allowlisting is supported with [APAR IT14385](#) and IBM WebSphere MQ Version 7.5.0, Fix Pack 8. For more information, see [“ClassName allowlisting in JMS ObjectMessage”](#) on page 699.

Here is an example of two entries in a policy configuration file that allow WebSphere MQ classes for JMS to run successfully under the default security manager:

```
grant codeBase "file:MQ_INSTALLATION_PATH/java/lib/com.ibm.mq.jmqi.jar" {
    //Required
    permission java.util.PropertyPermission "user.name","read";
    permission java.util.PropertyPermission "os.name","read";
    //Required if mqclient.ini/mqs.ini configuration files are used
    permission java.io.FilePermission "/var/mqm/mqclient.ini","read";
    permission java.io.FilePermission "/var/mqm/mqs.ini","read";
    //For the client transport type.
    permission java.net.SocketPermission "*","connect";
    //For the bindings transport type.
    permission java.lang.RuntimePermission "loadLibrary.*";
    //For applications that use CCDT tables (access to the CCDT
    AMQCLCHL.TAB)
    permission java.io.FilePermission
    "/var/mqm/qmgrs/QMGR/@ipcc/AMQCLCHL.TAB","read";
    //For applications that use User Exits
    permission java.io.FilePermission "/var/mqm/exits/*","read";
    permission java.lang.RuntimePermission "createClassLoader";
    //Required for the z/OS platform
    permission java.util.PropertyPermission
    "com.ibm.vm.bitmode","read";
};
grant codeBase "file:MQ_INSTALLATION_PATH/java/lib/com.ibm.mqjms.jar" {
    permission java.util.PropertyPermission "user.name","read";
    permission java.util.PropertyPermission "os.name","read";
    permission java.util.PropertyPermission "console.encoding","read";
    permission java.lang.RuntimePermission "setContextClassLoader";
    //tracing permissions
    permission java.util.PropertyPermission "com.ibm.msg.client.commonservices.*","read";
    permission java.util.PropertyPermission "MQJMS_TRACE_LEVEL","read";
    permission java.util.logging.LoggingPermission "control";
    //Wherever trace output is expected
    permission java.io.FilePermission "/tmp/*","read,write";
    //Required for the z/OS platform
    permission java.util.PropertyPermission
    "com.ibm.vm.bitmode","read";
};
```

Notes:

- `MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.
- The first `grant` statement contains the permissions required by WebSphere MQ classes for JMS, and the second `grant` statement contains the permissions required by a WebSphere MQ classes for JMS application.
- To allow WebSphere MQ classes for JMS to access the Java archive (JAR) files of an application, add the following permission to the first `grant` statement:

```
permission java.io.FilePermission "/path_to_your_app/-", "read";
```

- To use these `grant` statements in your policy configuration file, you might need to modify the path names depending on where you have installed WebSphere MQ classes for JMS and where you store your applications.
- The sample applications supplied with WebSphere MQ classes for JMS, and scripts to run them, do not enable the security manager.

V7.5.0.8 **ClassName allowlisting in JMS ObjectMessage**

V7.5.0.8 In WebSphere MQ classes for JMS, support for allowlisting of classes in the implementation of the JMS ObjectMessage interface provides a potential mitigation against some of the security risks that potentially relate to the Java object serialization and deserialization mechanism.

Note: Wherever possible, the term *allowlist* has replaced the term *whitelist*. One exception is the Java system property names mentioned in this topic.

The Java object serialization and deserialization mechanism has been identified as a potential security risk because deserialization instantiates arbitrary Java objects, where there is the potential for maliciously sent data to cause various problems. One notable application of serialization is in Java Message Service (JMS) ObjectMessages that use serialization to encapsulate and transfer arbitrary objects.

Serialization allowlisting is a potential mitigation against some of the risks that serialization poses. By explicitly specifying which classes can be encapsulated in, and extracted from, ObjectMessages, allowlisting provides some protection against some serialization risks.

Allowlisting in WebSphere MQ classes for JMS

With [APAR IT14385](#) and IBM WebSphere MQ Version 7.5.0, Fix Pack 8, WebSphere MQ classes for JMS supports allowlisting of classes in the implementation of the JMS ObjectMessage interface. The allowlist defines which Java classes might be serialized with ObjectMessage.setObject() and deserialized with ObjectMessage.getObject().

Attempts to serialize or deserialize an instance of a class not included in the allowlist with ObjectMessage cause a javax.jms.MessageFormatException to be thrown, with a java.io.InvalidClassException as its cause.

Producing the allowlist

Important: WebSphere MQ classes for JMS cannot be distributed with an allowlist. The choice of classes to be transferred by using ObjectMessages is an application design choice and IBM WebSphere MQ cannot pre-empt that.

For this reason, the allowlisting mechanism allows for two modes of operation:

DISCOVERY

In this mode, the mechanism produces a listing of fully qualified class names, reporting all classes that have been observed to be serialized or deserialized in ObjectMessages.

ENFORCEMENT

In this mode, the mechanism enforces allowlisting, rejecting attempts to serialize or deserialize classes that are not in the allowlist.

If you want to use this mechanism, you must initially run in DISCOVERY mode to gather the list of currently serialized and deserialized classes, review the list and use it as a basis for your allowlist. It might even be appropriate to use the list unchanged, but the list must be reviewed first before you decide to do this.

Controlling the allowlisting mechanism

Three system properties are available to control the allowlisting mechanism:

com.ibm.mq.jms.whitelist

This property can be specified in either of the following ways:

- The path name of the file that contains the allowlist, in file URI format (that is, starting with `file:`). In DISCOVERY mode, this file is written to by the allowlisting mechanism. The file must not exist. If the file does exist, the mechanism throws an exception rather than overwrite it. In ENFORCEMENT mode, this file is read by the allowlisting mechanism.

- A comma-separated list of fully qualified class names that constitute the allowlist.

If this property is unset, the allowlist mechanism is inactive.

If you are using a Java security manager, you must ensure that the WebSphere MQ classes for JMS JAR files have read and write access to this file.

com.ibm.mq.jms.whitelist.discover

- If this property is unset or set to false, the allowlist mechanism runs in ENFORCEMENT mode.
- If this property is set to true and the allowlist has been specified as a file URI, the allowlist mechanism runs in DISCOVERY mode.
- If this property is set to true and the allowlist has been specified as a list of class names, the allowlist mechanism throws a suitable exception.
- If this property is set to true and the allowlist has not been specified by using the `com.ibm.mq.jms.whitelist` property, the allowlist mechanism is inactive.
- If this property is set to true and the allowlist file already exists, the allowlist mechanism throws a `java.io.InvalidClassException` and entries are not added to the file.

com.ibm.mq.jms.whitelist.mode

This string property can be specified in any of three ways:

- If this property is set to `SERIALIZE`, then ENFORCEMENT mode performs allowlist validation only on the `ObjectMessage.setObject()` method.
- If this property is set to `DESERIALIZE`, then ENFORCEMENT mode performs allowlist validation only on the `ObjectMessage.getObject()` method.
- If this property is unset, or set to any other value, then ENFORCEMENT mode performs allowlist validation on both the `ObjectMessage.getObject()` and the `ObjectMessage.setObject()` methods.

Format of the allowlist file

These are the main features of the format of the allowlist file:

- The allowlist file is in default platform file encoding with platform-appropriate line-endings.
- **Note:** The file might need conversion if you move it between heterogeneous systems.
- Each non-empty line contains a fully qualified class name. Empty lines are ignored.
- Comments can be included - anything following a '#' character, to the end of the line, is ignored.
- There is a very basic wildcarding mechanism:
 - '*' can be the **last** element of a class name.
 - '*' matches a **single** element of a class name, that is, the class, but no part of the package.

So `com.ibm.mq.*` would match `com.ibm.mq.MQMessage` but not `com.ibm.mq.jmqi.remote.api.RemoteFAP`.

Wildcarding does not work for classes in the default package that is for classes without an explicit package name, so a class name of "" is rejected.

- Badly formatted allowlist files, for example, files that contain an entry such as `com.ibm.mq.*.Message`, where the wildcard is not the last element, result in a `java.lang.IllegalArgumentException` being thrown.
- An empty allowlist file has the effect of totally disabling the use of `ObjectMessage`.

Format of the allowlist as a comma-separated list

The same wildcarding mechanism is available for an allowlist as a comma-separated list.

- The '*' can be expanded by the operating system if specified on a command line or in a shell script or batch file, so it might need special handling.

- The '#' comment character is only applicable when a file is specified. If the allowlist is specified as a comma-separated list of class names, then assuming that the operating system or shell doesn't process it, as it is the default comment character in many UNIX or Linux shells, it is treated as a normal character.

When does allowlisting happen?

Allowlisting is initiated when the application first runs an `ObjectMessage setMessage()` or `getMessage()` method.

The system properties are evaluated, the allowlist file is opened and in ENFORCEMENT mode, the list of allowlisted classes are loaded when the mechanism is initialized. At this point, an entry is written in to the IBM WebSphere MQ JMS log file for the application.

When the mechanism is initialized, its parameters might not be changed. As the time of initialization is not easily predicted as it depends on application behavior. The system property settings and the allowlist file contents should therefore be regarded as fixed from the time that the application is started. Do not change the properties or the contents of the allowlist file while the application is running, as the results are not guaranteed.

Points to consider

The best approach to mitigating the risks intrinsic to the Java serialization mechanism would be to explore alternative approaches to data transfer such as using JSON instead of `ObjectMessage`. Using IBM WebSphere MQ Advanced Message Security (AMS) mechanisms can add further security by ensuring that messages come from trusted sources.

If you use the Java security manager mechanism with your application, you must grant the following permissions:

- `FilePermission` on any allowlist file that you use, with read permission for ENFORCEMENT mode, write permission for DISCOVER mode.
- `PropertyPermission` (read) on the `com.ibm.mq.jms.whitelist`, `com.ibm.mq.jms.whitelist.discover`, and `com.ibm.mq.jms.whitelist.mode` properties.

Related concepts

[“Running WebSphere MQ classes for JMS applications under the Java security manager” on page 697](#)
WebSphere MQ classes for JMS can run with the Java security manager enabled. To run applications successfully with the security manager enabled, you must configure your Java virtual machine (JVM) with a suitable policy configuration file.

The IBM WebSphere MQ resource adapter

The resource adapter allows applications running in an application server to access IBM WebSphere MQ resources. It supports inbound and outbound communication.

The Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) provides a standard way of connecting applications running in a Java EE environment to an Enterprise Information System (EIS) such as IBM WebSphere MQ or Db2. The IBM WebSphere MQ resource adapter implements the JCA 1.5 interfaces and contains the IBM WebSphere MQ classes for JMS. It allows JMS applications and message driven beans (MDBs), running in an application server, to access the resources of a IBM WebSphere MQ queue manager. The resource adapter supports both the point-to-point domain and the publish/subscribe domain.

The IBM WebSphere MQ resource adapter supports two types of communication between an application and a queue manager:

Outbound communication

An application starts a connection to a queue manager, and then sends JMS messages to JMS destinations and receives JMS messages from JMS destinations in a synchronous manner.

Inbound communication

A JMS message arriving at a JMS destination is delivered to an MDB, which processes the message asynchronously.

For more information on the IBM WebSphere MQ classes for JMS, see [“Using WebSphere MQ classes for JMS” on page 684](#).

The resource adapter also contains the IBM WebSphere MQ classes for Java. The classes are automatically available to applications running in an application server that the resource adapter has been deployed into, and allow applications running in that application server to use the IBM WebSphere MQ classes for Java API when accessing resources of a IBM WebSphere MQ queue manager. For more information about the IBM WebSphere MQ classes for Java, see [“Using WebSphere MQ classes for Java” on page 625](#).

The use of the IBM WebSphere MQ classes for Java within a Java EE environment is supported with restrictions. For information about these restrictions, see [“Running IBM WebSphere MQ classes for Java applications within Java platform Enterprise Edition” on page 682](#).

Other required documentation to support a JCA resource adapter

Refer to the documentation for your application server for information about how to configure a JCA resource adapter.

Every application server provides its own set of administration interfaces. Some application servers provide graphical user interfaces to define JCA resources, but others require the administrator to write XML deployment plans. It is therefore beyond the scope of this documentation to provide information about how to configure the WebSphere MQ resource adapter for each application server. This documentation focuses only on what you need to configure. Refer to documentation supplied with your application server for information about how to configure a JCA resource adapter.

To understand this documentation, you must be familiar with JMS and WebSphere MQ classes for JMS. Many of the properties used to configure the WebSphere MQ resource adapter are equivalent to properties of WebSphere MQ classes for JMS objects and have the same function.

Installation of the WebSphere MQ resource adapter

The WebSphere MQ resource adapter is supplied as a resource archive (RAR) file. Install the RAR file in your application server. You might need to add directories to the system path.

The WebSphere MQ resource adapter is supplied as a resource archive (RAR) file called `wmq.jmsra.rar`. This file is installed with WebSphere MQ classes for JMS in the directory shown in [Table 94 on page 702](#).

Table 94. The directory containing <code>wmq.jmsra.rar</code> for each platform	
Platform	Directory
AIX, HP-UX, Linux , and Solaris	<code>MQ_INSTALLATION_PATH /java/lib/jca</code>
IBM i	<code>/QIBM/ProdData/mqm/java/lib/jca</code>
Windows	<code>MQ_INSTALLATION_PATH \java\lib\jca</code>
<code>MQ_INSTALLATION_PATH</code> represents the high-level directory in which WebSphere MQ is installed.	

The RAR file contains WebSphere MQ classes for JMS and the WebSphere MQ implementation of the JCA interfaces.

You must install the WebSphere MQ resource adapter RAR file in your application server, but the way you do this depends on the application server. See the documentation for your application server for information about how to install a resource adapter RAR file.

For bindings connections on UNIX and Linux systems, you must ensure that the directory containing the Java Native Interface (JNI) libraries is in the system path. For the location of this directory, which also contains the WebSphere MQ classes for JMS libraries, see [Table 93 on page 693](#). On Windows, this

directory is automatically added to the system path during the installation of WebSphere MQ classes for JMS .

Transactions are supported in both client and bindings mode.

The WebSphere MQ resource adapter and the version of WebSphere MQ classes for JMS used by the resource adapter must be at the same release level.

WebSphere Application Server and the WebSphere MQ resource adapter

Do not use the WebSphere MQ resource adapter with WebSphere Application Server Version 6.

WebSphere Application Server, V7, includes a version of the WebSphere MQ V7 Resource Adapter.

Do not use the WebSphere MQ resource adapter within WebSphere Application Server, V6. To access the resources of a WebSphere MQ queue manager from within WebSphere Application Server from a JMS application, use the WebSphere MQ messaging provider. The WebSphere MQ messaging provider contains a version of the WebSphere MQ classes for JMS.

WebSphere Application Server, V7, includes a version of the WebSphere MQ V7 Resource Adapter.

For more information, see the technote [Which version of WebSphere MQ Resource Adapter \(RA\) is shipped with WebSphere Application Server ?](#).

WebSphere Application Server Liberty and the IBM WebSphere MQ resource adapter

The IBM WebSphere MQ Version 7.5 resource adapter can be installed into WebSphere Application Server Liberty Version 8.5.5, Fix Pack 2 or later, by using the `wmqJmsClient-1.1` feature. Alternatively you can, subject to some restrictions, install the resource adapter by using generic Java Platform, Enterprise Edition Connector Architecture (Java EE JCA) support.

General restrictions when installing the resource adapter into Liberty

The following restrictions apply to the Version 7.5 resource adapter when using the `wmqJmsClient-1.1` feature and also when using generic JCA support:

- The IBM WebSphere MQ classes for Java are not supported in Liberty. They must not be used with either the IBM WebSphere MQ Liberty messaging feature or with the generic JCA support. For more information, see [Using WebSphere MQ Java Interfaces in J2EE/JEE Environments](#).
- The IBM WebSphere MQ resource adapter has a transport type of `BINDINGS_THEN_CLIENT`. This transport type is not supported within the IBM WebSphere MQ Liberty messaging feature.
- The IBM WebSphere MQ Advanced Message Security (IBM WebSphere MQ AMS) feature is not included in the IBM WebSphere MQ Liberty messaging feature.

The IBM WebSphere MQ Version 7.5 resource adapter cannot be used with the `wmqJmsClient-2.0` feature.

Configuration of the WebSphere MQ resource adapter

To configure the WebSphere MQ resource adapter, you define various JCA resources and system properties.

Define JCA resources in the following categories:

- The properties of the `ResourceAdapter` object, which represent the global properties of the resource adapter, such as the level of diagnostic tracing. These properties are described in [“Configuration of the ResourceAdapter object”](#) on page 704.
- The properties of an `ActivationSpec` object, which determine how an MDB is activated for inbound communication. These properties are described in [“Configuring the resource adapter for inbound communication”](#) on page 706.
- The properties of a `ConnectionFactory` object, which the application server uses to create a JMS `ConnectionFactory` object for outbound communication. These properties are described in [“Configuring the resource adapter for outbound communication”](#) on page 718.

- The properties of an administered destination object, which the application server uses to create a JMS Queue object or JMS Topic object for outbound communication. These properties are also described in [“Configuring the resource adapter for outbound communication”](#) on page 718.

The WebSphere MQ resource adapter RAR file contains a file called META-INF/ra.xml, which contains a deployment descriptor for the resource adapter. This deployment descriptor is defined by the XML schema at https://java.sun.com/xml/ns/j2ee/connector_1_5.xsd and contains information about the resource adapter and the services that it provides. An application server might also require a deployment plan for the resource adapter. This deployment plan is specific to the application server. For example, WebSphere Application Server Community Edition requires a deployment plan called geronimo-ra.xml.

If you are using Secure Sockets Layer (SSL), specify the locations of the keystore file and truststore file as JVM system properties, as in the following example:

```
java ... -Djavax.net.ssl.keyStore=
key_store_location
-Djavax.net.ssl.trustStore=trust_store_location
-Djavax.net.ssl.keyStorePassword=key_store_password
```

These properties cannot be properties of an ActivationSpec or ConnectionFactory object, and you cannot specify more than one keystore for an application server. The properties apply to the whole JVM, and might therefore affect the application server if other applications, running in the application server, are using SSL connections. The application server might also reset these properties to different values. For more information about using SSL with WebSphere MQ classes for JMS, see [“Using Secure Sockets Layer \(SSL\) with WebSphere MQ classes for JMS”](#) on page 867.

An installation verification test (IVT) program is supplied with the WebSphere MQ resource adapter, but you must configure the resource adapter before you can run the program. For information about what you need to configure in order to run the IVT program, see [“The installation verification test program for the WebSphere MQ resource adapter”](#) on page 747.

The resource adapter logs, warning, and error messages use the same mechanism as the IBM WebSphere MQ classes for JMS, for details, see [“Logging and IBM WebSphere MQ classes for JMS”](#) on page 761. For WebSphere Application Server, these messages are automatically redirected to the application server output log. For other application servers, such as WAS CE and JBoss, these will, by default, go to a file called mqjms.log. To configure the resource adapter to additionally log warning messages to your application servers standard output log, set the following JVM system property for your application server:

```
-Dcom.ibm.msg.client.commonservices.log.outputName=mqjms.log,stdout
```

For details on how to set a JVM system property, see your application server documentation.

Configuration of the ResourceAdapter object

The ResourceAdapter object encapsulates the global properties of the WebSphere MQ resource adapter. Define these properties using the facilities of your resource adapter.

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.

For more information about defining properties associated with diagnostic trace, see [Tracing the IBM WebSphere MQ resource adapter](#).

The resource adapter manages an internal connection pool of JMS connections that are used to deliver messages to MDBs. [Table 95 on page 705](#) lists the properties of the ResourceAdapter object that are associated with the connection pool.

Table 95. Properties of the ResourceAdapter object that are associated with the connection pool

Name of property	Type	Default value	Description
maxConnections	String	50	The maximum number of connections to a WebSphere MQ queue manager and the maximum number of MDBs deployed.
connectionConcurrency	String	1	The maximum number of MDBs to share a JMS connection. Sharing connections is not possible and this property always has the value 1.
reconnectionRetryCount	String	5	The maximum number of attempts made by the resource adapter to reconnect to a WebSphere MQ queue manager if a connection fails.
reconnectionRetryInterval	String	300 000	The time, in milliseconds, that the resource adapter waits before trying to reconnect to a WebSphere MQ queue manager.
startupRetryCount	String	0	The default number of times to try and connect a MDB on startup, if the queue manager is not running when the application server is started.
startupRetryInterval	String	30 000	The default sleep time between startup connection attempts (in milliseconds).

When an MDB is deployed in the application server, a new JMS connection is created and a conversation started with the queue manager, provided the maximum number of connections specified by the maxConnection property is not exceeded. The maximum number of MDBs therefore equals the maximum number of connections. If the number of deployed MDBs reaches this maximum, any attempt to deploy another MDB fails. If an MDB is stopped, its connection can be used by another MDB.

In general, if many MDBs are to be deployed, you must increase the value of the maxConnections property.

The reconnectionRetryCount and reconnectionRetryInterval properties govern the behavior of the resource adapter when connections to a WebSphere MQ queue manager fail, because of a network failure for example. When a connection fails, the resource adapter suspends the delivery of messages to all MDBs supplied by that connection for an interval specified by the reconnectionRetryInterval property. The resource adapter then attempts to reconnect to the queue manager. If the attempt fails, the resource adapter makes further attempts to reconnect at intervals specified by the reconnectionRetryInterval property until the limit imposed by the reconnectionRetryCount property is reached. If all attempts fail, delivery is stopped permanently until the MDBs are restarted manually.

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 WebSphere 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=6
```

You do not need to define all the properties of the ResourceAdapter object. Any properties left 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.

Configuring the resource adapter for inbound communication

To configure inbound communication, define the properties of one or more ActivationSpec objects.

The properties of an ActivationSpec object determine how a message driven bean (MDB) receives JMS messages from a WebSphere MQ queue. The transactional behavior of the MDB is defined in its deployment descriptor.

An ActivationSpec object has two sets of properties:

- Properties that are used to create a JMS connection to a WebSphere MQ queue manager
- Properties that are used to create a JMS connection consumer that delivers messages asynchronously as they arrive on a specified queue

The way in which you define the properties of an ActivationSpec object depends on the administration interfaces provided by your application server.

Table 96 on page 706 lists the properties of an ActivationSpec object that are used to create a JMS connection to a WebSphere MQ queue manager.

<i>Table 96. Properties of an ActivationSpec object that are used to create a JMS connection</i>			
Name of property	Type	Valid values (default value in bold)	Description
applicationName	String	<ul style="list-style-type: none"> • The invoking class name, if it is available, adjusted to be no longer than 28 characters. If it is not available, the string WebSphere MQ Client for Java is used. 	The name by which an application is registered with the queue manager. This application name is shown by the DISPLAY CONN MQSC/PCF command (where the field is called APPLTAG) or in the IBM WebSphere MQ Explorer Application Connections display (where the field is called App name).
brokerCCDurSubQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE • A queue name 	The name of the queue from which a connection consumer receives durable subscription messages
brokerCCSubQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE • A queue name 	The name of the queue from which a connection consumer receives nondurable subscription messages
brokerControlQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.BROKER.CONTROL.QUEUE • A queue name 	The name of the broker control queue
brokerQueueManager ¹	String	<ul style="list-style-type: none"> • "" (empty string) • A queue manager name 	The name of the queue manager on which the broker is running
brokerSubQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.JMS.ND.SUBSCRIBER.QUEUE • A queue name 	The name of the queue from which a nondurable message consumer receives messages

Table 96. Properties of an *ActivationSpec* object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerVersion ¹	String	<ul style="list-style-type: none"> • unspecified - After the broker is migrated from V6 to V7, set this property so that RFH2 headers are no longer used. After migration, this property is no longer relevant. • V1 - To use a WebSphere MQ publish/subscribe broker. Or to use a broker of WebSphere MQ Integrator, WebSphere Event Broker, WebSphere Business Integration Event Broker, or WebSphere Business Integration Message Broker in compatibility mode. This value is the default value if TRANSPORT is set to BIND or CLIENT. • V2 - To use a broker of WebSphere MQ Integrator, WebSphere Event Broker, WebSphere Business Integration Event Broker, or WebSphere Business Integration Message Broker in native mode. This value is the default value if TRANSPORT is set to DIRECT or DIRECTHTTP. 	The version of the broker being used
ccdtURL	String	<ul style="list-style-type: none"> • null • A uniform resource locator (URL) 	A URL that identifies the name and location of the file containing the client channel definition table (CCDT) and specifies how the file can be accessed
CCSID	String	<ul style="list-style-type: none"> • 819 • A coded character set identifier supported by the Java virtual machine (JVM) 	The coded character set identifier for a connection
channel	String	<ul style="list-style-type: none"> • SYSTEM.DEF.SVRCONN • The name of an MQI channel 	The name of the MQI channel to use
cleanupInterval ¹	int	<ul style="list-style-type: none"> • 3 600 000 • A positive integer 	The interval, in milliseconds, between background runs of the publish/subscribe cleanup utility
cleanupLevel ¹	String	<ul style="list-style-type: none"> • SAFE • NONE • STRONG • FORCE • NONDUR 	The cleanup level for a broker-based subscription store
clientID	String	<ul style="list-style-type: none"> • null • A client identifier 	The client identifier for a connection

Table 96. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
cloneSupport	String	<ul style="list-style-type: none"> • DISABLED - Only one instance of a durable topic subscriber can run at a time. • ENABLED - Two or more instances of the same durable topic subscriber can run simultaneously, but each instance must run in a separate Java virtual machine (JVM). 	Whether two or more instances of the same durable topic subscriber can run simultaneously
connectionNameList	String	<ul style="list-style-type: none"> • localhost(1414) • A string composed of items separated by commas where each item takes the format: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <code>HOSTNAME (PORT)</code> </div> where <i>HOSTNAME</i> is either a DNS name or an IP address. 	<p>A list of TCP/IP connection names used for inbound communications.</p> <p>When specified, connectionNameList supersedes the hostname and port properties.</p> <p>This property is used to reconnect to multi-instance queue managers.</p> <p>connectionNameList is similar in form to localAddress, but must not be confused with it. localAddress specifies the characteristics of the local communications, whereas connectionNameList specifies how to reach a remote queue manager.</p>
failIfQuiesce	Boolean	<ul style="list-style-type: none"> • true • false 	Whether calls to certain methods fail if the queue manager is in a quiescing state
headerCompression	String	<ul style="list-style-type: none"> • NONE • SYSTEM - RLE message header compression is performed 	A list of the techniques that can be used for compressing header data on a connection
hostName	String	<ul style="list-style-type: none"> • localhost • A host name • An IP address 	<p>The host name or IP address of the system on which the queue manager resides.</p> <p>The hostname and port properties are superseded by the connectionNameList property when it is specified.</p>

Table 96. Properties of an *ActivationSpec* object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
localAddress	String	<ul style="list-style-type: none"> • null • A string in the format: <div>[host_name][(low_port[,high_port])]</div> where <i>host_name</i> is a host name or IP address, <i>low_port</i> and <i>high_port</i> are TCP port numbers, and brackets denote an optional component 	<p>For a connection to a queue manager, this property specifies either or both of the following things:</p> <ul style="list-style-type: none"> • The local network interface to be used • The local port, or range of local ports, to be used <p>localAddress is similar in form to connectionNameList, but must not be confused with it. localAddress specifies the characteristics of the local communications, whereas connectionNameList specifies how to reach a remote queue manager.</p>
messageCompression	String	<ul style="list-style-type: none"> • NONE • A list of one or more of the following values separated by blank characters: RLE ZLIBFAST ZLIBHIGH 	A list of the techniques that can be used for compressing message data on a connection
messageRetention ¹	Boolean	<ul style="list-style-type: none"> • true - Unwanted messages remain on the input queue • false - Unwanted messages are dealt with according to their disposition options 	Whether the connection consumer keeps unwanted messages on the input queue
messageSelection ¹	String	<ul style="list-style-type: none"> • CLIENT • BROKER 	Determines whether message selection is done by WebSphere MQ classes for JMS or by the broker. Message selection by the broker is not supported when <i>brokerVersion</i> has the value 1.
password	String	<ul style="list-style-type: none"> • null • A password 	The default password to use when creating a connection to the queue manager

Table 96. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
pollingInterval ¹	int	<ul style="list-style-type: none"> • 5000 • Any positive integer 	If each message listener within a session has no suitable message on its queue, this value is the maximum interval, in milliseconds, that elapses before each message listener tries again to get a message from its queue. If it frequently happens that no suitable message is available for any of the message listeners in a session, consider increasing the value of this property. This property is relevant only if TRANSPORT has the value BIND or CLIENT.
port	int	<ul style="list-style-type: none"> • 1414 • A TCP port number 	<p>The port on which the queue manager listens.</p> <p>The hostname and port properties are superseded by the connectionNameList property when it is specified.</p>
providerVersion	string	<ul style="list-style-type: none"> • unspecified • A string in one of the following formats <ul style="list-style-type: none"> – V.R.M.F – V.R.M – V.R – V <p>where V, R, M, and F are integer values greater than or equal to zero.</p>	The version, release, modification level and fix pack of the queue manager to which the MDB intends to connect.
queueManager	String	<ul style="list-style-type: none"> • "" (empty string) • A queue manager name 	The name of the queue manager to connect to
receiveExit ³	String	<ul style="list-style-type: none"> • null • A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the WebSphere MQ classes for Java interface, MQReceiveExit 	Identifies a channel receive exit program, or a sequence of receive exit programs to be run in succession
receiveExitInit	String	<ul style="list-style-type: none"> • null • A string comprising one or more items of user data separated by commas 	The user data that is passed to channel receive exit programs when they are called

Table 96. Properties of an *ActivationSpec* object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
rescanInterval ¹	int	<ul style="list-style-type: none"> • 5000 • Any positive integer 	When a message consumer in the point-to-point domain uses a message selector to select which messages it wants to receive, WebSphere MQ classes for JMS searches the WebSphere MQ queue for suitable messages in the sequence determined by the <i>MsgDeliverySequence</i> attribute of the queue. When WebSphere MQ classes for JMS finds a suitable message and delivers it to the consumer, WebSphere MQ classes for JMS resumes the search for the next suitable message from its current position in the queue. WebSphere MQ classes for JMS continues to search the queue in this way until it reaches the end of the queue, or until the interval of time in milliseconds, as determined by the value of this property, has expired. In each case, WebSphere MQ classes for JMS returns to the beginning of the queue to continue its search, and a new time interval commences.
securityExit ³	String	<ul style="list-style-type: none"> • null • The fully qualified name of a class that implements the WebSphere MQ classes for Java interface, <i>MQSecurityExit</i> 	Identifies a channel security exit program
securityExitInit	String	<ul style="list-style-type: none"> • null • A string of user data 	The user data that is passed to a channel security exit program when it is called
sendExit ³	String	<ul style="list-style-type: none"> • null • A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the WebSphere MQ classes for Java interface, <i>MQSendExit</i> 	Identifies a channel send exit program, or a sequence of send exit programs to be run in succession

Table 96. Properties of an *ActivationSpec* object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
sendExitInit	String	<ul style="list-style-type: none"> null A string comprising one or more items of user data separated by commas 	The user data that is passed to channel send exit programs when they are called
shareConvAllowed	Boolean	<ul style="list-style-type: none"> NO - A client connection cannot share its socket. YES - A client connection can share its socket. 	Whether a client connection can share its socket with other top-level JMS connections from the same process to the same queue manager, if the channel definitions match
sparseSubscriptions ¹	Boolean	<ul style="list-style-type: none"> false - Subscriptions receive frequent matching messages. true - Subscriptions receive infrequent matching messages. This value requires that the subscription queue can be opened for browse. 	Controls the message retrieval policy of a TopicSubscriber object
sslCertStores	String	<ul style="list-style-type: none"> null A string of one or more LDAP URLs separated by blanks. Each LDAP URL has the format: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <code>ldap://host_name[:port]</code> </div> where <i>host_name</i> is a host name or IP address, <i>port</i> is a TCP port number, and brackets denote an optional component. 	The Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs) for use on an SSL connection
sslCipherSuite	String	<ul style="list-style-type: none"> null The name of a CipherSuite 	The CipherSuite to use for an SSL connection
sslFipsRequired ²	Boolean	<ul style="list-style-type: none"> false true 	Whether an SSL connection must use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSSEFIPS)
sslPeerName	String	<ul style="list-style-type: none"> null A template for distinguished names 	For an SSL connection, a template that is used to check the distinguished name in the digital certificate provided by the queue manager
sslResetCount	int	<ul style="list-style-type: none"> 0 An integer in the range 0 - 999 999 999 	The total number bytes sent and received by an SSL connection before the secret keys used by SSL are renegotiated

Table 96. Properties of an ActivationSpec object that are used to create a JMS connection (continued)

Name of property	Type	Valid values (default value in bold)	Description
sslSocketFactory	String	A string representing the fully qualified class name of a class providing an implementation of the javax.net.ssl.SSLSocketFactory interface. Optionally including an argument to be passed to the constructor method, enclosed in parentheses.	Any connections established in the scope of the administered object use sockets obtained from this implementation of the SSLSocketFactory interface.
statusRefreshInterval ¹	int	<ul style="list-style-type: none"> • 60000 • Any positive integer 	The interval, in milliseconds, between refreshes of the long running transaction that detects when a subscriber loses its connection to the queue manager. This property is relevant only if subscriptionStore has the value QUEUE.
subscriptionStore ¹	String	<ul style="list-style-type: none"> • BROKER • MIGRATE • QUEUE 	Determines where WebSphere MQ classes for JMS stores persistent data about active subscriptions
transportType	String	<ul style="list-style-type: none"> • CLIENT • BINDINGS • BINDINGS_THEN_CLIENT 	Whether a connection to a queue manager uses client mode or bindings mode. If the value BINDINGS_THEN_CLIENT is specified, the resource adapter first tries to make a connection in bindings mode. If this connection attempt fails then the resource adapter tries to make a client mode connection.
username	String	<ul style="list-style-type: none"> • null • A user name 	The default user name to use when creating a connection to a queue manager
wildcardFormat	String	<ul style="list-style-type: none"> • CHAR- Recognizes character wildcards only, as used in broker version 1 • TOPIC - Recognizes topic level wildcards only, as used in broker version 2 	Which version of wildcard syntax is to be used

Notes:

1. This property can be used with Version 7.0 of WebSphere MQ classes for JMS. It does not affect an application connected to a Version 7.0 queue manager unless the providerVersion property is set to a version number less than 7.
2. For important information about using the sslFipsRequired property, see [“Limitations of the IBM WebSphere MQ resource adapter”](#) on page 736.

3. For information on how to configure the resource adapter so that it can locate an exit, see [“Configuring IBM WebSphere MQ classes for JMS to use channel exits”](#) on page 873.

Table 97 on page 714 lists the properties of an ActivationSpec object that are used to create a JMS connection consumer.

Table 97. Properties of an ActivationSpec object that are used to create a JMS connection consumer			
Name of property	Type	Valid values (default value in bold)	Description
destination	String	A destination name	The destination from which to receive messages. The useJNDI property determines how the value of this property is interpreted.
destinationType	String	<ul style="list-style-type: none"> • <code>javax.jms.Queue</code> • <code>javax.jms.Topic</code> 	The type of destination, a queue, or a topic
maxMessages	int	<ul style="list-style-type: none"> • 1 • A positive integer 	The maximum number of messages that can be assigned to a server session at one time. If the activation spec is delivering messages to an MDB in an XA transaction, a value of 1 is used regardless of the setting of this property.
maxPoolDepth	int	<ul style="list-style-type: none"> • 10 • A positive integer 	The maximum number of server sessions in the server session pool used by the connection consumer
messageSelector	String	<ul style="list-style-type: none"> • null • An SQL92 message selector expression 	A message selector expression specifying which messages are to be delivered
nonASFTimeout	int	<ul style="list-style-type: none"> • 0 • A positive integer 	<p>A positive value indicates that non-ASF delivery is used. The value is the time, in milliseconds, that a get request waits for messages that might not have yet arrived (a get with wait call). The default value, 0, indicates that ASF delivery is used.</p> <p>This parameter is valid only when the application is running on WebSphere Application Server version 7 or later.</p>
nonASFRollbackEnabled	Boolean	<ul style="list-style-type: none"> • false - The message is consumed even if the MDB fails • true - Failure within the MDB causes the message to rollback to the queue. 	Whether message delivery is within a WebSphere MQ syncpoint if the MDB is non-transacted. Ignored if the MDB is transacted or if nonASFTimeout is set to 0.
poolTimeout	int	<ul style="list-style-type: none"> • 300000 • A positive integer 	The time, in milliseconds, that an unused server session is held open in the server session pool before being closed due to inactivity

Table 97. Properties of an *ActivationSpec* object that are used to create a JMS connection consumer (continued)

Name of property	Type	Valid values (default value in bold)	Description
readAheadAllowed	int	<ul style="list-style-type: none"> • DESTINATION - Determine whether read ahead is allowed by referring to the queue or topic definition. • DISABLED - Read ahead is not allowed. • ENABLED - Read ahead is allowed. • QUEUE - Determine whether read ahead is allowed by referring to the queue definition. • TOPIC - Determine whether read ahead is allowed by referring to the topic definition. 	Whether the MDB is allowed to use read ahead to get nonpersistent messages from the destination into an internal buffer before receiving them
readAheadClosePolicy	int	<ul style="list-style-type: none"> • ALL - All messages in the internal read ahead buffer are delivered to the MDB before it stops. • CURRENT - Only the current MDB invocation completes, potentially leaving messages in the internal read ahead buffer, which are then discarded. 	What happens to messages in the internal read ahead buffer when the MDB is stopped by the administrator.
receiveCCSID	int	<ul style="list-style-type: none"> • 0 - Use JVM <code>Charset.defaultCharset</code> • 1208 - UTF-8 • A supported coded character set identifier 	Destination property that sets the target CCSID for queue manager message conversion. The value is ignored unless receiveConversion is set to QMGR
receiveConversion	String	<ul style="list-style-type: none"> • CLIENT_MSG • QMGR 	Destination property that determines whether data conversion is going to be performed by the queue manager.
startTimeout	int	<ul style="list-style-type: none"> • 10 000 • A positive integer 	The time, in milliseconds, within which delivery of a message to an MDB must start after the work to deliver the message has been scheduled. If this time elapses, the message is rolled back onto the queue.

Table 97. Properties of an ActivationSpec object that are used to create a JMS connection consumer (continued)

Name of property	Type	Valid values (default value in bold)	Description
subscriptionDurability	String	<ul style="list-style-type: none"> • NonDurable - A nondurable subscription is used to deliver messages to an MDB subscribing to the topic. • Durable - A durable subscription is used to deliver messages to an MDB subscribing to the topic. 	Whether a durable or nondurable subscription is used to deliver messages to an MDB subscribing to the topic
subscriptionName	String	<ul style="list-style-type: none"> • "" (empty string) • A subscription name 	The name of the durable subscription
useJNDI	Boolean	<ul style="list-style-type: none"> • false - The property called destination is interpreted as the name of a WebSphere MQ queue or a topic. • true - The property called destination is interpreted as the name of a javax.jms.Queue object or javax.jms.Topic object in the JNDI namespace of the application server. 	Determines how the value of the property called destination is interpreted

The ActivationSpec properties called destination and destinationType must be defined explicitly. All the other properties are optional.

An ActivationSpec object can have conflicting properties. For example, you can specify SSL properties for a connection in bindings mode. In this case, the behavior is determined by the transport type and the messaging domain, which is either point-to-point or publish/subscribe as determined by the destinationType property. Any properties that are not applicable to the specified transport type or messaging domain are ignored.

If you define a property that requires other properties to be defined, but you do not define these other properties, the ActivationSpec object throws an InvalidPropertyException exception when its validate() method is called during the deployment of an MDB. The exception is reported to the administrator of the application server in a manner that depends on the application server. For example, if you set the subscriptionDurability property to Durable, indicating that you want use durable subscriptions, you must also define the subscriptionName property.

If the properties called ccdtURL and channel are both defined, an InvalidPropertyException exception is thrown. However, if you define the ccdtURL property only, leaving the property called channel with its default value of SYSTEM.DEF.SVRCONN, no exception is thrown, and the client channel definition table identified by the ccdtURL property is used to start a JMS connection.

Most of the properties of an ActivationSpec object are equivalent to properties of WebSphere MQ classes for JMS objects or parameters of WebSphere MQ classes for JMS methods. However, three tuning properties, and one usability property, have no equivalents in WebSphere MQ classes for JMS:

startTimeout

The time, in milliseconds, that the work manager of the application server waits for resources to become available after the resource adapter schedules a Work object to deliver a message to an MDB. If this time elapses before delivery of the message starts, the Work object times out, the message is rolled back onto the queue, and the resource adapter can then attempt to deliver the message again. A warning is written to diagnostic trace, if enabled, but does not otherwise affect the process

of delivering messages. You might expect this condition to occur only at times when the application server is experiencing a very high load. If the condition occurs regularly, consider increasing the value of this property to give the work manager longer to schedule message delivery.

maxPoolDepth

The maximum number of server sessions in the server session pool used by a connection consumer. When a server session is created, it starts a conversation with a queue manager. The connection consumer uses a server session to deliver a message to an MDB. A larger pool depth allows more messages to be delivered concurrently in high volume situations, but uses more resources of the application server. If many MDBs are to be deployed, consider making the pool depth smaller in order to maintain the load on the application server at a manageable level. Each connection consumer uses its own server session pool, so that this property does not define the total number of server sessions available to all connection consumers.

poolTimeout

The time, in milliseconds, that an unused server session is held open in the server session pool before being closed due to inactivity. A transient increase in the message workload causes additional server sessions to be created in order to distribute the load but, after the message workload returns to normal, the additional server sessions remain in the pool and are not used.

Every time a server session is used, it is marked with a timestamp. Periodically a scavenger thread checks that each server session has been used within the period specified by this property. If a server session has not been used, it is closed and removed from the server session pool. A server session might not be closed immediately after the specified period has elapsed, this property represents the minimum period of inactivity before removal.

useJNDI

For a description of this property, see [Table 97 on page 714](#).

To deploy an MDB, first define the properties of an ActivationSpec object, specifying the properties that the MDB requires. The following example is a typical set of properties that you might define explicitly:

```
channel:          SYSTEM.DEF.SVRCONN
destination:      SYSTEM.DEFAULT.LOCAL.QUEUE
destinationType:  javax.jms.Queue
hostName:         192.168.0.42
messageSelector:  color='red'
port:            1414
queueManager:     ExampleQM
transportType:    CLIENT
```

The application server uses the properties to create an ActivationSpec object, which is then associated with an MDB. The properties of the ActivationSpec object determine how messages are delivered to the MDB. Deployment of the MDB fails if the MDB requires distributed transactions but the resource adapter does not support distributed transactions. For information about how to install the resource adapter so that distributed transactions are supported, see [“Installation of the WebSphere MQ resource adapter” on page 702](#).

If more than one MDB is receiving messages from the same destination, then a message sent in the point-to-point domain is received by only one MDB, even if other MDBs are eligible to receive the message. In particular, if two MDBs are using different message selectors, and an incoming message matches both message selectors, only one of the MDBs receives the message. The MDB chosen to receive a message is undefined, and you cannot rely on a specific MDB receiving the message. Messages sent in the publish/subscribe domain are received by all eligible MDBs.

Inbound poison message handling in the Resource Adapter

In some circumstances, a message delivered to an MDB might be rolled back onto a WebSphere MQ queue. This roll back can happen, for example, if a message is delivered within a unit of work that is then rolled back. A message that is rolled back is delivered again, but a badly formatted message might repeatedly cause an MDB to fail and therefore cannot be delivered. Such a message is called a poison message. You can configure WebSphere MQ so that WebSphere MQ classes for JMS automatically transfers a poison message to another queue for further investigation or discards the message.

For details on how to handle poison messages, see [“Handling poison messages in IBM WebSphere MQ classes for JMS” on page 850](#).

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](#)

Configuring the resource adapter for outbound communication

To configure outbound communication, define the properties of a `ConnectionFactory` object and an administered destination object.

When using outbound communication, an application running in an application server starts a connection to a queue manager, and then sends messages to its queues and receives messages from its queues in a synchronous manner. For example, the following servlet method, `doGet()`, uses outbound communication:

```
protected void doGet(HttpServletRequest request, HttpServletResponse response)
    throws ServletException, IOException {
    ...
    // Look up ConnectionFactory and Queue objects from the JNDI namespace
    InitialContext ic = new InitialContext();
    ConnectionFactory cf = (javax.jms.ConnectionFactory) ic.lookup("myCF");
    Queue q = (javax.jms.Queue) ic.lookup("myQueue");

    // Create and start a connection
    Connection c = cf.createConnection();
    c.start();

    // Create a session and message producer
    Session s = c.createSession(false, Session.AUTO_ACKNOWLEDGE);
    MessageProducer pr = s.createProducer(q);

    // Create and send a message
    Message m = s.createTextMessage("Hello, World!");
    pr.send(m);

    // Create a message consumer and receive the message just sent
    MessageConsumer co = s.createConsumer(q);
    Message mr = co.receive(5000);

    // Close the connection
    c.close();
}
```

When the servlet receives an HTTP GET request, it retrieves a `ConnectionFactory` object and a `Queue` object from the JNDI namespace, and uses the objects to send a message to a WebSphere MQ queue. The servlet then receives the message that it has sent.

To configure outbound communication, define JCA resources in the following categories:

- The properties of a `ConnectionFactory` object, which the application server uses to create a JMS `ConnectionFactory` object.
- The properties of an administered destination object, which the application server uses to create a JMS `Queue` object or JMS `Topic` object.

The way you define these properties depends on the administration interfaces provided by your application server. `ConnectionFactory`, `Queue`, and `Topic` objects created by the application server are bound into a JNDI namespace from where they can be retrieved by an application.

Typically, you define one `ConnectionFactory` object for each queue manager that applications might need to connect to. You define one `Queue` object for each queue that applications might need to access in the point-to-point domain. And you define one `Topic` object for each topic that applications

might want to publish or subscribe to. A `ConnectionFactory` object can be domain independent. Alternatively, it can be domain-specific, a `QueueConnectionFactory` object for the point-to-point domain or a `TopicConnectionFactory` object for the publish/subscribe domain.

Table 98 on page 719 lists the properties of a `ConnectionFactory` object.

Table 98. Properties of a <code>ConnectionFactory</code> object			
Name of property	Type	Valid values (default value in bold)	Description
<code>applicationName</code>	String	<ul style="list-style-type: none"> The invoking class name, if it is available, adjusted to be no longer than 28 characters. If it is not available, the string <code>WebSphere MQ Client for Java</code> is used. 	The name by which an application is registered with the queue manager. This application name is shown by the DISPLAY CONN MQSC/PCF command (where the field is called APPLTAG) or in the IBM WebSphere MQ Explorer Application Connections display (where the field is called App name).
<code>brokerCCSubQueue¹</code>	String	<ul style="list-style-type: none"> SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE A queue name 	The name of the queue from which a connection consumer receives nondurable subscription messages.
<code>brokerControlQueue¹</code>	String	<ul style="list-style-type: none"> SYSTEM.BROKER.CONTROL.QUEUE A queue name 	The name of the broker control queue.
<code>brokerPubQueue¹</code>	String	<ul style="list-style-type: none"> SYSTEM.BROKER.DEFAULT.STREAM A queue name 	The name of the queue where published messages are sent (the stream queue).
<code>brokerQueueManager¹</code>	String	<ul style="list-style-type: none"> "" (empty string) A queue manager name 	The name of the queue manager on which the broker is running.
<code>brokerSubQueue¹</code>	String	<ul style="list-style-type: none"> SYSTEM.JMS.ND.SUBSCRIBER.QUEUE A queue name 	<p>The name of the queue from which a nondurable message consumer receives messages.</p> <p>See <code>BROKERSUBQ</code> property for more information.</p>

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerVersion ¹	String	<ul style="list-style-type: none"> • unspecified - After the broker has been migrated from V6 to V7, set this property so that RFH2 headers are no longer used. After migration this property is no longer relevant. • V1 - To use a IBM WebSphere MQ Publish/Subscribe broker. Or to use a broker of IBM WebSphere MQ Integrator, WebSphere Event Broker, WebSphere Business Integration Event Broker, or WebSphere Business Integration Message Broker in compatibility mode. This value is the default value if TRANSPORT is set to BIND or CLIENT. • V2 - To use a broker of IBM WebSphere MQ Integrator, WebSphere Event Broker, WebSphere Business Integration Event Broker, or WebSphere Business Integration Message Broker in native mode. This value is the default value if TRANSPORT is set to DIRECT or DIRECTHTTP. 	The version of the broker being used.
ccdtURL	String	<ul style="list-style-type: none"> • null • A uniform resource locator (URL) 	A URL that identifies the name and location of the file containing the client channel definition table (CCDT) and specifies how the file can be accessed.
CCSID	String	<ul style="list-style-type: none"> • 819 • A coded character set identifier supported by the Java virtual machine (JVM) 	The coded character set identifier for a connection.
channel	String	<ul style="list-style-type: none"> • SYSTEM.DEF.SVRCONN • The name of an MQI channel 	The name of the MQI channel to use.
cleanupInterval ¹	int	<ul style="list-style-type: none"> • 3 600 000 • A positive integer 	The interval, in milliseconds, between background runs of the publish/subscribe cleanup utility.
cleanupLevel ¹	String	<ul style="list-style-type: none"> • SAFE • NONE • STRONG • FORCE • NONDUR 	The cleanup level for a broker-based subscription store.
clientID	String	<ul style="list-style-type: none"> • null • A client identifier 	The client identifier for a connection.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
cloneSupport	String	<ul style="list-style-type: none"> • DISABLED - Only one instance of a durable topic subscriber can run at a time. • ENABLED - Two or more instances of the same durable topic subscriber can run simultaneously, but each instance must run in a separate Java virtual machine (JVM). 	Whether two or more instances of the same durable topic subscriber can run simultaneously.
connectionNameList	String	<ul style="list-style-type: none"> • localhost(1414) • A string composed of items separated by commas where each item takes the format: <div style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;"><code>HOSTNAME (PORT)</code></div> where <i>HOSTNAME</i> is either a DNS name or an IP address. 	<p>A list of TCP/IP connection names used for outbound communications.</p> <p>connectionNameList supersedes the hostname and port properties.</p> <p>This property is used to reconnect to multi-instance queue managers.</p> <p>connectionNameList is similar in form to localAddress, but must not be confused with it. localAddress specifies the characteristics of the local communications, whereas connectionNameList specifies how to reach a remote queue manager.</p>
failIfQuiesce	Boolean	<ul style="list-style-type: none"> • true • false 	Whether calls to certain methods fail if the queue manager is in a quiescing state.
headerCompression	String	<ul style="list-style-type: none"> • NONE • SYSTEM - RLE message header compression is performed. 	A list of the techniques that can be used for compressing header data on a connection.
hostName	String	<ul style="list-style-type: none"> • localhost • A host name • An IP address 	<p>The host name or IP address of the system on which the queue manager resides.</p> <p>The hostname and port properties are superseded by the connectionNameList property when it is specified.</p>

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
localAddress	String	<ul style="list-style-type: none"> • null • A string in the format: <pre>[host_name] [(low_port[,high_port])]</pre> <p>where <i>host_name</i> is a host name or IP address, <i>low_port</i> and <i>high_port</i> are TCP port numbers, and brackets denote an optional component</p> 	<p>For a connection to a queue manager, this property specifies either or both of the following:</p> <ul style="list-style-type: none"> • The local network interface to be used • The local port, or range of local ports, to be used <p>localAddress is similar in form to connectionNameList, but must not be confused with it. localAddress specifies the characteristics of the local communications, whereas connectionNameList specifies how to reach a remote queue manager.</p>
messageCompression	String	<ul style="list-style-type: none"> • NONE • A list of one or more of the following values separated by blank characters: <pre>RLE ZLIBFAST ZLIBHIGH</pre> 	A list of the techniques that can be used for compressing message data on a connection.
messageSelection ¹	String	<ul style="list-style-type: none"> • CLIENT • BROKER 	Determines whether message selection is done by IBM WebSphere MQ classes for JMS or by the broker. Message selection by the broker is not supported when brokerVersion has the value 1.
password	String	<ul style="list-style-type: none"> • null • A password 	The default password to use when creating a connection to the queue manager.
pollingInterval ¹	int	<ul style="list-style-type: none"> • 5000 • Any positive integer 	If each message listener within a session has no suitable message on its queue, this value is the maximum interval, in milliseconds, that elapses before each message listener tries again to get a message from its queue. If it frequently happens that no suitable message is available for any of the message listeners in a session, consider increasing the value of this property. This property is relevant only if TRANSPORT has the value BIND or CLIENT.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
port	int	<ul style="list-style-type: none"> • 1414 • A TCP port number 	<p>The port on which the queue manager listens.</p> <p>The hostname and port properties are superseded by the connectionNameList property when it is specified.</p>
providerVersion	string	<ul style="list-style-type: none"> • unspecified • A string in one of the following formats <ul style="list-style-type: none"> – V.R.M.F – V.R.M – V.R – V <p>where V, R, M, and F are integer values greater than or equal to zero.</p>	The version, release, modification level and fix pack of the queue manager to which the application intends to connect.
pubAckInterval ¹	int	<ul style="list-style-type: none"> • 25 • A positive integer 	The number of messages published by a publisher before IBM WebSphere MQ classes for JMS requests an acknowledgment from the broker.
queueManager	String	<ul style="list-style-type: none"> • "" (empty string) • A queue manager name 	The name of the queue manager to connect to.
receiveExit ³	String	<ul style="list-style-type: none"> • null • A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the IBM WebSphere MQ classes for Java interface, MQReceiveExit 	Identifies a channel receive exit program, or a sequence of receive exit programs to be run in succession.
receiveExitInit	String	<ul style="list-style-type: none"> • null • A string comprising one or more items of user data separated by commas 	The user data that is passed to channel receive exit programs when they are called.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
rescanInterval ¹	int	<ul style="list-style-type: none"> • 5000 • Any positive integer 	When a message consumer in the point-to-point domain uses a message selector to select which messages it wants to receive, WebSphere MQ classes for JMS searches the IBM WebSphere MQ queue for suitable messages in the sequence determined by the <i>MsgDeliverySequence</i> attribute of the queue. When WebSphere MQ classes for JMS finds a suitable message and delivers it to the consumer, WebSphere MQ classes for JMS resumes the search for the next suitable message from its current position in the queue. WebSphere MQ classes for JMS continues to search the queue in this way until it reaches the end of the queue, or until the interval of time in milliseconds, as determined by the value of this property, has expired. In each case, WebSphere MQ classes for JMS returns to the beginning of the queue to continue its search, and a new time interval commences.
securityExit ³	String	<ul style="list-style-type: none"> • null • The fully qualified name of a class that implements the WebSphere MQ classes for Java interface, MQSecurityExit 	Identifies a channel security exit program.
securityExitInit	String	<ul style="list-style-type: none"> • null • A string of user data 	The user data that is passed to a channel security exit program when it is called.
sendCheckCount	int	<ul style="list-style-type: none"> • 0 • Any positive integer 	The number of send calls to allow between checking for asynchronous put errors, within a single non-transacted JMS session.
sendExit ³	String	<ul style="list-style-type: none"> • null • A string comprising one or more items separated by commas, where each item is the fully qualified name of a class that implements the WebSphere MQ classes for Java interface, MQSendExit 	Identifies a channel send exit program, or a sequence of send exit programs to be run in succession.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
sendExitInit	String	<ul style="list-style-type: none"> • null • A string comprising one or more items of user data separated by commas 	The user data that is passed to channel send exit programs when they are called.
shareConvAllowed	Boolean	<ul style="list-style-type: none"> • NO - A client connection cannot share its socket. • YES - A client connection can share its socket. 	Whether a client connection can share its socket with other top-level JMS connections from the same process to the same queue manager, if the channel definitions match.
sparseSubscriptions ¹	Boolean	<ul style="list-style-type: none"> • false - Subscriptions receive frequent matching messages. • true - Subscriptions receive infrequent matching messages. This value requires that the subscription queue can be opened for browse. 	Controls the message retrieval policy of a TopicSubscriber object.
sslCertStores	String	<ul style="list-style-type: none"> • null • A string of one or more LDAP URLs separated by blanks. Each LDAP URL has the format: <pre>ldap://host_name[:port]</pre> where <i>host_name</i> is a host name or IP address, <i>port</i> is a TCP port number, and brackets denote an optional component. 	The Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs) for use on an SSL connection.
sslCipherSuite	String	<ul style="list-style-type: none"> • null • The name of a CipherSuite 	The CipherSuite to use for an SSL connection.
sslFipsRequired ²	Boolean	<ul style="list-style-type: none"> • false • true 	Whether an SSL connection must use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (IBMJSEFIPS).
sslPeerName	String	<ul style="list-style-type: none"> • null • A template for distinguished names 	For an SSL connection, a template that is used to check the distinguished name in the digital certificate provided by the queue manager.
sslResetCount	int	<ul style="list-style-type: none"> • 0 • An integer in the range 0 - 999 999 999 	The total number bytes sent and received by an SSL connection before the secret keys used by SSL are renegotiated.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
sslSocketFactory	String	A string representing the fully qualified class name of a class providing an implementation of the javax.net.ssl.SSLSocketFactory interface, optionally including an argument to be passed to the constructor method, enclosed in parentheses.	Any connections established in the scope of the administered destination object use sockets obtained from this implementation of the SSLSocketFactory interface.
statusRefreshInterval ¹	int	<ul style="list-style-type: none"> • 60000 • Any positive integer 	The interval, in milliseconds, between refreshes of the long running transaction that detects when a subscriber loses its connection to the queue manager. This property is relevant only if SUBSTORE has the value QUEUE.
subscriptionStore ¹	String	<ul style="list-style-type: none"> • BROKER • MIGRATE • QUEUE 	Determines where WebSphere MQ classes for JMS stores persistent data about active subscriptions.
targetClientMatching	Boolean	<ul style="list-style-type: none"> • true • false 	Whether a reply message, sent to the queue identified by the JMSReplyTo header field of an incoming message, has an MQRFH2 header only if the incoming message has an MQRFH2 header.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
temporaryModel	String	<ul style="list-style-type: none"> • SYSTEM.DEFAULT.MODEL.QUEUE • SYSTEM.JMS.TEMPQ.MODEL • Any string 	<p>The name of the model queue from which JMS temporary queues are created.</p> <p>Use SYSTEM.DEFAULT.MODEL.QUEUE if both of the following are true:</p> <ul style="list-style-type: none"> • Your application uses a temporary queue that will accept non-persistent messages. • Only one application will create a temporary queue on the queue manager that the ConnectionFactory points to at a time. Note that SYSTEM.DEFAULT.MODEL.QUEUE can only be opened by one application at a time. <p>Use SYSTEM.JMS.TEMPQ.MODEL in the following situations:</p> <ul style="list-style-type: none"> • When your application uses a temporary queue that will accept persistent messages. • If multiple applications can connect to the queue manager that the ConnectionFactory points to and those applications need to create temporary queues at the same time. <p>Define a new model queue with the DEFPSIST attribute set to YES, and the DEFSOPT attribute set to SHARED in the following situation:</p> <ul style="list-style-type: none"> • When your application uses a temporary queue that will accept non-persistent messages, and multiple applications will connect to the queue manager that the ConnectionFactory points to, and those applications need to create temporary queues at the same time. <p>When the new model queue is created, set the temporaryModel property to the name of the new model queue.</p>

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
tempQPrefix	String	<ul style="list-style-type: none"> • "" (empty string) • A prefix that can be used to form the name of a IBM WebSphere MQ dynamic queue. The rules for forming the prefix are the same as the rules for forming the contents of the <i>DynamicQName</i> field in an IBM WebSphere MQ object descriptor, structure MQOD, but the last non-blank character must be an asterisk (*). If the value of the property is the empty string, WebSphere MQ classes for JMS uses the value AMQ.* when creating a dynamic queue. 	The prefix that is used to form the name of an IBM WebSphere MQ dynamic queue.
tempTopicPrefix	String	Any non-null string consisting only of valid characters for an IBM WebSphere MQ topic string	When creating temporary topics, JMS generates a topic string of the form "TEMP/TEMPTOPICPREFIX/ <i>unique_id</i> ", or if this property is left with the default value, just "TEMP/ <i>unique_id</i> ". Specifying a non-empty TEMPTOPICPREFIX allows specific model queues to be defined for creating the managed queues for subscribers to temporary topics created under this connection.
transportType	String	<ul style="list-style-type: none"> • CLIENT • BINDINGS • BINDINGS_THEN_CLIENT 	Whether a connection to a queue manager uses client mode or bindings mode. If the value BINDINGS_THEN_CLIENT is specified, the resource adapter first tries to make a connection in bindings mode. If this connection attempt fails, the resource adapter then tries to make a client mode connection.
username	String	<ul style="list-style-type: none"> • null • A user name 	The default user name to use when creating a connection to a queue manager.
wildcardFormat	int	<ul style="list-style-type: none"> • CHAR- Recognizes character wildcards only, as used in broker version 1 • TOPIC - Recognizes topic level wildcards only, as used in broker version 2 	Which version of wildcard syntax is to be used.

Table 98. Properties of a ConnectionFactory object (continued)

Name of property	Type	Valid values (default value in bold)	Description
Notes: <ol style="list-style-type: none"> 1. This property can be used with Version 7.0 of IBM WebSphere MQ classes for JMS but does not affect an application connected to a Version 7.0 queue manager unless the providerVersion property is set to a version number less than 7. 2. For important information about using the sslFipsRequired property, see “Limitations of the IBM WebSphere MQ resource adapter” on page 736. 3. For information on how to configure the resource adapter so that it can locate an exit, see “Configuring IBM WebSphere MQ classes for JMS to use channel exits” on page 873. 			

The following example shows a typical set of properties of a ConnectionFactory object:

```
channel:      SYSTEM.DEF.SVRCONN
hostName:    192.168.0.42
port:        1414
queueManager: ExampleQM
transportType: CLIENT
```

Table 99 on page 729 lists the properties that are common to a Queue object and a Topic object.

Table 99. Properties that are common to a Queue object and a Topic object

Name of property	Type	Valid values (default value in bold)	Description
CCSID	String	<ul style="list-style-type: none"> • 1208 • A coded character set identifier supported by the Java virtual machine (JVM) 	The coded character set identifier for the destination.
encoding	String	<ul style="list-style-type: none"> • NATIVE • A string of three characters: <ul style="list-style-type: none"> – The first character specifies the representation of binary integers: <ul style="list-style-type: none"> - <i>N</i> denotes normal encoding. - <i>R</i> denotes reverse encoding. – The second character specifies the representation of packed decimal integers: <ul style="list-style-type: none"> - <i>N</i> denotes normal encoding. - <i>R</i> denotes reverse encoding. – The third character specifies the representation of floating point numbers: <ul style="list-style-type: none"> - <i>N</i> denotes standard IEEE encoding. - <i>R</i> denotes reverse IEEE encoding. - <i>3</i> denotes zSeries encoding. <p>NATIVE is equivalent to the string NNN.</p>	The representation of binary integers, packed decimal integers, and floating point numbers for the destination.

Table 99. Properties that are common to a Queue object and a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
expiry	String	<ul style="list-style-type: none"> • APP - The expiry time of a message is determined by the message producer. • UNLIM - A message never expires. • 0 - A message never expires. • A positive integer representing the expiry time of a message in milliseconds. 	The expiry time of a message sent to the destination.
failIfQuiesce	String	<ul style="list-style-type: none"> • true • false 	Whether an attempt to access the destination fails if the queue manager is in a quiescing state.
persistence	String	<ul style="list-style-type: none"> • APP - The persistence of a message is determined by the message producer. • QDEF - The persistence of a message is determined by the <i>DefPersistence</i> attribute of the WebSphere MQ queue. • PERS - A message is persistent. • NON - A message is nonpersistent. • HIGH - The persistence of a message is determined by the <i>NonPersistentMessageClass</i> attribute of the WebSphere MQ queue according to the explanation in “JMS persistent messages” on page 866. 	The persistence of a message sent to the destination.
priority	String	<ul style="list-style-type: none"> • APP - The priority of a message is determined by the message producer. • QDEF - The priority of a message is determined by the <i>DefPriority</i> attribute of the IBM WebSphere MQ queue. • An integer in the range 0, lowest priority, to 9, highest priority. 	The priority of a message sent to the destination.
putAsyncAllowed	String	<ul style="list-style-type: none"> • QUEUE - Determine whether asynchronous puts are allowed by referring to the queue definition. • TOPIC - Determine whether asynchronous puts are allowed by referring to the topic definition. • DESTINATION - Determine whether asynchronous puts are allowed by referring to the queue or topic definition. • DISABLED - Asynchronous puts are not allowed. • ENABLED - Asynchronous puts are allowed. 	Whether message producers are allowed to use asynchronous puts to send messages to this destination.

Table 99. Properties that are common to a Queue object and a Topic object (continued)			
Name of property	Type	Valid values (default value in bold)	Description
readAheadAllowed	int	<ul style="list-style-type: none"> • DESTINATION - Determine whether read ahead is allowed by referring to the queue or topic definition. • DISABLED - Read ahead is not allowed. • ENABLED - Read ahead is allowed. • QUEUE - Determine whether read ahead is allowed by referring to the queue definition. • TOPIC - Determine whether read ahead is allowed by referring to the topic definition. 	Whether message consumers and queue browsers are allowed to use read ahead to get nonpersistent messages from the destination into an internal buffer before receiving them.
receiveCCSID	int	<ul style="list-style-type: none"> • 0 - Use JVM Charset.defaultCharset • 1208 - UTF-8 • A supported coded character set identifier 	Destination property that sets the target CCSID for queue manager message conversion. The value is ignored unless receiveConversion is set to QMGR
receiveConversion	String	<ul style="list-style-type: none"> • CLIENT_MSG • QMGR 	Destination property that determines whether data conversion is going to be performed by the queue manager.
targetClient	String	<ul style="list-style-type: none"> • JMS - The target of a message is a JMS application. • MQ - The target of a message is a non-JMS IBM WebSphere MQ application. 	Whether the target of a message sent to the destination is a JMS application. A message with a target that is a JMS application contains an MQRFH2 header.

Table 100 on page 731 lists the properties that are specific to a Queue object.

Table 100. Properties that are specific to a Queue object			
Name of property	Type	Valid values (default value in bold)	Description
baseQueueManagerName	String	<ul style="list-style-type: none"> • "" (empty string) • A queue manager name 	The name of the queue manager that owns the underlying IBM WebSphere MQ queue.
baseQueueName	String	<ul style="list-style-type: none"> • "" (empty string) • A queue name 	The name of the underlying IBM WebSphere MQ queue.

Table 101 on page 731 lists the properties that are specific to a Topic object.

Table 101. Properties that are specific to a Topic object			
Name of property	Type	Valid values (default value in bold)	Description
baseTopicName	String	<ul style="list-style-type: none"> • "" (empty string) • A topic name 	The name of the underlying topic.

Table 101. Properties that are specific to a Topic object (continued)

Name of property	Type	Valid values (default value in bold)	Description
brokerCCDurSubQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE • A queue name 	The name of the queue from which a connection consumer receives durable subscription messages.
brokerDurSubQueue ¹	String	<ul style="list-style-type: none"> • SYSTEM.JMS.D.SUBSCRIBER.QUEUE • A queue name 	The name of the queue from which a durable topic subscriber receives messages. See the BROKEDURRSUBQ property in the WebSphere MQ Explorer documentation for more information.
brokerPubQueue ¹	String	<ul style="list-style-type: none"> • Not set • A queue name 	The name of the queue where published messages are sent (the stream queue). The value of this property overrides the value of the brokerPubQueue property of the ConnectionFactory object. However, if you do not set the value of this property, the value of the brokerPubQueue property of the ConnectionFactory object is used instead.
brokerPubQueueManager ¹	String	<ul style="list-style-type: none"> • "" (empty string) • A queue manager name 	The name of the queue manager that owns the queue where messages published on the topic are sent.
brokerVersion ¹	String	<ul style="list-style-type: none"> • Not set • 1 • 2 	The version of the broker being used. The value of this property overrides the value of the brokerVersion property of the ConnectionFactory object. However, if you do not set the value of this property, the value of the brokerVersion property of the ConnectionFactory object is used instead.

Note:

1. This property can be used with Version 7.0 of IBM WebSphere MQ classes for JMS but does not affect an application connected to a Version 7.0 queue manager unless the providerVersion property of the ConnectionFactory object is set to a version number less than 7.

The following example shows a set of properties of a Queue object:

```
expiry:          UNLIM
persistence:    QDEF
baseQueueManagerName: ExampleQM
baseQueueName:  SYSTEM.JMS.TEMPQ.MODEL
```

The following example shows a set of properties of a Topic object:

```
expiry:          UNLIM
persistence:     NON
baseTopicName:   myTestTopic
```

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](#)

V7.5.0.9 *Configuring the targetClientMatching property for an activation specification*

You can configure the **targetClientMatching** property for an activation specification so that an MQRFH2 header is included on reply messages when request messages do not contain an MQRFH2 header. This means that any message properties that an application defines on a reply message are included when the message is sent.

About this task

If a message-driven bean (MDB) application consumes messages that do not contain an MQRFH2 header, through an IBM WebSphere MQ JCA resource adapter activation specification, and subsequently sends reply messages to the JMS Destination created from the JMSReplyTo field of the request message, the reply messages must include an MQRFH2 header, even if the request messages do not, otherwise any message properties that the application has defined on a reply message are lost.

The **targetClientMatching** property defines whether a reply message, sent to the queue identified by the JMSReplyTo header field of an incoming message, has an MQRFH2 header only if the incoming message has an MQRFH2 header. You can configure this property for an activation specification, in both WebSphere Application Server traditional and WebSphere Application Server Liberty.

If you set the value of the **targetClientMatching** property to `false`, an MQRFH2 header can be included in a reply message sent to a JMS Destination created from the JMSReplyTo header of an incoming request message that does not contain an MQRFH2. This is because the **targetClient** property on the JMS Destination is set to the value 0, which means that messages contain an MQRFH2 header. The presence of the MQRFH2 header in the outbound message permits the storage of user defined message properties on the message when sent to the IBM WebSphere MQ queue.

If the **targetClientMatching** property is set to `true` and a request message does not include an MQRFH2 header, an MQRFH2 header is not included in the reply message.

Procedure

- In WebSphere Application Server traditional, use the administration console to define the **targetClientMatching** property as a custom property on the IBM WebSphere MQ activation specification:
 - a) In the navigation pane, click **Resources -> JMS -> Activation specifications**.
 - b) Select the name of the activation specification that you want to view or change.
 - c) Click **Custom properties -> New** and then enter the details of the new custom property.
Set the name of the property to `targetClientMatching`, the type to `java.lang.Boolean` and the value to `false`.
- In WebSphere Application Server Liberty, specify the **targetClientMatching** property on the definition of an activation specification within the `server.xml`.

For example:

```
<jmsActivationSpec id="SimpleMDBApplication/SimpleEchoMDB/SimpleEchoMDB">
  <properties.wmqJms destinationRef="MDBRequestQ"
    queueManager="MY_QMGR" transportType="BINDINGS" targetClientMatching="false"/>
  <authData password="*****" user="tom"/>
</jmsActivationSpec>
```

Related concepts

[“Creating destinations in a JMS application” on page 840](#)

Instead of retrieving destinations as administered objects from a Java Naming and Directory Interface (JNDI) namespace, a JMS application can use a session to create destinations dynamically at run time. An application can use a uniform resource identifier (URI) to identify a WebSphere MQ queue or a topic and, optionally, to specify one or more properties of a Queue or Topic object.

[“Configuring the resource adapter for outbound communication” on page 718](#)

To configure outbound communication, define the properties of a ConnectionFactory object and an administered destination object.

ASF and non-ASF mode

Application Server Facilities (ASF) mode is the default method by which the message listener service in WebSphere Application Server processes messages.

The message listener service has two modes of operation, Application Server Facilities (ASF) and non-Application Server Facilities (non-ASF):

- ASF mode provides concurrency and transactional support for applications. For publish/subscribe message-driven beans, ASF mode provides better throughput and concurrency, because in the non-ASF mode the listener is single-threaded.
- Non-ASF mode is mainly for use with third-party messaging providers that do not support JMS ASF, which is an optional extension to the JMS specification. Non-ASF mode is also transactional but, because the path length is shorter than for ASF mode, usually provides improved performance.

To enable the non-ASF mode of operation for all message-driven bean listeners on the application server, set this property to a non-zero value.

Note:

Non-ASF mode cannot be selected on z/OS systems, so you must not set a non-zero value for this property in this case.

Message processing in ASF mode

In ASF mode, server sessions and threads are only allocated for work when a message that is suitable for the message-driven bean (MDB) is detected. The number of threads that an MDB can process concurrently is determined by the value of the **Maximum Sessions** property for the listener port or activation specification.

Message processing in non-ASF mode

In non-ASF mode threads are active from the moment that the listener port or activation specification is started. The number of active threads is dictated by the value specified for **Maximum Sessions** property. The number of threads specified in **Maximum Sessions** property are active, regardless of the number of messages that are available to be processed. Each active thread is an individual physical network connection.

IBM WebSphere MQ Version 7.0 or later allows you to have up to ten threads sharing a single physical network connection.

Related concepts

IBM WebSphere MQ classes for JMS Application Server Facilities

This topic describes how WebSphere MQ classes for JMS implements the ConnectionConsumer class and advanced functionality in the Session class. It also summarizes the function of a server session pool.

Related tasks

Configuring Activation Specifications for non-ASF mode

Activation specifications are the standardized way to manage and configure the relationship between a message driven bean (MDB) running in WebSphere Application Server and a destination in IBM

WebSphere MQ. This task explains how to configure WebSphere Application Server to use non-ASF mode to processes messages.

Related information

[Message processing in ASF mode and non-ASF mode](#)

Configuring activation specifications for non-ASF mode

Activation specifications are the standardized way to manage and configure the relationship between a message driven bean (MDB) running in WebSphere Application Server and a destination in IBM WebSphere MQ. This task explains how to configure WebSphere Application Server to use non-ASF mode to processes messages.

Before you begin

The way in which you define the properties of an activation specification depends on the administration interfaces provided by your application server. This task assumes that you are using WebSphere Application Server version 7 or later as your application server and IBM WebSphere MQ as your messaging provider.

Note:

Non-ASF mode cannot be selected on z/OS systems.

About this task

The properties of an activation specification determine how a message drive bean (MDB) receives JMS messages from an IBM WebSphere MQ queue. To configure non-ASF mode, define the properties of one or more activation specifications.

There are several IBM WebSphere MQ configurations that you can use in non-ASF mode. With the following configurations each thread uses a separate physical network connection:

- An IBM WebSphere MQ Version 7.x queue manager, using a connection factory that has the Provider version property set to 6.
- An IBM WebSphere MQ Version 7.x queue manager, using a connection factory that has the Provider version property set to 7 or unspecified, connecting over a IBM WebSphere MQ channel that has the **SHARECNV** (sharing conversations) parameter set to 0.

To configure non-ASF, set the ActivationSpec property **NON.ASF.RECEIVE.TIMEOUT** to a positive integer, that indicates that non-ASF delivery is used. The value is the time, in milliseconds, that a get request waits for messages that might not have yet arrived (a get with wait call). The default value, 0, indicates that ASF delivery is used. For further details see [Message listener service custom properties](#).

This parameter is valid only when the application is running on WebSphere Application Server version 7 or later.

Procedure

1. Start the WebSphere Application Server administrative console.
2. Display the listener service settings page:
 - a) In the navigation pane, select **Servers > Server Types > WebSphere application servers**.
 - b) In the content pane, click the name of the application server.
 - c) Under **Communications**, click **Messaging > Message Listener Service**.
3. Set the custom property **NON.ASF.RECEIVE.TIMEOUT** as a Custom properties of the message listener service.
 - a) Click **Custom properties**.
 - b) Click **New**.
 - c) Enter the name of the property, **NON.ASF.RECEIVE.TIMEOUT**, in the **Name** field.

- d) Enter the value you require, in the **Value** field.
 - e) Click **OK**.
4. Save your changes to the master configuration.
 5. To activate the changed configuration, stop and then restart the application server.

Results

You have configured the properties of the message listener service for WebSphere Application Server to use non-ASF mode.

Note: When using non-ASF mode, you must ensure that you allow a sufficient amount of time for processing to be completed before the total transaction lifetime timeout is reached, in order to avoid unwanted transaction time-outs. For further details see [**NON.ASF.RECEIVE.TIMEOUT**](#) in the WebSphere Application Server product documentation.

Related concepts

[“ASF and non-ASF mode” on page 734](#)

Application Server Facilities (ASF) mode is the default method by which the message listener service in WebSphere Application Server processes messages.

[Configuring the resource adapter for inbound communication](#)

To configure inbound communication, define the properties of one or more ActivationSpec objects.

Related information

[Message-driven beans](#)

[Message listener service](#)

[Message processing in ASF mode and non-ASF mode](#)

[How messages are processed in non-ASF mode](#)

Limitations of the IBM WebSphere MQ resource adapter

When you use the IBM WebSphere MQ resource adapter, some features of IBM WebSphere MQ are unavailable or limited.

The IBM WebSphere MQ resource adapter has the following limitations:

- The IBM WebSphere MQ resource adapter is supported on all IBM WebSphere MQ platforms, except z/OS.
- The IBM WebSphere MQ resource adapter does not support real-time connections to a broker. It supports only connections to an IBM WebSphere MQ queue manager in client or bindings mode.
- The IBM WebSphere MQ resource adapter does not support channel exit programs that are written in languages other than Java.
- While an application server is running, the value of the `sslFipsRequired` property must be true for all JCA resources or false for all JCA resources. This is a requirement even if the JCA resources are not used concurrently. If the `sslFipsRequired` property has different values for different JCA resources, IBM WebSphere MQ issues the reason code `MQRC_UNSUPPORTED_CIPHER_SUITE`, even if an SSL connection is not being used.
- You cannot specify more than one keystore for an application server. If connections are made to more than one queue manager, all the connections must use the same keystore. This limitation does not apply to WebSphere Application Server.
- If you use a client channel definition table (CCDT) with more than one suitable client connection channel definition, in the event of a failure the resource adapter might select a different channel definition and therefore a different queue manager from the CCDT, which would cause problems for transaction recovery. The resource adapter does not take any action to prevent such a configuration from being used, and it is your responsibility to avoid configurations that might cause problems for transaction recovery.
- The connection retry functionality introduced in IBM WebSphere MQ Version 7.0.1 is not supported for outbound connections when running in a JEE container (EJB/Servlet). Connection retry is not supported

at all for outbound JMS when the adapter is used in a JEE container context, regardless of transaction configuration or for non-transacted use.

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](#)

Post installation setup for WebSphere MQ classes for JMS applications

This topic tells you what authorities WebSphere MQ classes for JMS applications need in order to access the resources of a queue manager. It also introduces connection modes and describes how to configure a queue manager so that applications can connect in client mode.

Remember to check the WebSphere MQ readme file. It might contain information that supersedes the information in this topic.

Objects used by JMS that require authorization for non-privileged users

Non-privileged users need authorization granted to access the queues used by JMS. Every JMS application needs authorization to the queue manager with which it works.

For details about access control in IBM WebSphere MQ, see [Setting up security on Windows, UNIX and Linux systems](#).

WebSphere MQ classes for JMS applications need connect and inq authority to the queue manager. You can set appropriate authorizations using the **setmqaut** control command, for example:

```
setmqaut -m QM1 -t qmgr -g jmsappsgroup +connect +inq
```

For the point-to-point domain, the following authorities are required:

- Queues that are used by MessageProducer objects need put authority.
- Queues that are used by MessageConsumer and QueueBrowser objects need get, inq, and browse authorities.
- The QueueSession.createTemporaryQueue() method needs access to the model queue specified by the TEMPMODEL property of the QueueConnectionFactory object. By default this model queue is SYSTEM.TEMP.MODEL.QUEUE.

If any of these queues are alias queues, their target queues require inquire authority. If the target queue is a cluster queue it also requires browse authority.

For the publish/subscribe domain, the following queues are used if the WebSphere MQ classes for JMS are connecting to an IBM WebSphere MQ queue manager in IBM WebSphere MQ messaging provider migration mode:

- SYSTEM.JMS.ADMIN.QUEUE
- SYSTEM.JMS.REPORT.QUEUE
- SYSTEM.JMS.MODEL.QUEUE
- SYSTEM.JMS.PS.STATUS.QUEUE
- SYSTEM.JMS.ND.SUBSCRIBER.QUEUE
- SYSTEM.JMS.D.SUBSCRIBER.QUEUE
- SYSTEM.JMS.ND.CC.SUBSCRIBER.QUEUE
- SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE
- SYSTEM.BROKER.CONTROL.QUEUE

For further information on IBM WebSphere MQ messaging provider migration mode, see [When to use PROVIDERVERSION](#)

Additionally, if the WebSphere MQ classes for JMS are connecting to a queue manager in this mode, any application that publishes messages needs access to the stream queue specified by the TopicConnectionFactory or topic object. By default, this queue is SYSTEM.BROKER.DEFAULT.STREAM.

If you use ConnectionConsumer, IBM WebSphere MQ Resource Adapter, or the WebSphere Application Server IBM WebSphere MQ messaging provider, additional authorization might be needed.

Queues to be read by the ConnectionConsumer must have get, inq, and browse authorities. The system dead-letter queue, and any backout-requeue queue or report queue used by the ConnectionConsumer must have put and passall authorities.

When an application uses WebSphere MQ messaging provider normal mode to perform publish/subscribe messaging, the application makes use of the integrated publish/subscribe functionality provided by the queue manager. See [Publish/subscribe security](#) for information on securing the topics and queues that are used.

Connection modes for WebSphere MQ classes for JMS

A WebSphere MQ classes for JMS application can connect to a queue manager in either client or bindings mode. In client mode, WebSphere MQ classes for JMS connects to the queue manager over TCP/IP. In bindings mode, WebSphere MQ classes for JMS connects directly to the queue manager using the Java Native Interface (JNI).

An application running in WebSphere Application Server on z/OS can connect to a queue manager in either bindings or client mode, but an application running in any other environment on z/OS can connect to a queue manager only in bindings mode. An application running on any other platform can connect to a queue manager in either bindings or client mode.

You can use the current or any earlier supported version of WebSphere MQ classes for JMS with a current queue manager, and you can use a current or earlier supported version of queue manager with the current version of WebSphere MQ classes for JMS. If you mix different versions, function is limited to the level of the earlier version.

The following sections describe each of the connection modes in more detail.

Client mode

To connect to a queue manager in client mode, a WebSphere MQ classes for JMS application can run on the same system on which the queue manager is running, or on a different system. In each case, WebSphere MQ classes for JMS connects to the queue manager over TCP/IP.

Bindings mode

To connect to a queue manager in bindings mode, a WebSphere MQ classes for JMS application must run on the same system on which the queue manager is running.

The WebSphere MQ classes for JMS connects directly to the queue manager using the Java Native Interface (JNI). To use the bindings transport, the WebSphere MQ classes for JMS must be run in an environment that has access to the WebSphere MQ Java Native Interface libraries; see [“Configuring the Java Native Interface \(JNI\) libraries”](#) on page 693 for further information.

The WebSphere MQ classes for JMS support the following values for *ConnectOption* :

- MQCNO_FASTPATH_BINDING
- MQCNO_STANDARD_BINDING
- MQCNO_SHARED_BINDING
- MQCNO_ISOLATED_BINDING
- MQCNO_SERIALIZE_CONN_TAG_QSG
- MQCNO_RESTRICT_CONN_TAG_QSG
- MQCNO_SERIALIZE_CONN_TAG_Q_MGR
- MQCNO_RESTRICT_CONN_TAG_Q_MGR

To change the connection options used by the WebSphere MQ classes for JMS, modify the Connection Factory property `CONNOPT`.

For further information on connection options, see [“Connecting to a queue manager using the MQCONN call” on page 200](#)

To use the bindings transport, the Java Runtime Environment being used must support the Coded Character Set Identifier (CCSID) of the queue manager that the WebSphere MQ classes for JMS are connecting to.

Details on how to determine what CCSID's are supported by a Java Runtime Environment can be found in [WebSphere MQ FDC with Probe id 21 generated when using the WebSphere MQ V7 classes for Java or WebSphere MQ V7 classes for JMS](#).

Bindings, then client mode

This is the default. When connecting to a queue manager in this mode, a WebSphere MQ classes for JMS application will attempt to connect in bindings mode, which requires the queue manager to reside on the same machine as the application. Should the connection be unsuccessful, the application will then attempt to connect in client mode, allowing the queue manager to reside either locally on the same machine as the application, or remotely.

Configuring your queue manager so that WebSphere MQ classes for JMS applications can connect in client mode

To configure your queue manager so that WebSphere MQ classes for JMS applications can connect in client mode, you must create a server connection channel definition and start a listener.

On z/OS, the Client Attachment feature must be installed.

Creating a server connection channel definition

On all platforms, you can use the MQSC command `DEFINE CHANNEL` to create a server connection channel definition. See the following example:

```
DEFINE CHANNEL (JAVA.CHANNEL) CHLTYPE(SVRCONN) TRPTYPE(TCP)
```

On IBM i, you can use the CL command `CRTMQMCHL` instead, as in the following example:

```
CRTMQMCHL CHLNAME(JAVA.CHANNEL) CHLTYPE(*SVRCN)
          TRPTYPE(*TCP)
          QMNAME(QMGRNAME)
```

In this command, `QMGRNAME` is the name of your queue manager.

You can also create a server connection channel definition using IBM WebSphere MQ Explorer, which runs on Linux and Windows, or the operations and control panels on z/OS.

The name of the channel (`JAVA.CHANNEL` in the previous examples) must be the same as the channel name specified by the `CHANNEL` property of the connection factory that your application uses to connect to the queue manager. The default value of the `CHANNEL` property is `SYSTEM.DEF.SVRCONN`.

Starting a listener

You must start a listener for your queue manager if one is not already started.

On all platforms, you can use the MQSC command `START LISTENER` to start a listener but, except on z/OS, you must first create a listener object by using the MQSC command `DEFINE LISTENER`. See the following example:

```
DEFINE LISTENER(LISTENER.TCP) TRPTYPE(TCP) PORT(1414)
START LISTENER(LISTENER.TCP)
```

On UNIX, Linux, and Windows systems, you can also use the control command **runmqclsr** to start a listener, as in the following example:

```
runmqclsr -t tcp -p 1414 -m QMgrName
```

In this command, *QMgrName* is the name of your queue manager.

You can also start a listener using WebSphere MQ Explorer, which runs on Linux and Windows, or the operations and control panels on z/OS.

The number of the port on which the listener is listening must be the same as the port number specified by the PORT property of the connection factory that your application uses to connect to the queue manager. The default value of the PORT property is 1414.

The point-to-point installation verification test for WebSphere MQ classes for JMS

A point-to-point installation verification test (IVT) program is supplied with WebSphere MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, sends a message to the queue called SYSTEM.DEFAULT.LOCAL.QUEUE, and then receives the message from the queue. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

Run the installation verification test without using JNDI first because the test is self contained and does not require the use of a directory service. For a description of administered objects, see [“JMS object types”](#) on page 897.

The point-to-point installation verification test without using JNDI

In this test, the IVT program creates and configures all the objects that it requires dynamically at run time and does not use JNDI.

A script is provided to run the IVT program. The script is called IVTRun on UNIX and Linux systems and IVTRun.bat on Windows, and is in the bin subdirectory of the WebSphere MQ classes for JMS installation directory.

To run the test in bindings mode, enter the following command:

```
IVTRun -nojndi [-m qmgr] [-v providerVersion] [-t]
```

To run the test in client mode, first set up the queue manager as described in [“Preparing and running the sample programs”](#) on page 104. Note that the channel to be used defaults to **SYSTEM.DEF.SVRCONN** and the queue to be used is **SYSTEM.DEFAULT.LOCAL.QUEUE**, then enter the following command:

```
IVTRun -nojndi -client -m qmgr -host hostname [-port port] [-channel channel]  
[-v providerVersion] [-ccsid ccsid] [-t]
```

No equivalent script is provided on z/OS systems, but you can run the IVT in bindings mode by invoking the Java class directly, using the following command:

```
java com.ibm.mq.jms.MQJMSIVT -nojndi [-m qmgr] [-v providerVersion] [-t]
```

The classpath must contain com.ibm.mqjms.jar.

The parameters on the commands have the following meanings:

-m qmgr

The name of the queue manager to which the IVT program connects. If you run the test in bindings mode and omit this parameter, the IVT program connects to the default queue manager.

-host hostname

The host name or IP address of the system on which the queue manager is running.

-port port

The number of the port on which the listener of the queue manager is listening. The default value is 1414.

-channel channel

The name of the MQI channel that the IVT program uses to connect to the queue manager. The default value is `SYSTEM.DEF.SVRCONN`.

-v providerVersion

The release level of the queue manager to which the IVT program expects to connect.

This parameter is used to set the `PROVIDERVERSION` property of an `MQQueueConnectionFactory` object and has the same valid values as those of the `PROVIDERVERSION` property. For more information about this parameter therefore, including its valid values, see the description of the `PROVIDERVERSION` property in [Properties of IBM WebSphere MQ classes for JMS objects](#).

The default value is unspecified.

-ccsid ccsid

The identifier (CCSID) of the coded character set, or code page, to be used by the connection. The default value is 819.

-t

Tracing is switched on. By default, tracing is switched off.

A successful test produces output similar to the following sample output:

```
5724-H72, 5655-R36, 5724-L26, 5655-L82 (c) Copyright IBM Corp. 2008, 2024. All
Rights Reserved.
Websphere MQ classes for Java(tm) Message Service 7.0
Installation Verification Test
```

```
Creating a QueueConnectionFactory
Creating a Connection
Creating a Session
Creating a Queue
Creating a QueueSender
Creating a QueueReceiver
Creating a TextMessage
Sending the message to SYSTEM.DEFAULT.LOCAL.QUEUE
Reading the message back again
```

```
Got message
JMSMessage class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:414d5120514d5f6d627720202020202001edb14620005e03
JMSTimestamp: 1187170264000
JMSCorrelationID: null
JMSDestination: queue:///SYSTEM.DEFAULT.LOCAL.QUEUE
JMSReplyTo: null
JMSRedelivered: false
JMSXUserID: mwhite
JMS_IBM_Encoding: 273
JMS_IBM_PutApplType: 28
JMSXAppID: WebSphere MQ Client for Java
JMSXDeliveryCount: 1
JMS_IBM_PutDate: 20070815
JMS_IBM_PutTime: 09310400
JMS_IBM_Format: MQSTR
JMS_IBM_MsgType: 8
```

```
A simple text message from the MQJMSIVT
Reply string equals original string
Closing QueueReceiver
Closing QueueSender
Closing Session
Closing Connection
```

IVT completed OK
IVT finished

The point-to-point installation verification test using JNDI

In this test, the IVT program uses JNDI to retrieve administered objects from a directory service.

Before you can run the test, you must configure a directory service that is based on a Lightweight Directory Access Protocol (LDAP) server or the local file system. You must also configure the WebSphere MQ JMS administration tool so that it can use the directory service to store administered objects. For more information about these prerequisites, see [“Prerequisites for WebSphere MQ classes for JMS”](#) on page 686. For information about how to configure the WebSphere MQ JMS administration tool, see [“Configuring the JMS Administration tool”](#) on page 894.

The IVT program must be able to use JNDI to retrieve an `MQQueueConnectionFactory` object and an `MQQueue` object from the directory service. A script is provided to create these administered objects for you. The script is called `IVTSetup` on UNIX and Linux systems and `IVTSetup.bat` on Windows, and is in the `bin` subdirectory of the WebSphere MQ classes for JMS installation directory. To run the script, enter the following command:

```
IVTSetup
```

The script invokes the WebSphere MQ JMS administration tool to create the administered objects.

The `MQQueueConnectionFactory` object is bound with the name `ivtQCF` and is created with the default values for all its properties, which means that the IVT program runs in bindings mode and connects to the default queue manager. If you want the IVT program to run in client mode, or connect to a queue manager other than the default queue manager, you must use the WebSphere MQ JMS administration tool or WebSphere MQ Explorer to change the appropriate properties of the `MQQueueConnectionFactory` object. For information about how to use the WebSphere MQ JMS administration tool, see [“Using the WebSphere MQ JMS administration tool”](#) on page 893. For information about how to use WebSphere MQ Explorer, see the help provided with WebSphere MQ Explorer.

The `MQQueue` object is bound with the name `ivtQ` and is created with the default values for all its properties, except for the `QUEUE` property, which has the value `SYSTEM.DEFAULT.LOCAL.QUEUE`.

When you have created the administered objects, you can run the IVT program. To run the test using JNDI, enter the following command:

```
IVTRun -url "providerURL" [-icf initCtxFact] [-t]
```

The parameters on the command have the following meanings:

-url "providerURL"

The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName`, for a directory service based on an LDAP server
- `file:/directoryPath`, for a directory service based on the local file system

Note that you must enclose the URL in quotation marks (").

-icf *initCtxFact*

The class name of the initial context factory, which must be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server. This is the default value.
- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system.

-t

Tracing is switched on. By default, tracing is switched off.

A successful test produces output similar to that for a successful test without using JNDI. The main difference is that the output indicates that the test is using JNDI to retrieve an MQQueueConnectionFactory object and an MQQueue object.

Although not strictly necessary, it is good practice to tidy up after the test by deleting the administered objects created by the IVTSetup script. A script is provided for this purpose. The script is called IVTTidy on UNIX and Linux systems and IVTTidy.bat on Windows, and is in the bin subdirectory of the WebSphere MQ classes for JMS installation directory.

Problem determination for the point-to-point installation verification test

The installation verification test might fail for the following reasons:

- If the IVT program writes a message indicating that it cannot find a class, check that your class path is set correctly, as described in [“Environment variables used by IBM WebSphere MQ classes for JMS”](#) on page 691.
- The test might fail with the following message:

```
Failed to connect to queue manager 'qmgr' with connection mode 'connMode'
and host name 'hostname'
```

and an associated reason code of 2059. The variables in the message have the following meanings:

qmgr

The name of the queue manager to which the IVT program is trying to connect. This message insert is blank if the IVT program is trying to connect to the default queue manager in bindings mode.

connMode

The connection mode, which is either Bindings or Client.

hostname

The host name or IP address of the system on which the queue manager is running.

This message means that the queue manager to which the IVT program is trying to connect is not available. Check that the queue manager is running and, if the IVT program is trying to connect to the default queue manager, make sure that the queue manager is defined as the default queue manager for your system.

- The test might fail with the following message:

```
Failed to open MQ queue 'SYSTEM.DEFAULT.LOCAL.QUEUE'
```

This message means that the queue SYSTEM.DEFAULT.LOCAL.QUEUE does not exist on the queue manager to which the IVT program is connected. Alternatively, if the queue does exist, the IVT program cannot open the queue because it is not enabled for putting and getting messages. Check that the queue exists and that it is enabled for putting and getting messages.

- The test might fail with the following message:

```
Unable to bind to object
```

This message means that there is a connection to the LDAP server, but that the LDAP server is not correctly configured. Either the LDAP server is not configured for storing Java objects, or the permissions on the objects or the suffix are not correct. For more help in this situation, see the documentation for your LDAP server.

- The test might fail with the following message:

```
The security authentication was not valid that was supplied for
QueueManager 'qmgr' with connection mode 'Client' and host name 'hostname'
```

This message means that the queue manager is not correctly set up to accept a client connection from your system. See [“Preparing and running the sample programs”](#) on page 104 for details.

The publish/subscribe installation verification test for WebSphere MQ classes for JMS

A publish/subscribe installation verification test (IVT) program is supplied with WebSphere MQ classes for JMS. The program connects to a queue manager in either bindings or client mode, subscribes to a topic, publishes a message on the topic, and then receives the message that it has just published. The program can create and configure all the objects that it requires dynamically at run time, or it can use JNDI to retrieve administered objects from a directory service.

Run the installation verification test without using JNDI first because the test is self contained and does not require the use of a directory service. For a description of administered objects, see [“JMS object types”](#) on page 897.

The publish/subscribe installation verification test without using JNDI

In this test, the IVT program creates and configures all the objects that it requires dynamically at run time and does not use JNDI.

A script is provided to run the IVT program. The script is called PSIVTRun on UNIX and Linux systems and PSIVTRun.bat on Windows, and is in the bin subdirectory of the WebSphere MQ classes for JMS installation directory.

To run the test in bindings mode, enter the following command:

```
PSIVTRun -nojndi [-m qmgr] [-bqm brokerQmgr] [-v providerVersion] [-t]
```

To run the test in client mode, first set up the queue manager as described in [“Preparing and running the sample programs”](#) on page 104 noting that the channel to be used defaults to SYSTEM.DEF.SVRCONN, then enter the following command:

```
PSIVTRun -nojndi -client -m qmgr -host hostname [-port port] [-channel channel]
          [-bqm brokerQmgr] [-v providerVersion] [-ccsid ccsid] [-t]
```

The parameters on the commands have the following meanings:

-m qmgr

The name of the queue manager to which the IVT program connects. If you run the test in bindings mode and omit this parameter, the IVT program connects to the default queue manager.

-host hostname

The host name or IP address of the system on which the queue manager is running.

-port port

The number of the port on which the listener of the queue manager is listening. The default value is 1414.

-channel channel

The name of the MQI channel that the IVT program uses to connect to the queue manager. The default value is SYSTEM.DEF.SVRCONN.

-bqm brokerQmgr

The name of the queue manager on which the broker is running. The default value is the name of the queue manager to which the IVT program connects.

This parameter is relevant only if the -v parameter specifies a queue manager version number less than 7 and you are using WebSphere Event Broker or WebSphere Message Broker as the publish/subscribe broker.

-v providerVersion

The release level of the queue manager to which the IVT program expects to connect.

This parameter is used to set the PROVIDERVERSION property of an MQTopicConnectionFactory object and has the same valid values as those of the PROVIDERVERSION property. For more information about this parameter therefore, including its valid values, see the description of the PROVIDERVERSION property in [Properties of IBM WebSphere MQ classes for JMS objects](#).

The default value is unspecified.

-ccsid ccsid

The identifier (CCSID) of the coded character set, or code page, to be used by the connection. The default value is 819.

-t

Tracing is switched on. By default, tracing is switched off.

A successful test produces output similar to the following sample output:

```
5724-H72, 5655-R36, 5724-L26, 5655-L82 (c) Copyright IBM Corp. 2008, 2024. All
Rights Reserved.
Websphere MQ classes for Java(tm) Message Service 7.0
Publish/Subscribe Installation Verification Test
```

```
Creating a TopicConnectionFactory
Creating a Connection
Creating a Session
Creating a Topic
Creating a TopicPublisher
Creating a TopicSubscriber
Creating a TextMessage
Adding text
Publishing the message to topic://MQJMS/PSIVT/Information
Waiting for a message to arrive [5 secs max]...
```

```
Got message:
JMSMessage class: jms_text
JMSType: null
JMSDeliveryMode: 2
JMSExpiration: 0
JMSPriority: 4
JMSMessageID: ID:414d5120514d5f6d627720202020202001edb14620006706
JMSTimestamp: 1187182520203
JMSCorrelationID: ID:414d5120514d5f6d627720202020202001edb14620006704
JMSDestination: topic://MQJMS/PSIVT/Information
JMSReplyTo: null
JMSRedelivered: false
JMSXUserID: mwhite
JMS_IBM_Encoding: 273
JMS_IBM_PutApplType: 26
JMSXAppID: QM_mbw
JMSXDeliveryCount: 1
JMS_IBM_PutDate: 20070815
JMS_IBM_ConnectionID: 414D5143514D5F6D627720202020202001EDB14620006601
JMS_IBM_PutTime: 12552020
JMS_IBM_Format: MQSTR
JMS_IBM_MsgType: 8
```

```
A simple text message from the MQJMSPSIVT program
Reply string equals original string
Closing TopicSubscriber
Closing TopicPublisher
Closing Session
Closing Connection
PSIVT finished
```

The publish/subscribe installation verification test using JNDI

In this test, the IVT program uses JNDI to retrieve administered objects from a directory service.

Before you can run the test, you must configure a directory service that is based on a Lightweight Directory Access Protocol (LDAP) server or the local file system. You must also configure the WebSphere MQ JMS administration tool so that it can use the directory service to store administered objects. For more information about these prerequisites, see [“Prerequisites for WebSphere MQ classes for JMS” on page 686](#). For information about how to configure the WebSphere MQ JMS administration tool, see [“Configuring the JMS Administration tool” on page 894](#).

The IVT program must be able to use JNDI to retrieve an MQTopicConnectionFactory object and an MQTopic object from the directory service. A script is provided to create these administered objects for you. The script is called IVTSetup on UNIX and Linux systems and IVTSetup.bat on Windows, and is in the bin subdirectory of the WebSphere MQ classes for JMS installation directory. To run the script, enter the following command:

```
IVTSetup
```

The script invokes the WebSphere MQ JMS administration tool to create the administered objects.

The MQTopicConnectionFactory object is bound with the name ivtTCF and is created with the default values for all its properties, which means that the IVT program runs in bindings mode, connects to the default queue manager, and uses the embedded publish/subscribe function. If you want the IVT program to run in client mode, connect to a queue manager other than the default queue manager, or use WebSphere Event Broker or WebSphere Message Broker instead of the embedded publish/subscribe function, you must use the WebSphere MQ JMS administration tool or WebSphere MQ Explorer to change the appropriate properties of the MQTopicConnectionFactory object. For information about how to use the WebSphere MQ JMS administration tool, see [“Using the WebSphere MQ JMS administration tool” on page 893](#). For information about how to use WebSphere MQ Explorer, see the help provided with WebSphere MQ Explorer.

The MQTopic object is bound with the name ivtT and is created with the default values for all its properties, except for the TOPIC property, which has the value MQJMS/PSIVT/Information.

When you have created the administered objects, you can run the IVT program. To run the test using JNDI, enter the following command:

```
PSIVTRun -url "providerURL" [-icf initCtxFact] [-t]
```

The parameters on the command have the following meanings:

-url "providerURL"

The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName`, for a directory service based on an LDAP server
- `file:/directoryPath`, for a directory service based on the local file system

Note that you must enclose the URL in quotation marks (").

-icf initCtxFact

The class name of the initial context factory, which must be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server. This is the default value.
- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system.

-t

Tracing is switched on. By default, tracing is switched off.

A successful test produces output similar to that for a successful test without using JNDI. The main difference is that the output indicates that the test is using JNDI to retrieve an MQTopicConnectionFactory object and an MQTopic object.

Although not strictly necessary, it is good practice to tidy up after the test by deleting the administered objects created by the IVTSetup script. A script is provided for this purpose. The script is called IVTTidy on UNIX and Linux systems and IVTTidy.bat on Windows, and is in the bin subdirectory of the WebSphere MQ classes for JMS installation directory.

Problem determination for the publish/subscribe installation verification test

The installation verification test might fail for the following reasons:

- If the IVT program writes a message indicating that it cannot find a class, check that your class path is set correctly, as described in [“Environment variables used by IBM WebSphere MQ classes for JMS”](#) on page 691.
- The test might fail with the following message:

```
Failed to connect to queue manager 'qmgr' with
connection mode 'connMode' and host name 'hostname'
```

and an associated reason code of 2059. The variables in the message have the following meanings:

qmgr

The name of the queue manager to which the IVT program is trying to connect. This message insert is blank if the IVT program is trying to connect to the default queue manager in bindings mode.

connMode

The connection mode, which is either Bindings or Client.

hostname

The host name or IP address of the system on which the queue manager is running.

This message means that the queue manager to which the IVT program is trying to connect is not available. Check that the queue manager is running and, if the IVT program is trying to connect to the default queue manager, make sure that the queue manager is defined as the default queue manager for your system.

- The test might fail with the following message:

```
Unable to bind to object
```

This message means that there is a connection to the LDAP server, but that the LDAP server is not correctly configured. Either the LDAP server is not configured for storing Java objects, or the permissions on the objects or the suffix are not correct. For more help in this situation, see the documentation for your LDAP server.

- The test might fail with the following message:

```
The security authentication was not valid that was supplied for
QueueManager 'qmgr' with connection mode 'Client' and host name 'hostname'
```

This message means that the queue manager is not set up correctly to accept a client connection from your system. For more information, see [“Preparing and running the sample programs”](#) on page 104.

The installation verification test program for the WebSphere MQ resource adapter

The IVT program is supplied as an EAR file. To use the program, you must deploy it and define some objects as JCA resources.

The installation verification test (IVT) program is supplied as an enterprise archive (EAR) file called `wmq.jmsra.ivt.ear`. This file is installed with WebSphere MQ classes for JMS in the same directory as the WebSphere MQ resource adapter RAR file, `wmq.jmsra.rar`. For information about where these files are installed, see [“Installation of the WebSphere MQ resource adapter”](#) on page 702.

You must deploy the IVT program on your application server. The IVT program includes a servlet and an MDB that tests that a message can be sent to, and received from, a WebSphere MQ queue. Optionally, you can use the IVT program to verify that the WebSphere MQ resource adapter has been correctly configured to support distributed transactions.

Before you can run the IVT program, you must define a `ConnectionFactory` object, a `Queue` object and possibly an `Activation Specification` object as JCA resources, and ensure that your application server creates JMS objects from these definitions and binds them into a JNDI namespace. You can choose the properties of the objects, but the following set of properties is a simple example:

ConnectionFactory object

```
channel:          SYSTEM.DEF.SVRCONN
hostName:         localhost
port:            1414
queueManager:     ExampleQM
transportType:    CLIENT
```

Queue object

```
baseQueueManagerName: ExampleQM
baseQueueName:        TEST.QUEUE
```

By default, the IVT program expects a ConnectionFactory object to be bound in the JNDI namespace with the name `.jms/ivt/IVTCF` and a Queue object to be bound with the name `.jms/ivt/IVTQueue`. You can use different names, but if you do, you must enter the names of the objects on the initial page of the IVT program and modify the EAR file appropriately.

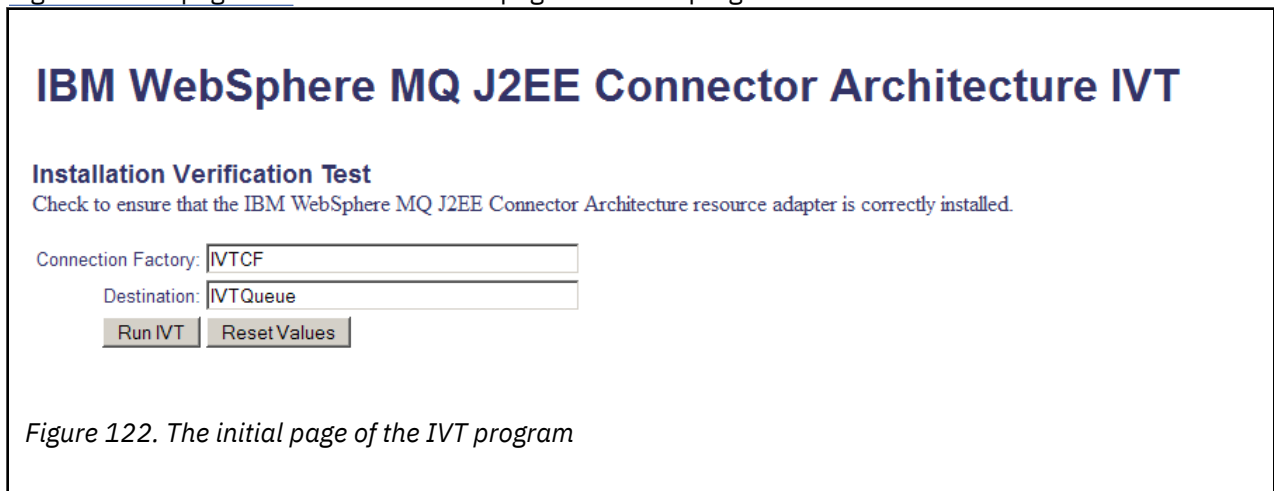
After you have deployed the IVT program, and the application server has created the JMS objects and bound them into the JNDI namespace, you can start the IVT program by entering a URL in the following format into your Web browser:

```
http://app_server_host:port/WMQ_IVT/
```

where *app_server_host* is the IP address or host name of the system on which your application server is running, and *port* is the number of the TCP port on which the application server is listening. Here is an example:

```
http://localhost:9080/WMQ_IVT/
```

Figure 122 on page 748 shows the initial page of the IVT program.



IBM WebSphere MQ J2EE Connector Architecture IVT

Installation Verification Test
Check to ensure that the IBM WebSphere MQ J2EE Connector Architecture resource adapter is correctly installed.

Connection Factory:

Destination:

Figure 122. The initial page of the IVT program

To run the test, click **Run IVT**. Figure 123 on page 749 shows the page that is displayed if the IVT is successful.

IBM WebSphere MQ J2EE Connector Architecture IVT

Running Installation Verification Test:

Using Connection Factory: `java:comp/env/IVTCF`
Using Destination: `java:comp/env/IVTQueue`

Creating initial context...	✓
Looking up MQ Connection Factory...	✓
Looking up Destination...	✓
Creating connection...	✓
Starting connection...	✓
Creating session...	✓
Creating a temporary reply queue...	✓
Creating message consumer...	✓
Creating message producer...	✓
Creating message...	✓
Sending message to the MDB...	✓
Receiving response message from the MDB...	✓
Closing connection...	✓

Installation Verification Test completed successfully!

[View Message Contents](#)

[Re-run Installation Verification Test](#)

Figure 123. Page showing the results of a successful IVT

If the IVT fails, a page like that shown in [Figure 124 on page 750](#) is displayed. To obtain further information about the cause of the failure, click **View Stack Trace**.

IBM WebSphere MQ J2EE Connector Architecture IVT

Running Installation Verification Test:

Using Connection Factory: `java:comp/env/IVTCF`
Using Destination: `java:comp/env/IVTQueue`

Creating initial context...	⊙
Looking up MQ Connection Factory...	⊙
Looking up Destination...	⊙
Creating connection...	⊙
Starting connection...	⊙
Creating session...	⊙
Creating a temporary reply queue...	⊙
Creating message consumer...	⊙
Creating message producer...	⊙
Creating message...	⊙
Sending message to the MDB... failed to send message!	⊗

Installation Verification Test failed!

Error received - JMS Exception:

```
com.ibm.msg.client.jms.DetailedIllegalStateException: JMSWMQ2007: Failed to send a message to destination 'TEST.QUEUE'.
```

JMS attempted to perform an MQPUT or MQPUT1; however WebSphere MQ reported an error.

Use the linked exception to determine the cause of this error.

[View Stack Trace](#)

Installation Verification Test failed!

[Retry Installation Verification Test](#)
[Change IVT parameters](#)

Figure 124. Page showing the results of an IVT that failed

For detailed instructions and information on utility scripts provided to deploy the IVT application on JBoss and WAS CE application servers see:

Related tasks

[“Installing and testing the MQ resource adapter in WAS CE” on page 750](#)

Installing the IBM WebSphere MQ resource adapter and running the installation verification test (IVT) application in WebSphere Application Server CE.

[“Installing and testing the resource adapter in JBoss AS 5.1 and 6” on page 753](#)

After installing the IBM WebSphere MQ resource adapter on JBoss AS 5.1 or 6, you can test the installation of the resource adapter by installing and running the installation verification test (IVT) application.

Installing and testing the MQ resource adapter in WAS CE

Installing the IBM WebSphere MQ resource adapter and running the installation verification test (IVT) application in WebSphere Application Server CE.

Before you begin

This task assumes that you have a running WebSphere Application Server CE server and that you are familiar with standard administration tasks for it. This task also assumes that you have a IBM WebSphere MQ installation on your local system and that you are familiar with standard administration tasks.

If you are using the resource adapter to connect to a IBM WebSphere MQ client and need to perform distributed XA transactions, you must follow the additional steps marked **Client XA only**.

1. Create a queue manager called ExampleQM, and set it up as described in [“Preparing and running the sample programs”](#) on page 104 noting that the listener should be started on port 1414, the channel to be used is called SYSTEM.DEF.SVRCONN and the queue used by the IVT application is named TEST.QUEUE. The model queue SYSTEM.DEFAULT.MODEL.QUEUE will also need to be granted DSP and PUT authority so this application can create a temporary reply queue. If you want to use a different queue manager, different connection details, or a different queue, see [“Deploying the IVT application on WAS CE with a custom MQ environment”](#) on page 752.
2. Obtain the resource adapter file (wmq.jmsra.rar), the IVT application (wmq.jmsra.ivt.ear), and the WAS_CE_jmsra_deployment_plan.xml and WAS_CE_jmsra_ivt_deployment_plan.xml deployment plan files. For details about the location of these files, see [“Installation of the WebSphere MQ resource adapter”](#) on page 702.

For a description of bindings and client mode connections, see [“Connection modes for WebSphere MQ classes for JMS”](#) on page 738.

If you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, see [“Deploying the IVT application on WAS CE with a custom MQ environment”](#) on page 752.

Procedure

1. **Client XA only:** Edit your copy of the WAS_CE_jmsra_deployment_plan.xml file.
 - a) Find the jms/ivt/IVTCF connection definition and modify it so that the connection factory is XA-transaction enabled.

- i) Comment out the NonXA section:

```
<conn:xa-transaction>
```

- ii) Uncomment the XA configuration section:

```
<conn:xa-transaction>
  <conn:transaction-caching/>
</conn:xa-transaction>
```

- b) Save your changes.
2. Optional: **Client XA only:** Modify the assembly descriptor of the MDB to require transactions. This forces the MDB in the IVT to participate in an XA transaction, although the IVT application still works without this modification.
 - a) Open the wmq.jmsra.ivt.ear file.
 - b) Open the WMQ_IVT_MDB.jar file within it.
 - c) Edit META-INF/ejb-jar.xml.

- i) Comment out or delete the line within the assembly descriptor:

```
<trans-attribute>NotSupported</trans-attribute>
```

- ii) Uncomment the line within the assembly descriptor:

```
<trans-attribute>Required</trans-attribute>
```

- iii) Save your changes and update the file within the WMQ_IVT_MDB.jar file.
 - iv) Update the wmq.jmsra.ivt.ear file with the modified WMQ_IVT_MDB.jar file.
3. Deploy the resource adapter to your server using the modified deployment plan file.
 - a) To do this on the command line, type the following WAS CE command:

```
deploy -u system -p manager deploy wmq.jmsra.rar WAS_CE_jmsra_deployment_plan.xml
```

- b) Using the web administration interface, go to **Applications > Deployer:**

- i) Set the Archive to be the `wmq.jmsra.rar` file.
 - ii) Set the Plan to be the `WAS_CE_jmsra_deployment_plan.xml` file.
 - iii) Ensure that 'Start application after install' is selected.
 - iv) Click **Install**.
- 4. Deploy the IVT application to your server using the provided deployment plan.
 - a) On the command line, this can be done using the following WAS CE command:


```
deploy -u system -p manager deploy wmq.jmsra.ivt.ear WAS_CE_jmsra_ivt_deployment_plan.xml
```
 - b) Using the web administration interface, go to **Applications > Deployer**:
 - i) Set the **Archive** to be the `wmq.jmsra.ivt.ear` file.
 - ii) Set the **Plan** to be the `WAS_CE_jmsra_ivt_deployment_plan.xml` file.
 - iii) Ensure that **Start application after install** is selected.
 - iv) Click **Install**.
- 5. Run the IVT application. For more details, see [“The installation verification test program for the WebSphere MQ resource adapter” on page 747](#). For WAS CE the default URL is `http://localhost:8080/WMQ_IVT/`.

Deploying the IVT application on WAS CE with a custom MQ environment

If you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, you must modify the IVT application and associated scripts in WebSphere Application Server CE before deploying the resource adapter or IVT application.

About this task

If you want to deploy to a different configuration from the one that is specified in [“Installing and testing the MQ resource adapter in WAS CE” on page 750](#), that is, if you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, then do the following steps before deploying the resource adapter or IVT application.

Procedure

1. If you want to specify a different queue manager and queue to use for the IVT application, set values for the queue manager and queue in `WAS_CE_jmsra_deployment_plan.xml`, for details, see [“Setting values for the queue manager and queue” on page 753](#).
2. If you want to specify a different queue manager and queue in the configuration for the message-driven bean (MDB), set values for the queue manager and queue that you are using in `WAS_CE_jmsra_ivt_deployment_plan.xml`, for details, see [“Setting values for the MDB configuration” on page 753](#).
3. If you are configuring the resource adapter to connect to IBM WebSphere MQ in bindings mode, ensure that the JNI libraries are on the system path or the path for WAS CE. For more information, see [“Installing and testing the MQ resource adapter in WAS CE” on page 750](#).
4. If you have already deployed the resource adapter, you can redeploy it with the modified deployment plan to change the settings using the following command:

```
deploy --user system --password manager redeploy wmq.jmsra.rar
WAS_CE_jmsra_deployment_plan.xml
```

What to do next

Continue deploying the resource adapter as described in [“Installing and testing the MQ resource adapter in WAS CE” on page 750](#).

Setting values for the queue manager and queue

Explains how to set values for the queue manager and queue that you are using in `WAS_CE_jmsra_deployment_plan.xml`.

Procedure

In `WAS_CE_jmsra_deployment_plan.xml`, set values for the queue manager and queue that you are using for the IVT application.

For the `jms/ivt/IVTCF` connection definition:

1. Set the value of the `queueManager` element to be the name of your queue manager.
2. If you are using a client connection, set the value of the various client connection elements to be appropriate for a connection to your queue manager.
3. If you are using a bindings connection:
 - a. Set the value of the `transportType` element to be `BINDINGS`.
 - b. Comment out or delete the various client connection elements.
4. For the `jms/ivt/IVTQueue` message destination, set the value of the `baseQueueName` element to be the name of the queue you have created for the IVT application
5. Save your changes.

Setting values for the MDB configuration

Explains how to set values for the MDB configuration in `WAS_CE_jmsra_deployment_plan.xml`.

Procedure

In `WAS_CE_jmsra_ivt_deployment_plan.xml`, set values for the queue manager and queue that you are using in the configuration for the MDB.

For the `WMQ_IVT_MDB` message driven bean:

1. Set the value of the `queueManager` element to be the name of your queue manager.
2. If you are using a client connection, set the value of the various client connection elements to be appropriate for a connection to your queue manager.
3. If you are using a bindings connection:
 - a. Set the value of the `transportType` element to be `BINDINGS`.
 - b. Comment out or delete the various client connection elements.
4. Save your changes.

Installing and testing the resource adapter in JBoss AS 5.1 and 6

After installing the IBM WebSphere MQ resource adapter on JBoss AS 5.1 or 6, you can test the installation of the resource adapter by installing and running the installation verification test (IVT) application.

Before you begin

Important: These instructions are for JBoss AS 5.1 and 6, they are not valid for JBoss AS 7.

For information about installing the resource adapter in JBoss EAP 6.3, see [“Installing and testing the resource adapter in JBoss EAP 6.3” on page 756](#).

This task assumes that you have a running JBoss server and are familiar with standard administration tasks for it. This task also assumes that you have an IBM WebSphere MQ installation on your local system and that you are familiar with standard administration tasks.

If you are using the resource adapter to connect to an IBM WebSphere MQ client and need to perform distributed XA transactions, you must follow the additional steps marked **Client XA only**. For a description

of bindings and client mode connections, see [“Connection modes for WebSphere MQ classes for JMS” on page 738](#).

Procedure

1. Create a queue manager called ExampleQM, and set it up as described in [“Preparing and running the sample programs” on page 104](#).

When setting up the queue manager, note the following points:

- The listener must be started on port 1414.
- The channel to be used is called SYSTEM.DEF.SVRCONN.
- The queue used by the IVT application is named TEST.QUEUE.

The model queue SYSTEM.DEFAULT.MODEL.QUEUE also needs to be granted DSP and PUT authority so this application can create a temporary reply queue.

If you want to use a different queue manager, different connection details, or a different queue, see [“Deploying the IVT application on WAS CE with a custom MQ environment” on page 752](#).

2. Obtain the resource adapter file (wmq.jmsra.rar), the IVT application (wmq.jmsra.ivt.ear) and the jboss-jmsra-ds.xml file.

For the location of these files, see [“Installation of the WebSphere MQ resource adapter” on page 702](#).

3. **Client XA only:** Edit the jboss-jmsra-ds.xml file to enable XA Transactions on the connection factory.
 - a) Comment out or delete the line within the connection factory definition `<local-transaction/>`.
 - b) Uncomment the line within the connection factory definition `<xa-transaction/>`.
 - c) Save your changes.
4. **Client XA only:** (optional) Modify the assembly descriptor of the MDB to require transactions. This forces the MDB in the IVT to participate in an XA transaction, although the IVT application still works without this modification.
 - a) Open the wmq.jmsra.ivt.ear file.
 - b) Open the WMQ_IVT_MDB.jar file within it.
 - c) Edit META-INF/ejb-jar.xml:
 - i) Comment out or delete the line within the assembly descriptor:

```
<trans-attribute>NotSupported</trans-attribute>
```

- ii) Uncomment the line within the assembly descriptor:

```
<trans-attribute>Required</trans-attribute>
```

- iii) Save your changes and update the file within the WMQ_IVT_MDB.jar file.
 - iv) Update the wmq.jmsra.ivt.ear file with the modified WMQ_IVT_MDB.jar.
5. Deploy the resource adapter to your server by copying the wmq.jmsra.rar file into the directory jboss/server/default/deploy.
 6. Create the JMS resources required for the IVT Application by copying the jboss-jmsra-ds.xml file into the directory jboss/server/default/deploy.
 7. Deploy the IVT application by copy the wmq.jmsra.ivt.ear file into the directory jboss/server/default/deploy.
 8. Run the IVT application. For more details, see [“The installation verification test program for the WebSphere MQ resource adapter” on page 747](#). For JBoss, the default URL is `http://localhost:8080/WMQ_IVT/`.

Deploying the IVT application in JBoss with a custom IBM WebSphere MQ environment

When installing the IBM WebSphere MQ resource adapter in JBoss, if you want to use a different queue, queue manager, port, host, channel, or use bindings mode instead of client mode, you must first modify the IVT application and associated scripts in JBoss before deploying the resource adapter or IVT application.

About this task

Important: These instructions are only applicable for Java EE Versions 6 and 5, not for Java EE Version 7. The use of these instructions for JBoss Version 8 (WildFly) is therefore not supported.

If you want to deploy to a different configuration from that specified in [“Installing and testing the resource adapter in JBoss AS 5.1 and 6”](#) on page 753, that is, if you want to use a different queue manager, queue, port, host, channel, or use bindings mode instead of client mode, then complete the following steps before deploying the resource adapter or IVT application.

Procedure

1. If you want to specify a different queue manager and queue to use for the IVT application, set values for the queue manager and queue.
 - a) For the `jms/ivt/IVTCF` connection definition:
 - i) Set the value of the `queueManager` config-property to be the name of your queue manager.
 - ii) If you are using a client connection, set the value of the various client connection elements to be appropriate for a connection to your queue manager.
 - iii) If you are using a bindings connection, set the value of the `transportType` element to be `BINDINGS`, and then comment out or delete the various client connection elements.
 - b) For the `jms/ivt/IVTQueue` mbean, set the value of the `baseQueueName` element to be the name of the queue you have created for the IVT application.
 - c) Save your changes.
2. If you want to specify a different queue manager and queue in the configuration for the message-driven bean (MDB), modify the configuration of the MDB to connect to the queue manager and queue.
 - a) Open the `wmq.jmsra.ivt.ear` file.
 - b) Open the `WMQ_IVT_MDB.jar` within it.
 - c) Edit `META-INF/ejb-jar.xml`:
 - i) Set the value of the `queueManager` activation-config-property to be the name of your queue manager.
 - ii) If you are using a client connection, set the value of the various client connection activation-config-properties to be appropriate for a connection to your queue manager.
 - iii) If you are using a bindings connection, set the value of the `transportType` activation-config-property to be `BINDINGS`, and then comment out or delete the various client connection elements.
 - d) Save the changes and update the file within the `WMQ_IVT_MDB.jar` file.
 - e) Update the `wmq.jmsra.ivt.ear` file with the modified `WMQ_IVT_MDB.jar`.
3. If you are configuring the resource adapter to connect to IBM WebSphere MQ in bindings mode, ensure that the JNI libraries are on the system path, or the path for JBoss. For details, see [“Configuring the Java Native Interface \(JNI\) libraries”](#) on page 693.

What to do next

Continue deploying the resource adapter as described in [“Installing and testing the resource adapter in JBoss AS 5.1 and 6”](#) on page 753.

Installing and testing the resource adapter in JBoss EAP 6.3

After installing the IBM WebSphere MQ resource adapter in JBoss Enterprise Application Platform (EAP) 6.3, either on a standalone server or on a server running in a managed domain, you can test the installation of the resource adapter by installing and running the installation verification test (IVT) application.

About this task

Important: These instructions are for JBoss EAP 6.3 only. For information about installing the resource adapter in JBoss AS 5.1 and 6, see [“Installing and testing the resource adapter in JBoss AS 5.1 and 6” on page 753](#).

This task assumes that you have a running JBoss server and are familiar with standard administration tasks for it. This task also assumes that you have an IBM WebSphere MQ installation on your local system and that you are familiar with standard administration.

Procedure

1. Create a queue manager called ExampleQM, and set it up as described in [“Preparing and running the sample programs” on page 104](#).

When setting up the queue manager, note the following points:

- The listener must be started on port 1414.
- The channel to be used is called SYSTEM.DEF.SVRCONN.
- The queue used by the IVT application is named TEST.QUEUE.

The model queue SYSTEM.DEFAULT.MODEL.QUEUE also needs to be granted DSP and PUT authority so this application can create a temporary reply queue.

If you want to use a different queue manager, different connection details, or a different queue, see [“Deploying the IVT application on WAS CE with a custom MQ environment” on page 752](#).

2. Obtain the resource adapter file (`wmq.jmsra.rar`) and the IVT application (`wmq.jmsra.ivt.ear`). For the location of these files, see [“Installation of the WebSphere MQ resource adapter” on page 702](#).
3. Install the resource adapter and then test the installation by running the installation verification test (IVT) application:
 - If you are installing the resource adapter on a standalone server, see [“Installing and testing on a standalone server” on page 756](#).
 - If you are installing the resource adapter on a server running in a managed domain, see [“Installing and testing on a server running in a managed domain” on page 758](#).

Installing and testing on a standalone server

After installing the IBM WebSphere MQ resource adapter on JBoss EAP 6.3 on a standalone server, you can test the installation of the resource adapter by installing and running the installation verification test (IVT) application.

About this task

The information in this task is for installing and testing the resource adapter on a standalone server. If you are installing the resource adapter on a server running in a managed domain, see [“Installing and testing on a server running in a managed domain” on page 758](#).

Important: These instructions are for JBoss EAP 6.3 only. For information about installing the resource adapter in JBoss AS 5.1 and 6, see [“Installing and testing the resource adapter in JBoss AS 5.1 and 6” on page 753](#).

Procedure

1. Deploy the resource adapter to your server by copying the `wmq.jmsra.rar` file into the directory `<EAP_HOME>/standalone/deployments`.
2. Create the JMS resources required for the IVT Application by adding the following entries to the `<resource-adapters>` section of the `<EAP_HOME>/standalone/configuration/standalone-full.xml` file:

```
<resource-adapter id="wmq.jmsra">
  <archive>
    wmq.jmsra.rar
  </archive>
  <transaction-support>NoTransaction</transaction-support>
  <connection-definitions>
    <connection-definition
      class-name="com.ibm.mq.connector.outbound.ManagedConnectionFactoryImpl"
      jndi-name="java:jboss/jms/ivt/IVTCF"
      enabled="true"
      use-java-context="true"
      pool-name="IVTCF">
      <config-property name="port">
        1414
      </config-property>
      <config-property name="hostName">
        localhost
      </config-property>
      <config-property name="channel">
        SYSTEM.DEF.SVRCONN
      </config-property>
      <config-property name="transportType">
        CLIENT
      </config-property>
      <config-property name="queueManager">
        ExampleQM
      </config-property>
    </connection-definition>
  </connection-definitions>
  <admin-objects>
    <admin-object class-name="com.ibm.mq.connector.outbound.MQQueueProxy"
      jndi-name="java:jboss/jms/ivt/IVTQueue"
      pool-name="IVTQueue">
      <config-property name="baseQueueName">
        TEST.QUEUE
      </config-property>
    </admin-object>
  </admin-objects>
</resource-adapter>
```

3. Add the following information to your application server's start up parameters:

```
-Dcom.ibm.mq.connector.IVTMDBCFJNDIName=java:jboss/jms/ivt/IVTCF
```

4. Deploy the IVT application by copying the `wmq.jmsra.ivt.ear` file into the directory `<EAP_HOME>/standalone/deployments`.
5. Start the application server.
6. Run the IVT application.

For more information, see “The installation verification test program for the WebSphere MQ resource adapter” on page 747. For JBoss, the default URL is http://localhost:8080/WMQ_IVT/.

Note: The JNDI names used for the JMS resources required for running IVT application are as follows:

Connection Factory	: IVTCF
JNDI name	: java:jboss/jms/ivt/IVTCF
Destination	: IVTQueue
JNDI name	: java:jboss/jms/ivt/IVTQueue

When you launch the IVT application using the URL specified above, enter the JNDI names of these resources in their respective fields and then run the application by clicking **Run IVT**.

Installing and testing on a server running in a managed domain

After installing the IBM WebSphere MQ resource adapter on JBoss EAP 6.3 on a server running in a managed domain, you can test the installation of the resource adapter by installing and running the installation verification test (IVT) application.

About this task

The information in this task is for installing and testing the resource adapter on a server running in a managed domain. If you are installing the resource adapter on a standalone server, see [“Installing and testing on a standalone server”](#) on page 756.

Important: These instructions are for JBoss EAP 6.3 only. For information about installing the resource adapter in JBoss AS 5.1 and 6, see [“Installing and testing the resource adapter in JBoss AS 5.1 and 6”](#) on page 753.

Procedure

1. Deploy the resource adapter to your server using the JBoss Management Console or Management CLI.
2. Create the JMS resources required for the IVT Application by adding the following entries to the <resource-adapters> section of the <EAP_HOME>/domain/configuration/domain.xml file:

```
<resource-adapter id="wmq.jmsra">
  <archive>
    wmq.jmsra.rar
  </archive>
  <transaction-support>NoTransaction</transaction-support>
  <connection-definitions>
    <connection-definition
      class-name="com.ibm.mq.connector.outbound.ManagedConnectionFactoryImpl"
      jndi-name="java:jboss/jms/ivt/IVTCF"
      enabled="true"
      use-java-context="true"
      pool-name="IVTCF">
      <config-property name="port">
        1414
      </config-property>
      <config-property name="hostName">
        localhost
      </config-property>
      <config-property name="channel">
        SYSTEM.DEF.SVRCONN
      </config-property>
      <config-property name="transportType">
        CLIENT
      </config-property>
      <config-property name="queueManager">
        ExampleQM
      </config-property>
    </connection-definition>
  </connection-definitions>
  <admin-objects>
    <admin-object class-name="com.ibm.mq.connector.outbound.MQQueueProxy"
      jndi-name="java:jboss/jms/ivt/IVTQueue"
      pool-name="IVTQueue">
      <config-property name="baseQueueName">
        TEST.QUEUE
      </config-property>
    </admin-object>
  </admin-objects>
</resource-adapter>
```

3. Add the following information to your application server's start up parameters:

```
-Dcom.ibm.mq.connector.IVTMBCFJNDIName=java:jboss/jms/ivt/IVTCF
```

4. Stop and restart the application server.
5. Deploy the IVT application using the JBoss Management Console or Management CLI.
6. Run the IVT application.

For more information, see [“The installation verification test program for the WebSphere MQ resource adapter” on page 747](#). For JBoss, the default URL is http://localhost:8080/WMQ_IVT/.

Note: The JNDI names used for the JMS resources required for running IVT application are as follows:

```
Connection Factory : IVTCF
JNDI name          : java:jboss/jms/ivt/IVTCF

Destination        : IVTQueue
JNDI name          : java:jboss/jms/ivt/IVTQueue
```

Enter the JNDI names of these resources in their respective fields when you launch the IVT application using the URL specified above and then run the application by clicking **Run IVT**.

Scripts provided with WebSphere MQ classes for JMS

A number of scripts are provided to assist with common tasks that need to be performed when using WebSphere MQ classes for JMS.

Table 102 on page 759 lists all the scripts and their uses. The scripts are in the bin subdirectory of the WebSphere MQ classes for JMS installation directory.

Table 102. Scripts provided with WebSphere MQ classes for JMS	
Utility	Use
Cleanup ¹	This script is maintained for compatibility with previous releases but performs no function. Manual cleanup of subscription information is no longer necessary
DefaultConfiguration	Runs the default configuration application on platforms other than Windows.
formatLog ¹	This script is maintained for compatibility with previous releases but performs no function. Log output is now produced in readable text.
IVTRun ¹ IVTSetup ¹ IVTTidy ¹	Used in the point-to-point installation verification test, as described in “The point-to-point installation verification test for WebSphere MQ classes for JMS” on page 740 .
JMSAdmin ¹	Runs the WebSphere MQ JMS administration tool, as described in “Invoking the IBM WebSphere MQ classes for JMS administration tool” on page 893 .
JMSAdmin.config	The configuration file for the WebSphere MQ JMS administration tool, as described in “Configuring the JMS Administration tool” on page 894 .
PSIVTRun ¹	Runs the publish/subscribe installation verification test program, as described in “The publish/subscribe installation verification test for WebSphere MQ classes for JMS” on page 744 .
PSReportDump.class	This class is maintained for compatibility with previous releases, but performs no function.
setjmsenv	Sets the environment variables for running a WebSphere MQ classes for JMS application in a 32-bit Java virtual machine (JVM) on UNIX and Linux systems, as described in “Environment variables used by IBM WebSphere MQ classes for JMS” on page 691 .
setjmsenv64	Sets the environment variables for running a WebSphere MQ classes for JMS application in a 64-bit JVM on UNIX and Linux systems, as described in “Environment variables used by IBM WebSphere MQ classes for JMS” on page 691 .

Table 102. Scripts provided with WebSphere MQ classes for JMS (continued)

Utility	Use
Note: 1. On Windows, the file name has the extension .bat .	

Support for OSGi

OSGi provides a framework that supports the deployment of applications as bundles. Nine OSGi bundles are supplied as part of the IBM WebSphere MQ classes for JMS .

OSGi provides a general purpose, secure, and managedJava framework, which supports the deployment of applications that come in the form of bundles. OSGi-compliant devices can download and install bundles, and remove them when they are no longer required. The framework manages the installation and update of bundles in a dynamic and scalable fashion.

The IBM WebSphere MQ classes for JMS. includes the following OSGi bundles.

com.ibm.msg.client.osgi.jms<version number>.jar

The common layer of code in the IBM WebSphere MQ classes for JMS. For information about the layered architecture of WebSphere MQ classes for JMS, see [“A layered architecture”](#) on page 763.

com.ibm.msg.client.osgi.jms.prereq_<version number>.jar

The prerequisite Java archive (JAR) files for the common layer.

com.ibm.msg.client.osgi.commonservices.j2se_<version number>.jar

Common services for Java Platform, Standard Edition (Java SE) applications.

com.ibm.msg.client.osgi.nls_<version number>.jar

Messages for the common layer.

com.ibm.msg.client.osgi.wmq_<version number>.jar

The IBM WebSphere MQ messaging provider in IBM WebSphere MQ classes for JMS. For information about the layered architecture of IBM WebSphere MQ classes for JMS , see [“A layered architecture”](#) on page 763.

com.ibm.msg.client.osgi.wmq.prereq_<version number>.jar

The prerequisite JAR files for the IBM WebSphere MQ messaging provider.

com.ibm.msg.client.osgi.wmq.nls_<version number>.jar

Messages for the IBM WebSphere MQ messaging provider.

com.ibm.mq.osgi.directip_<version number>.jar

The JAR files to allow the IBM WebSphere MQ messaging provider to create a real-time connection to a broker.

where <version number> is the version number of WebSphere MQ that has been installed.

The bundles are installed into the `java/lib/OSGi` subdirectory of your WebSphere MQ installation, or the `java\lib\OSGi` folder on Windows.

The bundle `com.ibm.mq.osgi.java <version number>.jar`, which is also installed into the `java/lib/OSGi` subdirectory of your WebSphere MQ installation, or the `java\lib\OSGi` folder on Windows, is part of the WebSphere MQ classes for Java. This bundle must not be loaded into an OSGi runtime environment that has the WebSphere MQ classes for JMS loaded.

The OSGi bundles for the WebSphere MQ classes for JMS have been written to the OSGi Release 4 specification. They do not work in an OSGi Release 3 environment.

You must set your system path or library path correctly so that the OSGi runtime environment can find any required DLL files or shared libraries.

If you use the OSGi bundles for the IBM WebSphere MQ classes for JMS, temporary topics do not work. In addition, channel exit classes written in Java are not supported because of an inherent problem in loading classes in a multiple class loader environment such as OSGi. A user bundle can be aware of the

IBM WebSphere MQ classes for JMS bundles, but the IBM WebSphere MQ classes for JMS bundles are not aware of any user bundle. As a result, the class loader used in a IBM WebSphere MQ classes for JMS bundle cannot load a channel exit class that is in a user bundle.

For more information about OSGi, see the [OSGi Alliance](#) website.

Solving problems with IBM WebSphere MQ classes for JMS

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 installation verification test for WebSphere MQ classes for JMS” on page 740 and “The publish/subscribe installation verification test for WebSphere MQ classes for JMS” on page 744, and follow the advice given in the diagnostic messages.

Logging and IBM WebSphere 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 WebSphere 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 `mqjms.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.
- `stdout` 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 WebSphere 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 WebSphere MQ classes for JMS libraries (see “Configuring the Java Native Interface (JNI) libraries” on page 693).

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.

Introduction to WebSphere MQ classes for JMS, for programmers

WebSphere MQ classes for JMS is the JMS provider that is supplied with WebSphere MQ. WebSphere MQ classes for JMS implements the interfaces defined in the `javax.jms` package, and also provides two sets of extensions to the JMS API. Both Java Platform, Standard Edition (Java SE) and Java Platform, Enterprise Edition (Java EE) applications can use WebSphere MQ classes for JMS.

The JMS specification defines a set of interfaces that applications can use to perform messaging operations. The latest version of the specification is Version 1.1. The `javax.jms` package specifies the

details of the JMS interfaces, and a JMS provider implements these interfaces for a specific messaging product. WebSphere MQ classes for JMS is a JMS provider that implements the JMS interfaces for WebSphere MQ.

The flow of logic within a JMS application starts with `ConnectionFactory` and `Destination` objects. The application uses a `ConnectionFactory` object to create a `Connection` object, which represents the active connection from the application to a messaging server. The application uses the `Connection` object to create a `Session` object, which is a single threaded context for producing and consuming messages. The application can then use the `Session` object and a `Destination` object to create a `MessageProducer` object, which the application uses to send messages to the specified destination. The destination is either a queue or a topic in the messaging system and is encapsulated by the `Destination` object. The application can also use the `Session` object and a `Destination` object to create a `MessageConsumer` object, which the application uses to receive messages that have been sent to the specified destination.

The JMS specification expects `ConnectionFactory` and `Destination` objects to be administered objects. An administrator creates and maintains administered objects in a central repository, and a JMS application retrieves these objects using the Java Naming and Directory Interface (JNDI). The repository of administered objects can range from a simple file to a Lightweight Directory Access Protocol (LDAP) directory.

WebSphere MQ classes for JMS supports the use of administered objects. An application can use all the features of WebSphere MQ that are exposed through WebSphere MQ classes for JMS without having any WebSphere MQ-specific information hard coded into the application itself. This arrangement provides the application with a degree of independence from the underlying WebSphere MQ configuration. To achieve this independence, the application can use JNDI to retrieve connection factories and destinations that are stored as administered objects, and use only the interfaces defined in the `javax.jms` package to perform messaging operations. An administrator can use the WebSphere MQ JMS administration tool or IBM WebSphere MQ Explorer to create and maintain administered objects in a central repository. An application server, however, typically provides its own repository for administered objects and its own tools for creating and maintaining the objects. A Java EE application can therefore use JNDI to retrieve administered objects either from the applications server repository or from a central repository.

WebSphere MQ classes for JMS also provides extensions to the JMS API. Previous releases of WebSphere MQ classes for JMS contain extensions that are implemented in `MQConnectionFactory`, `MQQueue`, and `MQTopic` objects. These objects have properties and methods that are specific to WebSphere MQ. The objects can be administered objects, or an application can create the objects dynamically at run time. This release of WebSphere MQ classes for JMS maintains these extensions, and you can continue to use, without change, any applications that use these extensions. These extensions are known as the *WebSphere MQ JMS extensions*. Note that, in this set of documentation, objects that are created dynamically by an application at run time are *not* considered to be administered objects.

In addition to the WebSphere MQ JMS extensions, this release of WebSphere MQ classes for JMS provides a more generic set of extensions to the JMS API. These extensions are known as the *IBM JMS extensions*, and have the following broad objectives:

- To provide a greater level of consistency across IBM JMS providers
- To make it easier to write a bridge application between two IBM messaging systems
- To make it easier to port an application from one IBM JMS provider to another

The main focus of these extensions concerns creating and configuring connection factories and destinations dynamically at run time, but the extensions also provide function that is not directly related to messaging, such as function for problem determination.

A connection factory, queue, or topic object created using the `javax.jms` interface or either set of JMS extensions can be addressed using any of these APIs; that is, it can be cast to any of the interfaces. To maintain application portability at the highest level, use the most generic API that is suitable for your requirements.

Both Java SE and Java EE applications can use WebSphere MQ classes for JMS. On the Java EE platform, WebSphere MQ classes for JMS supports two types of communication between a component of an application and a WebSphere MQ queue manager:

Outbound communication

Using the JMS API directly, an application component creates a connection to a queue manager, and then sends and receives messages.

For example, the application component can be an application client, a servlet, a JavaServer Page (JSP), an enterprise Java bean (EJB), or a message driven bean (MDB). In this type of communication, the application server container provides only low-level functions in support of messaging operations, such as connection pooling and thread management.

Inbound communication

A message arriving at a destination is delivered to an MDB, which then processes the message.

Java EE applications use MDBs to process messages asynchronously. An MDB acts as a JMS message listener and is implemented by an `onMessage()` method, which defines how a message is processed. An MDB is deployed in the EJB container of an application server. The precise way in which an MDB is configured depends on which application server you are using, but the configuration information must specify which queue manager to connect to, how to connect to the queue manager, which destination to monitor for messages, and the transactional behavior of the MDB. This information is then used by the EJB container. When a message satisfying the selection criteria of the MDB arrives at the specified destination, the EJB container uses WebSphere MQ classes for JMS to retrieve the message from the queue manager, and then delivers the message to the MDB by calling its `onMessage()` method.

IBM WebSphere MQ classes for JMS architecture

IBM WebSphere MQ classes for JMS, as supplied in IBM WebSphere MQ Version 7.0, and subsequent releases, contains a number of enhancements compared to previous releases. Some of these enhancements are as a result of changes to the implementation of IBM WebSphere MQ classes for JMS, and some are as a result of IBM WebSphere MQ classes for JMS exploiting changes to the underlying IBM WebSphere MQ function.

The following sections summarize the key enhancements.

A layered architecture

In previous releases of WebSphere MQ, the implementation of WebSphere MQ classes for JMS has been entirely specific to WebSphere MQ. Other IBM products that provide messaging systems have also included JMS providers, but these JMS providers have very little or nothing in common with the implementation of WebSphere MQ classes for JMS.

From WebSphere MQ V7.0, WebSphere MQ classes for JMS has a layered architecture. The top layer of code is a common layer that can be used by any IBM JMS provider. When an application calls a JMS method, any processing of the call that is not specific to a messaging system is performed by the common layer, which also provides a consistent response to the call. Any processing of the call that is specific to a messaging system is delegated to a lower layer. [Figure 125 on page 764](#) shows the layered architecture.




Figure 125. The layered architecture for IBM JMS providers

Moving to a layered architecture has following objectives:

- To improve the consistency of behavior of the various IBM JMS providers
- To make it easier to write an bridge application between two IBM messaging systems
- To make it easier to port an application from one IBM JMS provider to another

This implementation of WebSphere MQ classes for JMS also introduces a new set of extensions to the JMS API. These extensions are known as the *IBM JMS extensions*. The main focus of these extensions concerns creating and configuring connection factories and destinations dynamically at run time.

An application using the IBM JMS extensions starts by creating a `JmsFactoryFactory` object, specifying as a parameter a constant that identifies the chosen messaging system. The application uses the `JmsFactoryFactory` object to create connection factories and destinations that have the correct specialized classes for the chosen messaging system.

The application can then configure the connection factories and destinations by setting their properties. The IBM JMS extensions provide a set of methods to set properties. These methods are independent of any messaging system. Each data type has its own set method, and each property is identified by a name, which is defined as a static final member of the `WMQConstants` class. When an application calls one of these methods, one of the parameters on the call is the name of the property, and the other parameter is the value of the property.

For example, if WebSphere MQ is the messaging system, one of the properties of a connection factory is the name of the queue manager to connect to. Using the IBM JMS extensions, an application sets the name of the queue manager to JUPITER by calling the following method:

```
JmsConnectionFactory myCF;
...
myCF.setStringProperty(WMQConstants.WMQ_QUEUE_MANAGER, "JUPITER");
```

By contrast, an application can perform the same function by calling the following method:

```
MQConnectionFactory myCF;
...
myCF.setQueueManager("JUPITER");
```

This method is a WebSphere MQ JMS extension and is specific to WebSphere MQ as the messaging system. The use of this method therefore makes the application potentially less easy to port to another IBM JMS provider.

The relationship between WebSphere MQ classes for JMS and WebSphere MQ classes for Java

In releases of WebSphere MQ, prior to Version 7.0, WebSphere MQ classes for JMS was implemented almost entirely as a layer of code on top of WebSphere MQ classes for Java. This arrangement has caused some confusion among application developers because setting fields or calling methods in the `MQEnvironment` class can cause unwanted and unexpected effects on the runtime behavior of code that is written using WebSphere MQ classes for JMS. In addition, the implementation of WebSphere MQ classes for JMS had some constraints in areas where the JMS API is not a natural fit on top of WebSphere MQ classes for Java, and these constraints have led to some issues regarding runtime performance.

From WebSphere MQ V7.0, the implementation of WebSphere MQ classes for JMS is no longer dependent on WebSphere MQ classes for Java. WebSphere MQ classes for Java and WebSphere MQ classes for JMS are now peers that use a common Java interface to the MQI. This arrangement allows more scope for optimizing performance, and means that setting fields or calling methods in the `MQEnvironment` class has no effect on the runtime behavior of code that is written using WebSphere MQ classes for JMS. [Figure 126 on page 765](#) shows the relationship between WebSphere MQ classes for JMS and WebSphere MQ classes for Java in previous releases of WebSphere MQ and in WebSphere MQ V7.0 and subsequent releases.

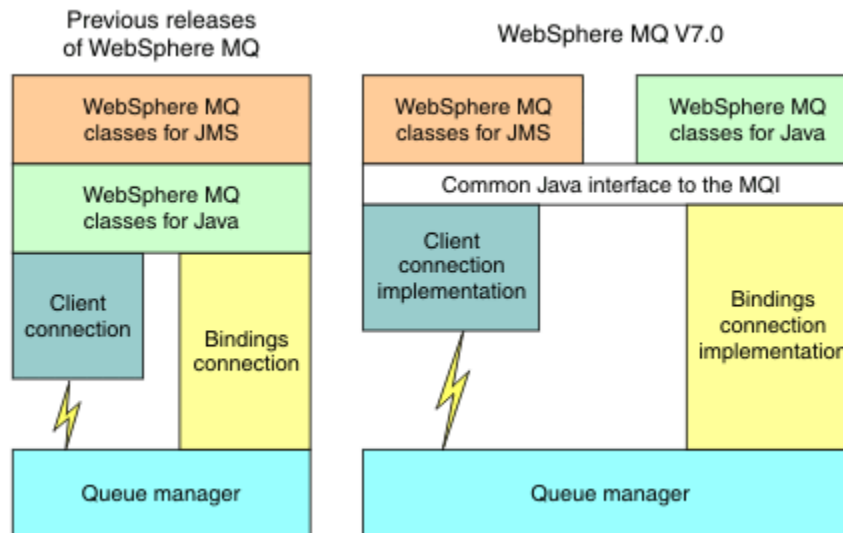


Figure 126. The relationship between WebSphere MQ classes for JMS and WebSphere MQ classes for Java

In order to maintain compatibility with earlier releases, channel exit classes written in Java can still use the WebSphere MQ classes for Java interfaces, even if the channel exit classes are called from WebSphere MQ classes for JMS. However, using the WebSphere MQ classes for Java interfaces means that your applications are still dependent on the WebSphere MQ classes for Java JAR file, `com.ibm.mq.jar`. If you do not want `com.ibm.mq.jar` in your class path, you can use the new set of interfaces in the `com.ibm.mq.exits` package instead.

You can now create and configure JMS administered objects with the WebSphere MQ Explorer.

Publish/subscribe messaging

WebSphere MQ V7.0, and subsequent releases, contain embedded publish/subscribe function. This function replaces WebSphere MQ Publish/Subscribe, which was supplied with WebSphere MQ V6.0.

WebSphere MQ classes for JMS applications can use the embedded publish/subscribe function, and can use it instead of using WebSphere Event Broker or WebSphere Message Broker for publish/subscribe messaging with WebSphere MQ as the transport. Configuring WebSphere MQ classes for JMS to use the new function is simpler than configuring WebSphere MQ classes for JMS to use WebSphere MQ Publish/Subscribe, WebSphere Event Broker, or WebSphere Message Broker. Administrators and application

developers no longer need to manage publication queues, subscriber queues, subscription stores, and subscriber cleanup. In addition, ConnectionFactory and Topic objects have a smaller number of properties.

The embedded publish/subscribe function also provides some additional features such as retained publications and a choice of two wildcard schemes for specifying a range of topics to which an application wants to subscribe.

An application can still use a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker for publish/subscribe messaging. This support is unchanged.

Applications using WebSphere MQ Publish/Subscribe can use the embedded publish/subscribe function without change when the queue manager to which they are connected is upgraded. Properties that are set by an application, but are not required by the embedded publish/subscribe function, are ignored.

WebSphere MQ messaging provider

The WebSphere MQ messaging provider has two modes of operation:

- *WebSphere MQ messaging provider normal mode*
- *WebSphere MQ messaging provider migration mode*

The WebSphere MQ messaging provider normal mode uses all the features of the WebSphere MQ Version 7.0, and subsequent release queue managers to implement JMS. This mode is used only to connect to a WebSphere MQ queue manager and can connect to WebSphere MQ Version 7.0, and subsequent release queue managers in either client or bindings mode. This mode is optimized to use the new WebSphere MQ Version 7.0 and subsequent release function.

The WebSphere MQ messaging provider migration mode is based on WebSphere MQ Version 6.0 function and uses only features that were available in the WebSphere MQ Version 6.0 queue manager to implement JMS. You can connect to a WebSphere MQ Version 7.0 and subsequent release queue managers using WebSphere MQ messaging provider migration mode but you cannot use any of the Version 7.0 optimizations. This mode allows connections to either of the following queue manager versions:

1. WebSphere MQ Version 7.0, and subsequent, queue manager in bindings or client mode, but this mode uses only those features that were available to a WebSphere MQ Version 6.0 queue manager
2. WebSphere MQ Version 6.0 or earlier queue manager in client mode

If you want to connect to WebSphere Event Broker or WebSphere Message Broker using either WebSphere MQ Enterprise Transport, use the WebSphere MQ messaging provider migration mode. If you use WebSphere MQ Real-Time Transport, the WebSphere MQ messaging provider migration mode is automatically selected, because you have explicitly selected properties in the connection factory object. Connection to WebSphere Event Broker or WebSphere Message Broker using the WebSphere MQ Enterprise Transport follows the general rules for mode selection described in [Rules for selecting the WebSphere MQ messaging provider mode](#).

Asynchronous message consumption

WebSphere MQ V7.0 and any subsequent release supports asynchronous message consumption. An application can register a callback function for a destination. When a suitable message is sent to the destination, WebSphere MQ calls the function and passes the message as a parameter. The function then processes the message asynchronously. In previous releases of WebSphere MQ, this feature was available only when using WebSphere MQ classes for JMS.

WebSphere MQ classes for JMS has been changed to exploit this new feature in WebSphere MQ V7.0 and any subsequent release. The implementation of JMS message listeners is now a more natural fit with WebSphere MQ, and WebSphere MQ classes for JMS no longer has to poll a destination to check whether a suitable message has been sent to the destination. The performance of JMS message listeners is improved as a result, particularly when an application uses multiple message listeners in a session to

monitor multiple destinations. Message throughput is increased, and the time taken to deliver a message to a message listener after it has arrived at a destination is reduced.

Message driven beans (MDBs) have similar performance improvements. In addition, because of another enhancement to WebSphere MQ function, multiple MDBs that are consuming messages from the same destination now experience reduced contention on the messages.

Message selection

With the exception of selecting messages by message identifier or correlation identifier, all message selection in releases of WebSphere MQ prior to Version 7.0 was done by WebSphere MQ classes for JMS. In WebSphere MQ V7.0, and any subsequent release, all message selection is done by the queue manager.

As a result, message throughput is increased for applications that consume messages using message selection. The performance improvement is greater for an application that connects in client mode because only those messages that satisfy the selection criteria are transported over the network, and WebSphere MQ classes for JMS handles only those messages that it delivers to the application.

Sharing a communications connection

In previous releases of WebSphere MQ, if a WebSphere MQ client application connected to a queue manager more than once using the same MQI channel, each instance of the MQI channel required a separate TCP connection. In WebSphere MQ V7.0 and any subsequent release, each connection to the queue manager using the same MQI channel can share a single TCP connection. This arrangement means that fewer network resources are required and the total time taken to create multiple connections to the queue manager is reduced, particularly when using SSL because the SSL handshake takes place only once at the start of the TCP connection.

WebSphere MQ classes for JMS exploits this enhancement. For an application that connects to a queue manager in client mode, WebSphere MQ classes for JMS might create more than one connection to a queue manager using the MQI channel with the name that is specified as a property of the ConnectionFactory object. Each of these connections to the queue manager can now share a single TCP connection.

Read ahead on client connections

If an application uses a client connection to consume nonpersistent messages from a destination, the destination can be configured so that WebSphere MQ classes for JMS uses a buffer to store the messages of interest before delivering them to the application. This optimization is called *read ahead* and can be used by applications that consume messages synchronously by calling the `receive()` method, and by message listeners and MDBs, which consume messages asynchronously. Read ahead is particularly effective for destinations with a large number of messages that need to be consumed rapidly.

Read ahead does not apply to persistent messages because, if persistent messages were read into a buffer, the queue manager would no longer be able to recover the messages following a failure. However, an application that consumes messages from a destination with a mixture of persistent and nonpersistent messages can still use read ahead. The order of the messages is preserved, but the runtime benefits of read ahead apply only to the nonpersistent messages.

When deciding whether to use read ahead, consider the following points:

- If an application is consuming messages from a destination that is configured for read ahead, and the application ends for any reason, any nonpersistent messages that are currently stored in the buffer are discarded.
- If all the following conditions are true, messages sent to a queue in a session might not be received in the order in which they were sent:
 - An application uses two message consumers in the same session to consume the messages from the queue.

- Each message consumer uses a different Destination object for the queue.
- Any or both of the Destination objects are configured for read ahead.

Sending messages

When an application sends messages to a destination, the destination can be configured so that, when the application calls `send()`, WebSphere MQ classes for JMS forwards the message to the queue manager and returns control back to the application without determining whether the queue manager has received the message safely. WebSphere MQ classes for JMS can work in this way only for nonpersistent messages and for persistent messages sent in a transacted session.

For persistent messages sent in a transacted session, the application ultimately determines whether the queue manager has received the messages safely when it calls `commit()`. For any messages sent in a session that is not transacted, the `SENDCHECKCOUNT` property of the `ConnectionFactory` object specifies how many messages are to be sent before WebSphere MQ classes for JMS checks that the queue manager has received the messages safely.

This optimization is of most benefit to an application that connects to a queue manager in client mode and needs to send a sequence of messages in rapid succession, but does not require immediate feedback from the queue manager for each message sent.

Channel exits

When called from WebSphere MQ classes for JMS, channel exit programs written in C or C++ now behave in the same way as when they are called from a WebSphere MQ MQI client. The performance of channel exit classes written in Java has been improved, and you can now write channel exit classes using a new set of interfaces in the `com.ibm.mq.exits` package instead of using the interfaces in WebSphere MQ classes for Java.

Message properties

A JMS message consists of a set of header fields, a set of properties, and a body that contains the application data. As a minimum, a WebSphere MQ message consists of a message descriptor and the application data.

When a WebSphere MQ classes for JMS application sends a JMS message, WebSphere MQ classes for JMS maps the JMS message into a WebSphere MQ message. Some of the JMS header fields and properties are mapped into fields in the message descriptor, and some are mapped into fields in an additional WebSphere MQ header called an MQRFH2 header. When a WebSphere MQ classes for JMS application receives a JMS message, WebSphere MQ classes for JMS performs the reverse mapping.

An application that is using the MQI to receive messages from a WebSphere MQ classes for JMS application must therefore be able to handle an MQRFH2 header. If the application cannot handle an MQRFH2 header, the `TARGETCLIENT` property of the Destination object can be set to tell WebSphere MQ classes for JMS not to include an MQRFH2 header in the WebSphere MQ messages. However, by excluding the MQRFH2 header, the information held in some of the JMS header fields and properties is lost.

Similarly, an application that is using the MQI to send messages to a WebSphere MQ classes for JMS application must include an MQRFH2 header in each message. If an MQRFH2 header is not included, WebSphere MQ classes for JMS can set only those JMS header fields and properties that can be derived from the fields in a message descriptor.

WebSphere MQ V7.0 provides some additional support for applications that use the MQI to receive messages from, and send messages to, WebSphere MQ classes for JMS applications.

When an application calls `MQGET` to receive a message from a WebSphere MQ classes for JMS application, the application can choose to receive the message in one of the following ways:

1. The message is delivered with a message descriptor, an MQRFH2 header that contains data derived from JMS header fields and properties, and the application data.

2. The message is delivered with a message descriptor, the application data, and a set of message properties.

In option 2, each message property represents a JMS header field or property that was originally mapped by WebSphere MQ classes for JMS into a field in an MQRFH2 header. After the MQGET call, the application can use the MQINQMP call to get the values of the message properties. Using option 2 instead of option 1 to receive a message simplifies the application logic in the following ways:

- The application does not have to parse the variable portion of the MQRFH2 header, which contains the JMS header field and property data encoded in an XML-like format.
- The application does not have to convert the character data in the variable portion of the MQRFH2 header.

Correspondingly, before an application calls MQPUT to send a message to a WebSphere MQ classes for JMS application, the application can use the MQSETMP call to set the values of message properties instead of constructing an MQRFH2 header.

Serviceability

WebSphere MQ classes for JMS contains a number of improvements related to serviceability:

- Tracing.

WebSphere MQ classes for JMS contains a class that an application can use to control tracing. An application can start and stop tracing, specify the required level of detail in a trace, and customize trace output in various ways..

- Logging.

WebSphere MQ classes for JMS maintains a log file, which contains messages about errors that you need to correct. The messages are written in plain text. WebSphere MQ classes for JMS contains a class that an application can use to specify the location of the log file and its maximum size.

- First Failure Support Technology (FFST).

If a serious failure occurs, WebSphere MQ classes for JMS generates an FFST report in an FDC file. The FFST report contains information that IBM Service can use to diagnose the problem more quickly.

- Version information.

WebSphere MQ classes for JMS contains a class that an application can use to query the version of WebSphere MQ classes for JMS.

- Exception messages.

Exception messages have been enhanced to provide more information about the causes of errors and the actions required to correct errors.

- Application servers.

The integration of the serviceability features of WebSphere MQ classes for JMS with those of WebSphere Application Server has been improved.

MQC is replaced by MQConstants

A new package, `com.ibm.mq.constants`, is supplied with IBM WebSphere MQ Version 7.0. This package contains the class `MQConstants`, which implements a number of interfaces. `MQConstants` contains definitions of all the constants that were in the `MQC` interface and a number of new constants. The interfaces in this package closely follow the names of the constants header files used in IBM WebSphere MQ.

For example, the interface `CMQC` contains a constant `MQOO_INPUT_SHARED`; this interface and constant correspond to the header file `cmqc . h` and the constant `MQOO_INPUT_SHARED`.

`com.ibm.mq.constants` can be used with both IBM WebSphere MQ classes for Java and IBM WebSphere MQ classes for JMS.

MQC is still present, and has the constants it previously had. However, for any new applications, you must use the `com.ibm.mq.constants` package.

Writing WebSphere MQ classes for JMS applications

After a brief introduction to the JMS model, this topic provides detailed guidance on how to write WebSphere MQ classes for JMS applications.

The JMS model

The JMS model defines a set of interfaces that Java applications can use to perform messaging operations. WebSphere MQ classes for JMS, as a JMS provider, defines how JMS objects are related to WebSphere MQ concepts. The JMS specification expects certain JMS objects to be administered objects.

The JMS specification and the `javax.jms` package define a set of interfaces that Java applications can use to perform messaging operations. The following list summarizes the main JMS interfaces:

Destination

A destination is where an application sends messages, or it is a source from which an application receives messages, or both.

ConnectionFactory

A `ConnectionFactory` object encapsulates a set of configuration properties for a connection. An application uses a connection factory to create a connection.

Connection

A `Connection` object encapsulates an application's active connection to a messaging server. An application uses a connection to create sessions.

Session

A session is a single threaded context for sending and receiving messages. An application uses a session to create messages, message producers, and message consumers. A session is either transacted or not transacted.

Message

A `Message` object encapsulates a message that an application sends or receives.

MessageProducer

An application uses a message producer to send messages to a destination.

MessageConsumer

An application uses a message consumer to receive messages sent to a destination.

[Figure 127 on page 770](#) shows these objects and their relationships.




Figure 127. JMS objects and their relationships

A Destination, ConnectionFactory, or Connection object can be used concurrently by different threads of a multithreaded application, but a Session, MessageProducer, or MessageConsumer object cannot be used concurrently by different threads. The simplest way of ensuring that a Session, MessageProducer, or MessageConsumer object is not used concurrently is to create a separate Session object for each thread.

JMS support two styles of messaging:

- Point-to-point messaging
- Publish/subscribe messaging

These styles of messaging are also referred to as *messaging domains*, and you can combine both styles of messaging in an application. In the point-to-point domain, a destination is a queue and, in the publish/subscribe domain, a destination is a topic.

With versions of JMS before JMS 1.1, programming for the point-to-point domain uses one set of interfaces and methods, and programming for the publish/subscribe domain uses another set. The two sets are similar, but separate. With JMS 1.1, you can use a common set of interfaces and methods that support both messaging domains. The common interfaces provide a domain independent view of each messaging domain. [Table 103 on page 771](#) lists the JMS domain independent interfaces and their corresponding domain specific interfaces.

<i>Table 103. The JMS domain independent and domain specific interfaces</i>		
Domain independent interfaces	Domain specific interfaces for the point-to-point domain	Domain specific interfaces for the publish/subscribe domain
ConnectionFactory	QueueConnectionFactory	TopicConnectionFactory
Connection	QueueConnection	TopicConnection
Destination	Queue	Topic
Session	QueueSession	TopicSession
MessageProducer	QueueSender	TopicPublisher
MessageConsumer	QueueReceiver QueueBrowser	TopicSubscriber

JMS 1.1 retains all the domain specific interfaces, and so existing applications can still use these interfaces. For new applications, however, consider using the domain independent interfaces.

In WebSphere MQ classes for JMS, JMS objects are related to WebSphere MQ concepts in the following ways:

- A Connection object has properties that are derived from the properties of the connection factory that was used to create the connection. These properties control how an application connects to a queue manager. Examples of these properties are the name of the queue manager and, for an application that connects to the queue manager in client mode, the host name or IP address of the system on which the queue manager is running.
- A Session object encapsulates a WebSphere MQ connection handle, which therefore defines the transactional scope of the session.
- A MessageProducer object and a MessageConsumer object each encapsulates a WebSphere MQ object handle.

When using WebSphere MQ classes for JMS, all the normal rules of WebSphere MQ apply. Note, in particular, that an application can send a message to a remote queue but it can receive a message only from a queue that is owned by the queue manager to which the application is connected.

The JMS specification expects ConnectionFactory and Destination objects to be administered objects. An administrator creates and maintains administered objects in a central repository, and a JMS application retrieves these objects using the Java Naming and Directory Interface (JNDI).

In WebSphere MQ classes for JMS, the implementation of the Destination interface is an abstract superclass of Queue and Topic, and so an instance of Destination is either a Queue object or a Topic object. The domain independent interfaces treat a queue or a topic as a destination. The messaging domain for a MessageProducer or MessageConsumer object is determined by whether the destination is a queue or a topic.

In WebSphere MQ classes for JMS therefore, objects of the following types can be administered objects:

- ConnectionFactory
- QueueConnectionFactory
- TopicConnectionFactory
- Queue
- Topic
- XAConnectionFactory
- XAQueueConnectionFactory
- XATopicConnectionFactory

JMS messages

JMS messages are composed of a header, properties, and a body. JMS defines five types of message body.

JMS messages are composed of the following parts:

Header

All messages support the same set of header fields. Header fields contain values that are used by both clients and providers to identify and route messages.

Properties

Each message contains a built-in facility to support application-defined property values. Properties provide an efficient mechanism to filter application-defined messages.

Body

JMS defines several types of message body that cover the majority of messaging styles currently in use.

JMS defines five types of message body:

Stream

A stream of Java primitive values. It is filled and read sequentially.

Map

A set of name-value pairs, where names are strings and values are Java primitive types. The entries can be accessed sequentially or randomly by name. The order of the entries is undefined.

Text

A message containing a java.lang.String.

Object

A message that contains a serializable Java object

Bytes

A stream of uninterpreted bytes. This message type is for literally encoding a body to match an existing message format.

The JMSCorrelationID header field is used to link one message with another. It typically links a reply message with its requesting message. JMSCorrelationID can hold a provider-specific message ID, an application-specific String, or a provider-native byte[] value.

Message selectors in JMS

Messages can contain application-defined property values. An application can use message selectors to have a JMS provider filter messages.

A message contains a built-in facility to support application-defined property values. In effect, this provides a mechanism to add application-specific header fields to a message. Properties allow an application, using message selectors, to have a JMS provider select or filter messages on its behalf, using application-specific criteria. Application-defined properties must obey the following rules:

- Property names must obey the rules for a message selector identifier.
- Property values can be Boolean, byte, short, int, long, float, double, and String.
- The JMSX and JMS_ name prefixes are reserved.

Property values are set before sending a message. When a client receives a message, the message properties are read-only. If a client attempts to set properties at this point, a `MessageNotWriteableException` is thrown. If `clearProperties` is called, the properties can now be both read from, and written to.

A property value might duplicate a value in a message body. JMS does not define a policy for what might be made into a property. However, application developers must be aware that JMS providers probably handle data in a message body more efficiently than data in message properties. For best performance, applications must use message properties only when they need to customize a message header. The primary reason for doing this is to support customized message selection.

A JMS message selector allows a client to specify the messages that it is interested in by using the message header. Only messages with headers that match the selector are delivered.

Message selectors cannot refer to message body values.

A message selector matches a message when the selector evaluates to true when the message header field and property values are substituted for their corresponding identifiers in the selector.

A message selector is a String, with syntax that is based on a subset of the SQL92 conditional expression syntax. The order in which a message selector is evaluated is from left to right within a precedence level. You can use parentheses to change this order. Predefined selector literals and operator names are written here in uppercase; however, they are not case-sensitive.

A selector can contain:

- Literals
 - A string literal is enclosed in quotation marks. A doubled quotation mark represents a quotation mark. Examples are 'literal' and 'literal's'. Like Java string literals, these use the Unicode character encoding.
 - An exact numeric literal is a numeric value without a decimal point, such as 57, -957, and +62. Numbers in the range of Java long are supported.
 - An approximate numeric literal is a numeric value in scientific notation, such as 7E3 or -57.9E2, or a numeric value with a decimal, such as 7., -95.7, or +6.2. Numbers in the range of Java double are supported.
 - The Boolean literals TRUE and FALSE.
- Identifiers:
 - An identifier is an unlimited length sequence of Java letters and Java digits, the first of which must be a Java letter. A letter is any character for which the method `Character.isJavaLetter` returns true. This includes `_` and `$`. A letter or digit is any character for which the method `Character.isJavaLetterOrDigit` returns true.
 - Identifiers cannot be the names NULL, TRUE, or FALSE.
 - Identifiers cannot be NOT, AND, OR, BETWEEN, LIKE, IN, or IS.
 - Identifiers are either header field references or property references.
 - Identifiers are case sensitive.

- Message header field references are restricted to:
 - JMSDeliveryMode
 - JMSPriority
 - JMSMessageID
 - JMSTimestamp
 - JMSCorrelationID
 - JMSType

JMSMessageID, JMSTimestamp, JMSCorrelationID, and JMSType values can be null, and if so, are treated as a NULL value.
- Any name beginning with JMSX is a JMS-defined property name.
- Any name beginning with JMS_ is a provider-specific property name.
- Any name that does not begin with JMS is an application-specific property name. If there is a reference to a property that does not exist in a message, its value is NULL. If it does exist, its value is the corresponding property value.
- White space is the same as it is defined for Java: space, horizontal tab, form feed, and line terminator.
- Expressions:
 - A selector is a conditional expression. A selector that evaluates to true matches; a selector that evaluates to false or unknown does not match.
 - Arithmetic expressions are composed of themselves, arithmetic operations, identifiers (with a value that is treated as a numeric literal), and numeric literals.
 - Conditional expressions are composed of themselves, comparison operations, and logical operations.
- Standard bracketing (), to set the order in which expressions are evaluated, is supported.
- Logical operators in precedence order: NOT, AND, OR.
- Comparison operators: =, >, >=, <, <=, <> (not equal).
 - Only values of the same type can be compared. One exception is that it is valid to compare exact numeric values and approximate numeric values. (The type conversion required is defined by the rules of Java numeric promotion.) If there is an attempt to compare different types, the selector is always false.
 - String and Boolean comparison is restricted to = and <>. Two strings are equal only if they contain the same sequence of characters.
- Arithmetic operators in precedence order:
 - +, - unary.
 - *, /, multiplication, and division.
 - +, -, addition, and subtraction.
 - Arithmetic operations on a NULL value are not supported. If they are attempted, the complete selector is always false.
 - Arithmetic operations must use Java numeric promotion.
- arithmetic-expr1 [NOT] BETWEEN arithmetic-expr2 and arithmetic-expr3 comparison operator:
 - Age BETWEEN 15 and 19 is equivalent to age >= 15 AND age <= 19.
 - Age NOT BETWEEN 15 and 19 is equivalent to age < 15 OR age > 19.
 - If any of the expressions of a BETWEEN operation are NULL, the value of the operation is false. If any of the expressions of a NOT BETWEEN operation are NULL, the value of the operation is true.
- identifier [NOT] IN (string-literal1, string-literal2,...) comparison operator where identifier has a String or NULL value.
 - Country IN ('UK', 'US', 'France') is true for 'UK' and false for 'Peru'. It is equivalent to the expression (Country = 'UK') OR (Country = 'US') OR (Country = 'France').

- Country NOT IN ('UK', 'US', 'France') is false for 'UK' and true for 'Peru'. It is equivalent to the expression NOT ((Country = 'UK') OR (Country = 'US') OR (Country = 'France')).
- If the identifier of an IN or NOT IN operation is NULL, the value of the operation is unknown.
- identifier [NOT] LIKE pattern-value [ESCAPE escape-character] comparison operator, where identifier has a string value. pattern-value is a string literal, where _ stands for any single character and % stands for any sequence of characters (including the empty sequence). All other characters stand for themselves. The optional escape-character is a single character string literal, with a character that is used to escape the special meaning of the _ and % in pattern-value.
 - phone LIKE '12%3' is true for 123 and 12993 and false for 1234.
 - word LIKE 'l_se' is true for "lose" and false for "loose".
 - underscored LIKE '_%' ESCAPE '\' is true for "_foo" and false for "bar".
 - phone NOT LIKE '12%3' is false for 123 and 12993 and true for 1234.
 - If the identifier of a LIKE or NOT LIKE operation is NULL, the value of the operation is unknown.
- identifier IS NULL comparison operator tests for a null header field value, or a missing property value.
 - prop_name IS NULL.
- identifier IS NOT NULL comparison operator tests for the existence of a non-null header field value or a property value.
 - prop_name IS NOT NULL.

The following message selector selects messages with a message type of car, color of blue, and weight greater than 2500 lbs:

```
"JMSType = 'car' AND color = 'blue' AND weight > 2500"
```

As noted in the preceding list, property values can be NULL. The evaluation of selector expressions that contain NULL values is defined by SQL 92 NULL semantics. The following list gives a brief description of these semantics:

- SQL treats a NULL value as unknown.
- Comparison or arithmetic with an unknown value always yields an unknown value.
- The IS NULL operator converts an unknown value into a TRUE value.
- The IS NOT NULL operator converts an unknown value into a FALSE value.

Although SQL supports fixed decimal comparison and arithmetic, JMS message selectors do not. This is why exact numeric literals are restricted to those without a decimal. It is also why there are numerics with a decimal as an alternative representation for an approximate numeric value.

SQL comments are not supported.

Mapping JMS messages onto WebSphere MQ messages

WebSphere MQ messages are composed of a Message Descriptor, an optional MQRFH2 header, and a body. The contents of a JMS message are partly mapped and partly copied to a WebSphere MQ message.

This topic describes how the JMS message structure that is described in the first part of this section is mapped onto a WebSphere MQ message. It is of interest to programmers who want to transmit messages between JMS and traditional WebSphere MQ applications. It is also of interest to people who want to manipulate messages transmitted between two JMS applications, for example, in a WebSphere Message Broker implementation.

This section does not apply if an application uses a real-time connection to a broker. When an application uses a real-time connection, all communication is performed directly over TCP/IP; no WebSphere MQ queues or messages are involved.

WebSphere MQ messages are composed of three components:

- The WebSphere MQ Message Descriptor (MQMD)

- A WebSphere MQ MQRFH2 header
- The message body.

The MQRFH2 is optional, and its inclusion in an outgoing message is governed by a flag in the JMS Destination class. You can set this flag using the WebSphere MQ JMS administration tool. Because the MQRFH2 carries JMS-specific information, always include it in the message when the sender knows that the receiving destination is a JMS application. Normally, omit the MQRFH2 when sending a message directly to a non-JMS application. This is because such an application does not expect an MQRFH2 in its WebSphere MQ message.

If an incoming message does not have an MQRFH2 header, the Queue or Topic object derived from the JMSReplyTo header field of the message, by default, has this flag set so that a reply message sent to the queue or topic also does not have an MQRFH2 header. You can switch off this behavior of including an MQRFH2 header in a reply message only if the original message has an MQRFH2 header, by setting the TARGCLIENTMATCHING property of the connection factory to NO.

Figure 128 on page 776 shows how the structure of a JMS message is transformed to a WebSphere MQ message and back again:




Figure 128. How messages are transformed between JMS and WebSphere MQ using the MQRFH2 header

The structures are transformed in two ways:

Mapping

Where the MQMD includes a field that is equivalent to the JMS field, the JMS field is mapped onto the MQMD field. Additional MQMD fields are exposed as JMS properties, because a JMS application might need to get or set these fields when communicating with a non-JMS application.

Copying

Where there is no MQMD equivalent, a JMS header field or property is passed, possibly transformed, as a field inside the MQRFH2.

The MQRFH2 header and JMS

This collection of topics describes the MQRFH Version 2 header, which carries JMS-specific data that is associated with the message content. The MQRFH2 Version 2 is an extensible header, and can also carry additional information that is not directly associated with JMS. However, this section covers only its use by JMS. For a full description see [MQRFH2 - Rules and formatting header 2](#).

There are two parts of the header, a fixed portion and a variable portion.

Fixed portion

The fixed portion is modeled on the *standard* WebSphere MQ header pattern and consists of the following fields:

StrucId (MQCHAR4)

Structure identifier.

Must be MQRFH_STRUC_ID (value: "RFH ") (initial value).

MQRFH_STRUC_ID_ARRAY (value: "R","F","H"," ") is also defined.

Version (MQLONG)

Structure version number.

Must be MQRFH_VERSION_2 (value: 2) (initial value).

StrucLength (MQLONG)

Total length of MQRFH2, including the NameValueData fields.

The value set into StrucLength must be a multiple of 4 (the data in the NameValueData fields can be padded with space characters to achieve this).

Encoding (MQLONG)

Data encoding.

Encoding of any numeric data in the portion of the message following the MQRFH2 (the next header, or the message data following this header).

CodedCharSetId (MQLONG)

Coded character set identifier.

Representation of any character data in the portion of the message following the MQRFH2 (the next header, or the message data following this header).

Format (MQCHAR8)

Format name.

Format name for the portion of the message following the MQRFH2.

Flags (MQLONG)

Flags.

MQRFH_NO_FLAGS = 0. No flags set.

NameValueCCSID (MQLONG)

The coded character set identifier (CCSID) for the NameValueData character strings contained in this header. The NameValueData can be coded in a character set that differs from the other character strings that are contained in the header (StrucID and Format).

If the NameValueCCSID is a 2 byte Unicode CCSID (1200, 13488, or 17584), the byte order of the Unicode is the same as the byte ordering of the numeric fields in the MQRFH2. (For example, Version, StrucLength, and NameValueCCSID itself.)

<i>Table 104. Possible values for NameValueCCSID field</i>	
Value	Meaning
1200	UCS2 open-ended
1208	UTF8
13488	UCS2 2.0 subset
17584	UCS2 2.1 subset (includes Euro symbol)

Variable portion

The variable portion follows the fixed portion. The variable portion contains a variable number of MQRFH2 folders. Each folder contains a variable number of elements or properties. Folders group related properties. The MQRFH2 headers created by JMS can contain any of the following folders:

The <mcd> folder

mcd contains properties that describe the format of the message. For example, the message service domain Msd property identifies a JMS message as being JMSTextMessage, JMSBytesMessage, JMSStreamMessage, JMSMapMessage, JMSObjectMessage, or null.

The mcd folder is always present in a JMS message containing an MQRFH2.

It is always present in a message containing an MQRFH2 sent from WebSphere Message Broker. It describes the domain, format, type, and message set of a message.

Table 105. mcd property name, synonym, data type, and folder			
Property synonym	Property name	Data type	Folder
	mcd.Msd	string	<mcd><Msd>messageDomain</Msd></mcd>
	mcd.Set	string	<mcd><Set>messageDomain</Set></mcd>
	mcd.Type	string	<mcd><Type>messageDomain</Type></mcd>
	mcd.Fmt	string	<mcd><Fmt>messageDomain</Fmt></mcd>

Do not add your own properties in the mcd folder.

The <jms> folder

jms contains JMS header fields, and JMSX properties that cannot be fully expressed in the MQMD. The jms folder is always present in a JMS MQRFH2.

The <usr> folder

usr contains application-defined JMS properties associated with the message. The usr folder is present only if an application has set an application-defined property.

The <mqext> folder

mqext contains properties that are used only by WebSphere Application Server. The folder is present only if the application has set at least one of the IBM defined properties.

Table 106. mqext property name, synonym, data type, and folder			
Property synonym	Property name	Data type	Folder
JMSArmCorrelator	mqext.Arm	string	<mqext><Arm>armCorrelator</Arm></mqext>
JMSRMCorrelator	mqext.Wrm	string	<mqext><Wrm>wrmCorrelator</Wrm></mqext>

Do not add your own properties in the mqext folder.

The <mmps> folder

mmps contains properties that are used only by IBM WebSphere MQ publish/subscribe. The folder is present only if the application has set at least one of the integrated publish/subscribe properties.

Table 107. mmps property name, synonym, data type, and folder			
Property synonym	Property name	Data type	Folder
MQTopicString	mmps.Top	string	<mmps><Top>topicString</Top></mmps>
MQSubUserData	mmps.Sud	string	<mmps><Sud>subscriberUserData</Sud></mmps>
MQIsRetained	mmps.Ret	boolean	<mmps><Ret>isRetained</Ret></mmps>

Table 107. mqs property name, synonym, data type, and folder (continued)			
Property synonym	Property name	Data type	Folder
MQPubOptions	mqs.Pub	i8	<mqs><Pub>publicationOptions</Pub></mqs>
MQPubLevel	mqs.Pbl	i8	<mqs><Pbl>publicationLevel</Pbl></mqs>
MQPubTime	mqs.Pts	string	<mqs><Pts>publicationTime</Pts></mqs>
MQPubSeqNum	mqs.Seq	i8	<mqs><Seq>publicationSequenceNumber</Seq></mqs>
MQPubStrIntData	mqs.Sid	string	<mqs><Sid>publicationData</Sid></mqs>
MQPubFormat	mqs.Pfmt	i8	<mqs><Pfmt>messageFormat</Pfmt></mqs>

Do not add your own properties in the mqs folder.

Table 108 on page 779 shows a full list of property names.

Table 108. MQRFH2 folders and properties used by JMS				
JMS field name	Java type	MQRFH2 folder name	Property name	Type/values
JMSDestination	Destination	jms	Dst	string
JMSExpiration	long	jms	Exp	i8
JMSPriority	int	jms	Pri	i4
JMSDeliveryMode	int	jms	Dlv	i4
JMSCorrelationID	String	jms	Cid	string
JMSReplyTo	Destination	jms	Rto	string
JMSTimestamp	long	jms	Tms	i8
JMSType	String	mcd	Type, Set, Fmt	string
JMSXGroupID	String	jms	Gid	string
JMSXGroupSeq	int	jms	Seq	i4
xxx (user defined)	Any	usr	xxx	any
		mcd	Msd	jms_none jms_text jms_bytes jms_map jms_stream jms_object

NameValueLength (MQLONG)

Length in bytes of the NameValueData string that immediately follows this length field (it does not include its own length).

NameValueData (MQCHARn)

A single character string, whose length in bytes is given by the preceding NameValueLength field. It contains a folder holding a sequence of properties. Each property is a name/type/value triplet, contained within an XML element whose name is the folder name, as follows:

```
<foldername>
triplet1 triplet2 ..... tripletn </foldername>
```

The closing `</foldername>` tag can be followed by spaces as padding characters. Each triplet is encoded using an XML-like syntax:

```
<name dt='datatype'>value</name>
```

The `dt='datatype'` element is optional and is omitted for many properties, because the data type is predefined. If it is included, one or more space characters must be included before the `dt=` tag.

name

is the name of the property; see [Table 108 on page 779](#).

datatype

must match, after folding, one of the data types listed in [Table 109 on page 780](#).

value

is a string representation of the value to be conveyed, using the definitions in [Table 109 on page 780](#).

A null value is encoded using the following syntax:

```
<name dt='datatype' xsi:nil='true'></name>
```

Do not use `xsi:nil='false'`.

Table 109. Property data types	
Data type	Definition
string	Any sequence of characters excluding < and &
boolean	The character 0 or 1 (0 = false, 1 = true)
bin.hex	Hexadecimal digits representing octets
i1	A number, expressed using digits 0 . . 9, with optional sign (no fractions or exponent). Must lie in the range -128 to 127 inclusive
i2	A number, expressed using digits 0 . . 9, with optional sign (no fractions or exponent). Must lie in the range -32768 to 32767 inclusive
i4	A number, expressed using digits 0 . . 9, with optional sign (no fractions or exponent). Must lie in the range -2147483648 to 2147483647 inclusive
i8	A number, expressed using digits 0 . . 9, with optional sign (no fractions or exponent). Must lie in the range -9223372036854775808 to 92233720368547750807 inclusive
int	A number, expressed using digits 0 . . 9, with optional sign (no fractions or exponent). Must lie in the same range as i8. This can be used in place of one of the i* types if the sender does not want to associate a particular precision with the property
r4	Floating point number, magnitude <= 3.40282347E+38, >= 1.175E-37 expressed using digits 0 . . 9, optional sign, optional fractional digits, optional exponent
r8	Floating point number, magnitude <= 1.7976931348623E+308, >= 2.225E-307 expressed using digits 0 . . 9, optional sign, optional fractional digits, optional exponent

A string value can contain spaces. You must use the following escape sequences in a string value:

- & for the & character
- < for the < character

You can use the following escape sequences, but they are not required:

- > for the > character
- ' for the ' character
- " for the " character

JMS fields and properties with corresponding MQMD fields

These tables show the MQMD fields equivalent to JMS header fields, JMS properties, and JMS provider-specific properties.

Table 110 on page 781 lists the JMS header fields and Table 111 on page 781 lists the JMS properties that are mapped directly to MQMD fields. Table 112 on page 781 lists the provider-specific properties and the MQMD fields that they are mapped to.

Table 110. JMS header fields mapping to MQMD fields			
JMS header field	Java type	MQMD field	C type
JMSDeliveryMode	int	Persistence	MQLONG
JMSExpiration	long	Expiry	MQLONG
JMSPriority	int	Priority	MQLONG
JMSMessageID	String	MsgID	MQBYTE24
JMSTimestamp	long	PutDate PutTime	MQCHAR8 MQCHAR8
JMSCorrelationID	String	CorrelId	MQBYTE24

Table 111. JMS properties mapping to MQMD fields			
JMS property	Java type	MQMD field	C type
JMSXUserID	String	UserIdentifier	MQCHAR12
JMSXAppID	String	PutApplName	MQCHAR28
JMSXDeliveryCount	int	BackoutCount	MQLONG
JMSXGroupID	String	GroupId	MQBYTE24
JMSXGroupSeq	int	MsgSeqNumber	MQLONG

Table 112. JMS provider-specific properties mapping to MQMD fields			
JMS provider-specific property	Java type	MQMD field	C type
JMS_IBM_Report_Exception	int	Report	MQLONG
JMS_IBM_Report_Expiration	int	Report	MQLONG
JMS_IBM_Report_COA	int	Report	MQLONG
JMS_IBM_Report_COD	int	Report	MQLONG

Table 112. JMS provider-specific properties mapping to MQMD fields (continued)

JMS provider-specific property	Java type	MQMD field	C type
JMS_IBM_Report_PAN	int	Report	MQLONG
JMS_IBM_Report_NAN	int	Report	MQLONG
JMS_IBM_Report_Pass_Msg_ID	int	Report	MQLONG
JMS_IBM_Report_Pass_Correl_ID	int	Report	MQLONG
JMS_IBM_Report_Discard_Msg	int	Report	MQLONG
JMS_IBM_MsgType	int	MsgType	MQLONG
JMS_IBM_Feedback	int	Feedback	MQLONG
JMS_IBM_Format	String	Format“1” on page 782	MQCHAR8
JMS_IBM_PutApplType	int	PutApplType	MQLONG
JMS_IBM_Encoding	int	Encoding	MQLONG
JMS_IBM_Character_Set	String	CodedCharacterSetId“2” on page 782	MQLONG
JMS_IBM_PutDate	String	PutDate	MQCHAR8
JMS_IBM_PutTime	String	PutTime	MQCHAR8
JMS_IBM_Last_Msg_In_Group	boolean	MsgFlags	MQLONG
Note: <ol style="list-style-type: none"> 1. JMS_IBM_Format represents the format of the message body. This can be defined by the application setting the JMS_IBM_Format property of the message (note that there is an 8 character limit), or can default to the WebSphere MQ format of the message body appropriate to the JMS message type. JMS_IBM_Format maps to the MQMD Format field only if the message contains no RFH or RFH2 sections. In a typical message, it maps to the Format field of the RFH2 immediately preceding the message body. 2. JMS_IBM_Character_Set property value is a String value that contains the Java character set equivalent for the numeric CodedCharacterSetId value. MQMD field CodedCharacterSetId is a numeric value that contains the equivalent of the Java character set string specified by the JMS_IBM_Character_Set property. 			

Mapping JMS fields onto WebSphere MQ fields (outgoing messages)

These tables show how JMS header and property fields are mapped into MQMD and MQRFH2 fields at send() or publish() time.

Table 113 on page 783 shows how the JMS header fields are mapped into MQMD/RFH2 fields at send() or publish() time. Table 114 on page 783 shows how JMS properties are mapped into MQMD/RFH2 fields at send() or publish() time. Table 115 on page 783 shows how JMS provider-specific properties are mapped to MQMD fields at send() or publish() time,

For fields marked Set by Message Object, the value transmitted is the value held in the JMS message immediately before the send() or publish() operation. The value in the JMS message is left unchanged by the operation.

For fields marked Set by Send Method, a value is assigned when the send() or publish() is performed (any value held in the JMS message is ignored). The value in the JMS message is updated to show the value used.

Fields marked as Receive-only are not transmitted and are left unchanged in the message by send() or publish().

<i>Table 113. Outgoing message field mapping</i>			
JMS header field name	MQMD field used for transmission	Header	Set by
JMSDestination		MQRFH2	Send Method
JMSDeliveryMode	Persistence	MQRFH2	Send Method
JMSExpiration	Expiry	MQRFH2	Send Method
JMSPriority	Priority	MQRFH2	Send Method
JMSMessageID	MsgID		Send Method
JMSTimestamp	PutDate/PutTime		Send Method
JMSCorrelationID	CorrelId	MQRFH2	Message Object
JMSReplyTo	ReplyToQ/ReplyToQMgr	MQRFH2	Message Object
JMSType		MQRFH2	Message Object
JMSRedelivered			Receive-only
Note: 1. MQMD field CodedCharacterSetId is a numeric value that contains the equivalent of the Java character set string specified by the JMS_IBM_Character_Set property.			

<i>Table 114. Outgoing message JMS property mapping</i>			
JMS property name	MQMD field used for transmission	Header	Set by
JMSXUserID	UserIdentifier		Send Method
JMSXAppID	PutApplName		Send Method
JMSXDeliveryCount			Receive-only
JMSXGroupID	GroupId	MQRFH2	Message Object
JMSXGroupSeq	MsgSeqNumber	MQRFH2	Message Object

<i>Table 115. Outgoing message JMS provider-specific property mapping</i>			
JMS provider-specific property name	MQMD field used for transmission	Header	Set by
JMS_IBM_Report_Exception	Report		Message Object
JMS_IBM_Report_Expiration	Report		Message Object
JMS_IBM_Report_COA/COD	Report		Message Object
JMS_IBM_Report_NAN/PAN	Report		Message Object
JMS_IBM_Report_Pass_Msg_ID	Report		Message Object
JMS_IBM_Report_Pass_Correl_ID	Report		Message Object
JMS_IBM_Report_Discard_Msg	Report		Message Object
JMS_IBM_MsgType	MsgType		Message Object

Table 115. Outgoing message JMS provider-specific property mapping (continued)

JMS provider-specific property name	MQMD field used for transmission	Header	Set by
JMS_IBM_Feedback	Feedback		Message Object
JMS_IBM_Format	Format		Message Object
JMS_IBM_PutApplType	PutApplType		Send Method
JMS_IBM_Encoding	Encoding		Message Object
JMS_IBM_Character_Set	CodedCharacterSetId		Message Object
JMS_IBM_PutDate	PutDate		Send Method
JMS_IBM_PutTime	PutTime		Send Method
JMS_IBM_Last_Msg_In_Group	MsgFlags		Message Object

Mapping JMS header fields at send() or publish()

These notes relate to the mapping of JMS fields at send() or publish().

JMSDestination to MQRFH2

This is stored as a string that serializes the salient characteristics of the destination object, so that a receiving JMS can reconstitute an equivalent destination object. The MQRFH2 field is encoded as URI (see “Uniform resource identifiers (URIs)” on page 841 for details of the URI notation).

JMSReplyTo to MQMD.ReplyToQ, ReplyToQMGr, MQRFH2

The queue name is copied to the MQMD.ReplyToQ field, and the queue manager name is copied to the ReplyToQMGr fields. The destination extension information (other useful details that are kept in the destination object) is copied into the MQRFH2 field. The MQRFH2 field is encoded as a URI (see “Uniform resource identifiers (URIs)” on page 841 for details of the URI notation).

JMSDeliveryMode to MQMD.Persistence

The JMSDeliveryMode value is set by the send() or publish() Method or MessageProducer, unless the Destination Object overrides it. The JMSDeliveryMode value is mapped to the MQMD.Persistence field as follows:

- JMS value PERSISTENT is equivalent to MQPER_PERSISTENT
- JMS value NON_PERSISTENT is equivalent to MQPER_NOT_PERSISTENT

If the MQQueue persistence property is not set to WMQConstants.WMQ_PER_QDEF, the delivery mode value is also encoded in the MQRFH2.

JMSExpiration to/from MQMD.Expiry, MQRFH2

JMSExpiration stores the time to expire (the sum of the current time and the time to live), whereas MQMD stores the time to live. Also, JMSExpiration is in milliseconds, but MQMD.Expiry is in tenths of a second.

- If the send() method sets an unlimited time to live, MQMD.Expiry is set to MQEI_UNLIMITED, and no JMSExpiration is encoded in the MQRFH2.
- If the send() method sets a time to live that is less than 214748364.7 seconds (about 7 years), the time to live is stored in MQMD.Expiry, and the expiration time (in milliseconds), is encoded as an i8 value in the MQRFH2.
- If the send() method sets a time to live greater than 214748364.7 seconds, MQMD.Expiry is set to MQEI_UNLIMITED. The true expiration time in milliseconds is encoded as an i8 value in the MQRFH2.

JMSPriority to MQMD.Priority

Directly map JMSPriority value (0-9) onto MQMD priority value (0-9). If JMSPriority is set to a non-default value, the priority level is also encoded in the MQRFH2.

JMSMessageID from MQMD.MessageID

All messages sent from JMS have unique message identifiers assigned by WebSphere MQ. The value assigned is returned in the MQMD.MessageID field after the MQPUT call, and is passed back to the application in the JMSMessageID field. The WebSphere MQ messageId is a 24-byte binary value, whereas the JMSMessageID is a string. The JMSMessageID is composed of the binary messageId value converted to a sequence of 48 hexadecimal characters, prefixed with the characters ID:. JMS provides a hint that can be set to disable the production of message identifiers. This hint is ignored, and a unique identifier is assigned in all cases. Any value that is set into the JMSMessageID field before a send() is overwritten.

If you do require the ability to specify the MQMD.MessageID, you can do this with one of the WebSphere MQ JMS extensions described in [“Reading and writing the message descriptor from a WebSphere MQ classes for JMS application”](#) on page 857.

JMSTimestamp to MQRFH2

During a send, the JMSTimestamp field is set according to the JVM's clock. This value is set into the MQRFH2. Any value that is set into the JMSTimestamp field before a send() is overwritten. See also the JMS_IBM_PutDate and JMS_IBM_PutTime properties.

JMSType to MQRFH2

This string is set into the MQRFH2 mcd.Type field. If it is in URI format, it can also affect mcd.Set and mcd.Fmt fields. See also [“Using a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker”](#) on page 883.

JMSCorrelationID to MQMD.CorrelId, MQRFH2

The JMSCorrelationID can hold one of the following:

A provider specific message ID

This is a message identifier from a message previously sent or received, and so should be a string of 48 lowercase hexadecimal digits that are prefixed with ID: The prefix is removed, the remaining characters are converted into binary, and then they are set into the MQMD.CorrelId field. No CorrelId value is encoded in the MQRFH2.

A provider-native byte[] value

The value is copied into the MQMD.CorrelId field - padded with nulls, or truncated to 24 bytes if necessary. No CorrelId value is encoded in the MQRFH2.

An application-specific string

The value is copied into the MQRFH2. The first 24 bytes of the string, in UTF8 format, are written into the MQMD.CorrelID.

Mapping JMS property fields

These notes refer to the mapping of JMS property fields in WebSphere MQ messages.

JMSXUserID from MQMD UserIdentifier

JMSXUserID is set on return from send call.

JMSXAppID from MQMD PutApplName

JSMXAppID is set on return from send call.

JMSXGroupID to MQRFH2 (point-to-point)

For point-to-point messages, the JMSXGroupID is copied into the MQMD GroupID field. If the JMSXGroupID starts with the prefix ID:, it is converted into binary. Otherwise, it is encoded as a UTF8 string. The value is padded or truncated if necessary to a length of 24 bytes. The MQMF_MSG_IN_GROUP flag is set.

JMSXGroupID to MQRFH2 (publish/subscribe)

For publish/subscribe messages, the JMSXGroupID is copied into the MQRFH2 as a string.

JMSXGroupSeq MQMD MsgSeqNumber (point-to-point)

For point-to-point messages, the JMSXGroupSeq is copied into the MQMD MsgSeqNumber field. The MQMF_MSG_IN_GROUP flag is set.

JMSXGroupSeq MQMD MsgSeqNumber (publish/subscribe)

For publish/subscribe messages, the JMSXGroupSeq is copied into the MQRFH2 as an i4.

Mapping JMS provider-specific fields

The following notes refer to the mapping of JMS provider-specific fields into IBM WebSphere MQ messages.

JMS_IBM_Report_<name> to MQMD Report

A JMS application can set the MQMD Report options, using the following JMS_IBM_Report_XXX properties. The single MQMD is mapped to several JMS_IBM_Report_XXX properties. The application must set the value of these properties to the standard IBM WebSphere MQ MQRO_constants (included in com.ibm.mq.MQC). So, for example, to request COD with full Data, the application must set JMS_IBM_Report_COD to the value CMQC.MQRO_COD_WITH_FULL_DATA.

JMS_IBM_Report_Exception

MQRO_EXCEPTION or
MQRO_EXCEPTION_WITH_DATA or
MQRO_EXCEPTION_WITH_FULL_DATA

JMS_IBM_Report_Expiration

MQRO_EXPIRATION or
MQRO_EXPIRATION_WITH_DATA or
MQRO_EXPIRATION_WITH_FULL_DATA

JMS_IBM_Report_COA

MQRO_COA or
MQRO_COA_WITH_DATA or
MQRO_COA_WITH_FULL_DATA

JMS_IBM_Report_COD

MQRO_COD or
MQRO_COD_WITH_DATA or
MQRO_COD_WITH_FULL_DATA

JMS_IBM_Report_PAN

MQRO_PAN

JMS_IBM_Report_NAN

MQRO_NAN

JMS_IBM_Report_Pass_Msg_ID

MQRO_PASS_MSG_ID

JMS_IBM_Report_Pass_Correl_ID

MQRO_PASS_CORREL_ID

JMS_IBM_Report_Discard_Msg

MQRO_DISCARD_MSG

JMS_IBM_MsgType to MQMD MessageType

Value maps directly onto MQMD MessageType. If the application has not set an explicit value of JMS_IBM_MsgType, a default value is used. This default value is determined as follows:

- If JMSReplyTo is set to an IBM WebSphere MQ queue destination, MSGType is set to the value MQMT_REQUEST
- If JMSReplyTo is not set, or is set to anything other than an IBM WebSphere MQ queue destination, MsgType is set to the value MQMT_DATAGRAM

JMS_IBM_Feedback to MQMD Feedback

Value maps directly onto MQMD Feedback.

JMS_IBM_Format to MQMD Format

Value maps directly onto MQMD Format.

JMS_IBM_Encoding to MQMD Encoding

If set, this property overrides the numeric encoding of the Destination Queue or Topic.

JMS_IBM_Character_Set to MQMD CodedCharacterSetId

If set, this property overrides the coded character set property of the Destination Queue or Topic.

JMS_IBM_PutDate from MQMD PutDate

The value of this property is set, during send, directly from the PutDate field in the MQMD. Any value that is set into the JMS_IBM_PutDate property before a send is overwritten. This field is a String of eight characters, in the IBM WebSphere MQ Date format of YYYYMMDD. This property can be used with the JMS_IBM_PutTime property to determine the time the message was put according to the queue manager.

JMS_IBM_PutTime from MQMD PutTime

The value of this property is set, during send, directly from the PutTime field in the MQMD. Any value that is set into the JMS_IBM_PutTime property before a send is overwritten. This field is a String of eight characters, in the IBM WebSphere MQ Time format of HHMMSSSTH. This property can be used with the JMS_IBM_PutDate property to determine the time the message was put according to the queue manager.

JMS_IBM_Last_Msg_In_Group to MQMD MsgFlags

For point-to-point messaging, this Boolean value maps to the MQMF_LAST_MSG_IN_GROUP flag in the MQMD MsgFlags field. It is normally used with the JMSXGroupID and JMSXGroupSeq properties to indicate to a legacy IBM WebSphere MQ application that this message is the last in a group. This property is ignored for publish/subscribe messaging.

Mapping WebSphere MQ fields onto JMS fields (incoming messages)

These tables show how JMS header and property fields are mapped into MQMD and MQRFH2 fields at get() or receive() time.

Table 116 on page 787 shows how JMS header fields are mapped onto MQMD/MQRFH2 fields at get() or receive() time. Table 117 on page 788 shows how JMS property fields are mapped onto MQMD/MQRFH2 fields at get() or receive() time. Table 118 on page 788 shows how JMS provider-specific properties are mapped.

Table 116. Incoming message JMS header field mapping		
JMS header field name	MQMD field retrieved from	MQRFH2 field retrieved from
JMSDestination		jms.Dst or mqps.Top ^{“1”} on page 788
JMSDeliveryMode	Persistence ^{“2”} on page 788	jms.Dlv ^{“2”} on page 788
JMSExpiration		jms.Exp
JMSPriority	Priority	
JMSMessageID	MsgID	
JMSTimestamp	PutDate ^{“2”} on page 788 PutTime ^{“2”} on page 788	jms.Tms ^{“2”} on page 788
JMSCorrelationID	CorrelId ^{“2”} on page 788	jms.Cid ^{“2”} on page 788
JMSReplyTo	ReplyToQ ^{“2”} on page 788 ReplyToQMGr ^{“2”} on page 788	jms.Rto ^{“2”} on page 788
JMSType		mcd.Type, mcd.Set, mcd.Fmt
JMSRedelivered	BackoutCount	

Table 116. Incoming message JMS header field mapping (continued)

JMS header field name	MQMD field retrieved from	MQRFH2 field retrieved from
Note: <ol style="list-style-type: none"> 1. If both jms.Dst and mqps.Top are set, the value in jms.Dst is used. 2. For properties that can have values retrieved from the MQRFH2 or the MQMD, if both are available, the setting in the MQRFH2 is used. 3. JMS_IBM_Character_Set property value is a String value that contains the Java character set equivalent for the numeric CodedCharacterSetId value. 		

Table 117. Incoming message property mapping

JMS property name	MQMD field retrieved from	MQRFH2 field retrieved from
JMSXUserID	UserIdentifier	
JMSXAppID	PutApplName	
JMSXDeliveryCount	BackoutCount	
JMSXGroupID	GroupId “1” on page 788	jms.Gid “1” on page 788
JMSXGroupSeq	MsgSeqNumber “1” on page 788	jms.Seq “1” on page 788
Note: <ol style="list-style-type: none"> 1. For properties that can have values retrieved from the MQRFH2 or the MQMD, if both are available, the setting in the MQRFH2 is used. The properties are set from the MQMD values only if the MQMF_MSG_IN_GROUP or MQMF_LAST_MSG_IN_GROUP message flags are set. 		

Table 118. Incoming message provider-specific JMS property mapping

JMS property name	MQMD field retrieved from	MQRFH2 field retrieved from
JMS_IBM_Report_Exception	Report	
JMS_IBM_Report_Expiration	Report	
JMS_IBM_Report_COA	Report	
JMS_IBM_Report_COD	Report	
JMS_IBM_Report_PAN	Report	
JMS_IBM_Report_NAN	Report	
JMS_IBM_Report_Pass_Msg_ID	Report	
JMS_IBM_Report_Pass_Correl_ID	Report	
JMS_IBM_Report_Discard_Msg	Report	
JMS_IBM_MsgType	MsgType	
JMS_IBM_Feedback	Feedback	
JMS_IBM_Format	Format	
JMS_IBM_PutApplType	PutApplType	

Table 118. Incoming message provider-specific JMS property mapping (continued)

JMS property name	MQMD field retrieved from	MQRFH2 field retrieved from
JMS_IBM_Encoding “1” on page 789	Encoding	
JMS_IBM_Character_Set “1” on page 789	CodedCharacterSetId	
JMS_IBM_PutDate	PutDate	
JMS_IBM_PutTime	PutTime	
JMS_IBM_Last_Msg_In_Group	MsgFlags	
1. Only set if the incoming message is a Bytes Message.		

Exchanging messages between a JMS application and a traditional WebSphere MQ application

This topic describes what happens when a JMS application exchanges messages with a traditional WebSphere MQ application that cannot process the MQRFH2 header.

. Figure 129 on page 789 shows the mapping.

The administrator indicates that the JMS application is communicating with a traditional WebSphere MQ application by setting the TARGCLIENT property of the destination to *MQ*. This indicates that no MQRFH2 header is to be produced. If this is not done, the receiving application must be able to handle the MQRFH2 header.

The mapping from JMS to MQMD targeted at a traditional WebSphere MQ application is the same as mapping from JMS to MQMD targeted at a JMS application. If WebSphere MQ classes for JMS receives a WebSphere MQ message with the MQMD *Format* field set to anything other than MQFMT_RFH2, data is being received from a non-JMS application. If the format is MQFMT_STRING, the message is received as a JMS text message. Otherwise, it is received as a JMS bytes message. Because there is no MQRFH2, only those JMS properties that are transmitted in the MQMD can be restored.

If WebSphere MQ classes for JMS receives a message that does not have an MQRFH2 header, the TARGCLIENT property of the Queue or Topic object derived from the JMSReplyTo header field of the message is set to MQ by default. This means that a reply message sent to the queue or topic also does not have an MQRFH2 header. You can switch off this behavior of including an MQRFH2 header in a reply message only if the original message has an MQRFH2 header, by setting the TARGCLIENTMATCHING property of the connection factory to NO.

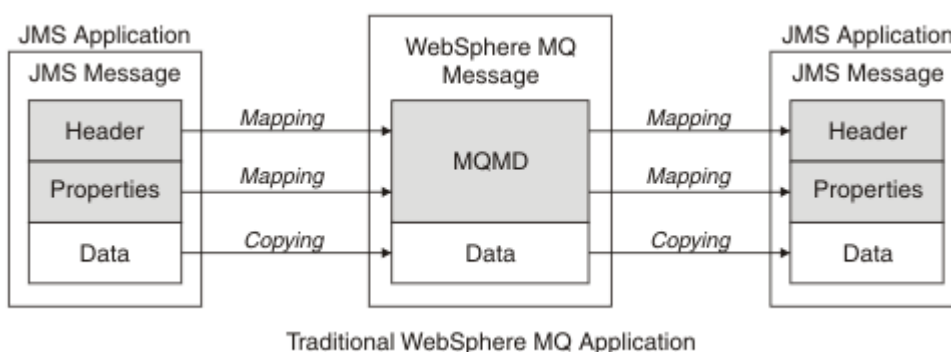


Figure 129. How JMS messages are transformed to WebSphere MQ messages with no MQRFH2 header

The JMS message body

This topic contains information about the encoding of the message body itself. The encoding depends on the type of JMS message.

ObjectMessage

An ObjectMessage is an object serialized by the Java Runtime in the normal way.

TextMessage

A TextMessage is an encoded string. For an outgoing message, the string is encoded in the character set given by the destination object. This defaults to UTF8 encoding (the UTF8 encoding starts with the first character of the message; there is no length field at the start). It is, however, possible to specify any other character set supported by WebSphere MQ classes for JMS. Such character sets are used mainly when you send a message to a non-JMS application.

If the character set is a double-byte set (including UTF16), the destination object's integer encoding specification determines the order of the bytes.

An incoming message is interpreted using the character set and encoding that are specified in the message itself. These specifications are in the last WebSphere MQ header (or MQMD if there are no headers). For JMS messages, the last header is usually the MQRFH2.

BytesMessage

A BytesMessage is, by default, a sequence of bytes as defined by the JMS 1.0.2 specification and associated Java documentation.

For an outgoing message that was assembled by the application itself, the destination object's encoding property can be used to override the encodings of integer and floating point fields contained in the message. For example, you can request that floating point values are stored in S/390® rather than IEEE format).

An incoming message is interpreted using the numeric encoding specified in the message itself. This specification is in the last WebSphere MQ header (or MQMD if there are no headers). For JMS messages, the last header is usually the MQRFH2.

If a BytesMessage is received, and is re-sent without modification, its body is transmitted byte for byte, as it was received. The destination object's encoding property has no effect on the body. The only string-like entity that can be sent explicitly in a BytesMessage is a UTF8 string. This is encoded in Java UTF8 format, and starts with a 2-byte length field. The destination object's character set property has no effect on the encoding of an outgoing BytesMessage. The character set value in an incoming WebSphere MQ message has no effect on the interpretation of that message as a JMS BytesMessage.

Non-Java applications are unlikely to recognize the Java UTF8 encoding. Therefore, for a JMS application to send a BytesMessage that contains text data, the application itself must convert its strings to byte arrays, and write these byte arrays into the BytesMessage.

MapMessage

A MapMessage is a string containing XML name/type/value triplets encoded as:

```
<map>
  <elt name="elementname1" dt="datatype1">value1</elt>
  <elt name="elementname2" dt="datatype2">value2</elt>
  ...
</map>
```

where datatype is one of the data types listed in [Table 109 on page 780](#). The default data type is string, and so the attribute dt="string" is omitted for string elements.

The character set used to encode or interpret the XML string that forms the body of a map message is determined according to the rules that apply to a text message.

Versions of WebSphere MQ classes for JMS earlier than Version 5.3 encoded the body of a map message in the following format:

```
<map>
  <elementname1 dt="datatype1">value1</elementname1>
  <elementname2 dt="datatype2">value2</elementname2>
  ...
</map>
```

Version 5.3 and later versions of WebSphere MQ classes for JMS can interpret either format, but versions of WebSphere MQ classes for JMS earlier than Version 5.3 cannot interpret the current format.

If an application needs to send map messages to another application that is using a version of WebSphere MQ classes for JMS earlier than Version 5.3, the sending application must call the connection factory method `setMapNameStyle(WMQConstants.WMQ_MAP_NAME_STYLE_COMPATIBLE)` to specify that the map messages are sent in the previous format. By default, all map messages are sent in the current format.

StreamMessage

A `StreamMessage` is like a map message, but without element names:

```
<stream>
  <elt dt="datatype1">value1</elt>
  <elt dt="datatype2">value2</elt>
  ...
</stream>
```

where `datatype` is one of the data types listed in [Table 109 on page 780](#). The default data type is `string`, and so the attribute `dt="string"` is omitted for string elements.

The character set used to encode or interpret the XML string that makes up the `StreamMessage` body is determined following the rules that apply to a `TextMessage`.

The `MQRFH2.format` field is set as follows:

MQFMT_NONE

for `ObjectMessage`, `BytesMessage`, or messages with no body.

MQFMT_STRING

for `TextMessage`, `StreamMessage`, or `MapMessage`.

JMS message conversion

Message data conversion in JMS is performed when sending and receiving messages. WebSphere MQ performs most data conversion automatically. It converts text and numeric data when transferring a message between JMS applications. Text is converted when exchanging a `JMSTextMessage` between a JMS application and a WebSphere MQ application.

If you are planning to do more complex message exchanges, the following topics are of interest you. Complex message exchanges include:

- Transferring non-text messages between a WebSphere MQ application and a JMS application.
- Exchanging text data in byte format.
- Converting the text in your application.

JMS message data

Data conversion is necessary to exchange text and numeric data between applications, even between two JMS applications. The internal representation of text and numbers must be encoded so they can be transferred in a message. Encoding forces a decision about how numbers and text are represented. WebSphere MQ manages the encoding of text and numbers in JMS messages, except for `JMSObjectMessage`, see [“JMSObjectMessage” on page 798](#). It uses three message attributes. The three attributes are `CodedCharacterSetId`, `Encoding`, and `Format`.

These three message attributes are normally stored in the JMS header, `MQRFH2`, fields of a JMS message. If the message type is an MQ, rather than JMS type of message, the attributes are stored in the message descriptor, `MQMD`. The attributes are used to convert the JMS message data. JMS message data is transferred in the message data part of a WebSphere MQ message.

JMS message properties

JMS message properties, such as `JMS_IBM_CHARACTER_SET`, are exchanged in the `MQRFH2` header part of a JMS message, unless the message has been sent without an `MQRFH2`. Only `JMSTextMessage` and `JMSBytesMessage` can be sent without an `MQRFH2`. If a JMS property is stored as a WebSphere MQ message property in the message descriptor, `MQMD`, it is converted as part of the `MQMD`

conversion. If a JMS property is stored in the MQRFH2, it is stored in the character set specified by MQRFH2.NameValueCCSID. When a message is sent or received, message properties are converted to and from their internal representation in the JVM. The conversion is to and from the character set of the message descriptor or MQRFH2.NameValueCCSID. Numeric data is converted to text.

JMS message conversion

The following topics contain examples and tasks that are useful if you plan to exchange more complex messages that require conversion.

JMS message conversion approaches

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by WebSphere MQ.

You can ask a number of questions about how to approach message conversion:

Is it necessary to think about message conversion at all?

In some cases, such as JMS to JMS message transfers, and exchanging text messages with IBM WebSphere MQ programs, IBM WebSphere MQ performs the necessary conversions for you, automatically. You might want to control data conversion for performance reasons, or you might be exchanging complex messages that have a predefined format. In cases such as these you must understand message conversion, and read the following topics.

What kinds of conversion are there?

There are four main types of conversion, which are explained in the following sections:

1. [“JMS client data conversion” on page 792](#)
2. [“Application data conversion” on page 793](#)
3. [“Queue manager data conversion” on page 793](#)
4. [“Message channel data conversion” on page 794](#)

Where should conversion be performed?

The section, [“Choosing an approach to message conversion: receiver makes good” on page 795](#), describes the usual approach of “receiver makes good”. “Receiver makes good” also applies to JMS data conversion.

JMS client data conversion

JMS client⁴ data conversion is the conversion of Java primitives and objects into bytes in a JMS message as it is sent to a destination, and conversion back again, when it is received. JMS client data conversion uses the methods of the JMSMessage classes. The methods are listed by JMSMessage class type in [Table 119 on page 796](#).

Conversion to and from the internal JVM representation of numbers and text is performed for read, get, set, and write methods. The conversion is performed when the message is sent, and when any of the read or get methods is called on a message that has been received.

The code page and numeric encoding used to write or set the contents of a message are defined as attributes of the destination. The destination code page and numeric encoding can be changed administratively. An application can also override the destination code page and encoding by setting the message properties that control writing or setting message content.

If you want to convert number encoding when a JMSBytesMessage message is sent to a destination that is not defined as Native encoding, you must set the message property JMS_IBM_ENCODING before sending the message. If you are following the “receiver makes good” pattern, or if you are exchanging

⁴ “JMS Client” refers to the WebSphere MQ classes for JMS that implement the JMS interface, which runs either in client or bindings mode.

messages between JMS applications, the application does not need to set `JMS_IBM_ENCODING`. In most cases you can leave the Encoding property as Native.

For `JMSStreamMessage`, `JMSMapMessage`, and `JMSTextMessage` messages, the character set identifier properties of the destination are used. Encoding is ignored on sending as numbers are written out in text format. The JMS client application program does not have to set `JMS_IBM_CHARACTER_SET` before sending the message if the destination character set property to apply..

To get the data in a message an application calls the JMS message read or get methods. The methods refer to the code page and encoding defined in the previous message header to create the Java primitives and objects correctly.

JMS client data conversion meets the needs of most JMS applications that are exchanging messages between one JMS client and another. You do not code any explicit data conversion. You do not use the `java.nio.charset.Charset` class, which is typically used when writing text to a file. The `writeString` and `setString` methods do the conversion for you.

For more details on JMS client data conversion, see [“JMS client message conversion and encoding” on page 804](#).

Application data conversion

A JMS client application can perform explicit character data conversion by using the `java.nio.charset.Charset` class; see the examples in [Figure 132 on page 797](#) and [Figure 133 on page 797](#). String data is converted into bytes, using the `getBytes` method, and sent as bytes. The bytes are converted back into text by using a `String` constructor that takes a byte array and a `Charset`. Character data is converted using the `encode` and `decode` `Charset` methods. Typically the message is sent or received as `JMSBytesMessage`, because the message part of a `JMSBytesMessage` does not contain anything other than the data written by the application⁵. You can also send and receive bytes using `JMSStreamMessage`, `JMSMapMessage`, or `JMSObjectMessage`.

There are no Java methods to encode and decode bytes that contain numeric data represented in different encoding formats. Numeric data is encoded and decoded automatically using the numeric JMSMessage read and write methods. The read and write methods use the value of the `JMS_IBM_ENCODING` attribute of the message data.

A typical use for application data conversion is if a JMS client sends or receives a formatted message from a non-JMS application. A formatted message contains text, numeric, and bytes data organized by the length of the data fields. Unless the non-JMS application has specified the message format as "MQSTR ", the message is constructed as a `JMSBytesMessage`. To receive formatted message data in a `JMSBytesMessage` you must call a sequence of methods. The methods must be called in the same order the fields were written into the message. If the fields are numeric, you must know the encoding and length of the numeric data. If any of the fields contain byte or text data, you must know the length of any byte data in the message. There are two ways to convert a formatted message into a Java object that is easy to use.

1. Construct a Java class corresponding to the record, to encapsulate reading and writing the message. Access to the data in the record is with get and set methods of the class.
2. Construct a Java class corresponding to the record by extending the `com.ibm.mq.headers` class. Access to the data in the class is with type-specific accessors of the form, `getStringValue(fieldName)`;

See [“Exchanging a formatted record with a non-JMS application” on page 811](#).

Queue manager data conversion

In WebSphere MQ V7.0, code page conversion can be performed by the queue manager when a JMS client program gets a message. The conversion is the same as the conversion performed for a C program. A C program sets `MQGMO_CONVERT` as an `MQGET GetMsgOpts` parameter option; see [Figure 131 on page](#)

⁵ One exception: Data written using `writeUTF` starts with a 2 byte length field

797. A queue manager performs conversion for a JMS client program that is receiving a message, if the WMQ_RECEIVE_CONVERSION destination property is set to WMQ_RECEIVE_CONVERSION_QMGR. The JMS client program can also set the destination property; see [Figure 130 on page 794](#).

Before V7.0, conversions were always performed by the JMS client. JMS client data conversion is restricted to converting sequences of numbers and text of type and length known to the JMS client. It cannot convert data structures; see [“Exchanging a formatted record with a non-JMS application” on page 811](#). In V7.0, until fix pack 7.0.1.5, if the conversion can be performed by the queue manager, it is always performed by the queue manager. From 7.0.1.5 onwards, the default conversion behavior reverts to the same as V6.0, and all conversions are performed by the JMS client. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, you can set a new destination option, WMQ_RECEIVE_CONVERSION, to control where conversion is performed, and WMQ_RECEIVE_CCSD, to set the target code page; see [Figure 130 on page 794](#).

```
((MQDestination)destination).setIntProperty(  
    WMQConstants.WMQ_RECEIVE_CONVERSION,  
    WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Or,

```
((MQDestination)destination).setReceiveConversion  
    (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Figure 130. Enable queue manager data conversion

The main benefit of queue manager conversion comes when exchanging messages with non-JMS applications. If the Format field in the message is defined, and the target character set, or encoding, is different to the message, the queue manager performs data conversion for the target application, if the application requests it. The queue manager converts message data formatted according to one of the predefined WebSphere MQ message types, such as a CICS bridge header (MQCIH). If the Format field is user-defined, the queue manager looks for a data conversion exit with the name provided in the Format field.

Queue manager data conversion is used to best effect with the "receiver makes good" design pattern. A sending JMS client does not need to perform conversion. A non-JMS receiving program relies on the conversion exit to ensure that the message is delivered in the required code page and encoding. With a sending JMS client, and non-JMS receiver, the example applies to IBM WebSphere MQ pre- and post-V7.0. With IBM WebSphere MQ V7.0, the conversion exit can be called for a receiving JMS program as well.

You can create a data conversion exit, using the data conversion exit utility, **crtmqcvx**, to enable the queue manager to convert your own record formatted data. You can build your own record format, use the `com.ibm.mq.headers` to access it as a Java class, and use your own conversion exit to convert it. On z/OS the utility is called **CSQUCVX**, and on IBM i, **CVTMQMDTA**. See [“Exchanging a formatted record with a non-JMS application” on page 811](#).

Message channel data conversion

WebSphere MQ Sender, Server, Cluster-receiver, and Cluster-sender channels have a message conversion option, CONVERT. The contents of a message can optionally be converted when a message is sent. The conversion takes place at the sending end of the channel. The cluster-receiver definition is used to auto-define the corresponding cluster-sender channel.

Data conversion by message channels is typically used if it is not possible to use other forms of conversion.

Choosing an approach to message conversion: "receiver makes good"

The usual approach in WebSphere MQ application design for code conversion is "receiver makes good". "Receiver makes good" reduces the number of message conversions. It also avoids the problem of unexpected channel errors if message conversion fails on some intermediary queue manager during message transfer. The "receiver makes good" rule is only broken if there is some reason why the receiver cannot make good. The receiving platform might not have the right character set, for example.

"Receiver makes good" is also good general guidance for JMS client applications. But in specific cases, conversion to the correct character set at source can be more efficient. Conversion from the JVM internal representation must take place when a message containing text or numeric types is sent. Conversion to the character set required by the receiver, if the receiver is not a JMS client, might remove the need for the non-JMS recipient to perform conversion. If the recipient is a JMS client, it is going to convert again, anyway, to decode the message data and create Java primitives and objects.

The difference between JMS client applications, and applications written in a language such as C, is that Java must perform data conversion. A Java application must convert numbers and text from their internal representation to an encoded format used in messages.

By setting destination, or message properties, you can set the character set and encoding used by WebSphere MQ to encode numbers and text in messages. Normally, you would leave the character set as 1208 and encoding as Native.

WebSphere MQ does not convert byte arrays. To encode strings and character arrays into byte arrays use the `java.nio.charset` package. `Charset` specifies the character set used to convert a string or character array into a byte array. You can also decode a byte array into a string or character array using a `Charset`. It is not good practice to rely on `java.nio.charset.Charset.defaultCodePage` when encoding strings and character arrays. The default `Charset` is typically windows-1252 on Windows, and UTF-8 on UNIX. windows-1252 is a single-byte character set and UTF-8 is a multi-byte character set.

Generally leave the destination character set and encoding properties at their default values of UTF-8 and Native when exchanging messages with other JMS applications. If you are exchanging messages containing numbers or text with a JMS application, choose one of the `JMSTextMessage`, `JMSStreamMessage`, `JMSMapMessage`, or `JMSObjectMessage` message types that fit your purpose. There are no other conversion tasks to do.

If you are exchanging messages with non-JMS applications that use a record format, it is more complicated. Unless the entire record contains text and can be transferred as a `JMSTextMessage`, you must encode and decode text in the application. Set the destination message type to MQ, and use `JMSBytesMessage` to avoid the IBM WebSphere MQ classes for JMS adding additional header and tagging information to the message data. Use `JMSBytesMessage` methods to write numbers and bytes, and the `Charset` class convert text into byte arrays explicitly. A number of factors might influence your choice of character set:

- Performance: Can you reduce the number of conversions by transforming text into a character set that is used on the largest number of servers?
- Uniformity: Transfer all messages in the same character set.
- Richness: What character sets have all the code points that applications must use?
- Simplicity: Single-byte character sets are simpler to use than variable length and multibyte character sets.

See [“Exchanging a formatted record with a non-JMS application” on page 811](#). for examples of converting messages exchanged with non-JMS applications.

Examples

Table of message types and conversion types

<i>Table 119. Message types and conversion types</i>				
	Conversion type			
Message type	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

Calling data conversion from a C program

```
gmo.Options = MQGMO_WAIT          /* wait for new messages          */
              | MQGMO_NO_SYNCPOINT /* no transaction                */
              | MQGMO_CONVERT;    /* convert if necessary          */

while (CompCode != MQCC_FAILED) {
    buflen = sizeof(buffer) - 1; /* buffer size available for GET */
    memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
    memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
    md.Encoding = MQENC_NATIVE;
    md.CodedCharSetId = MQCCSI_Q_MGR;

    MQGET(Hcon,          /* connection handle          */
          Hobj,          /* object handle              */
          &md,            /* message descriptor          */
          &gmo,          /* get message options         */
          buflen,        /* buffer length               */
          buffer,        /* message buffer              */
          &messlen,      /* message length              */
          &CompCode,     /* completion code             */
          &Reason);      /* reason code                  */
}
```

Figure 131. Code snippet from `amqsget0.c`

Sending and receiving text in a `JMSBytesMessage`

The code in Figure 132 on page 797 sends a string in a `BytesMessage`. For simplicity, the example sends a single string, for which a `JMSTextMessage` is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called `TEXT_LENGTH` in Figure 133 on page 797. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

```
BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodepage(((MQDestination) destination)
    .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);
```

Figure 132. Sending a String in a `JMSBytesMessage`

```
BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);
```

Figure 133. Receiving a String from a `JMSBytesMessage`

Related concepts

JMS client message conversion and encoding

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

Queue manager data conversion

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

Related tasks

[Exchanging a formatted record with a non-JMS application](#)

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using `JMSBytesMessage`. The exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

Related reference

JMS message types and conversion

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, `JMSObjectMessage`, `JMSTextMessage`, `JMSMapMessage`, `JMSStreamMessage`, and `JMSBytesMessage`.

JMS message types and conversion

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, `JMSObjectMessage`, `JMSTextMessage`, `JMSMapMessage`, `JMSStreamMessage`, and `JMSBytesMessage`.

JMSObjectMessage

`JMSObjectMessage` contains one object, and any objects that it references, serialized into a byte stream by the JVM. Text is serialized into UTF-8, and limited to strings or character arrays of no more than 65534 bytes. An advantage of `JMSObjectMessage` is that applications are not involved in any data conversion issues as long as they use only the methods and attributes of the object. `JMSObjectMessage` provides data conversion for complex objects without the application programmer considering how to encode an object in a message. The disadvantage of using `JMSObjectMessage` is it can be exchanged only with other JMS applications. By choosing one of the other JMS message types, it is possible to exchange JMS messages with non-JMS applications.

[“Sending and receiving a JMSObjectMessage” on page 800](#) shows a `String` object being exchanged in a message.

A JMS client application can receive a `JMSObjectMessage` only in a message that has a JMS-style body. The destination must specify a JMS style body.

JMSTextMessage

`JMSTextMessage` contains a single text string. When a text message is sent, the text Format is set to "MQSTR", `WMQConstants.MQFMT_STRING`. The `CodedCharacterSetId` of the text is set to the coded character set identifier defined for its destination. The text is encoded into the `CodedCharacterSetId` by WebSphere MQ. The `CodedCharacterSetId` and Format fields are either set in the message descriptor, MQMD, or into the JMS fields in an MQRFH2. If the message is defined as having an `WMQ_MESSAGE_BODY_MQ` message body style, or the body style is unspecified, but the target destination is `WMQ_TARGET_DEST_MQ`, then the message descriptor fields are set. Otherwise the message has a JMS RFH2 and the fields are set in the fixed part of the MQRFH2.

An application can override the coded character set identifier defined for a destination. It must set the message property `JMS_IBM_CHARACTER_SET` to a coded character set identifier; see the example in [“Sending and receiving a JMSTextmessage” on page 800](#).

When the JMS client calls the `consumer.receive` method queue manager conversion is optional. Queue manager conversion is enabled by setting the destination property `WMQ_RECEIVE_CONVERSION` to `WMQ_RECEIVE_CONVERSION_QMGR`. The queue manager converts the text message from the `JMS_IBM_CHARACTER_SET` specified for the message before transferring the message to the JMS client. The character set of the converted message is 1208, UTF-8, unless the destination has a different `WMQ_RECEIVE_CCSID`. The `CodedCharacterSetId` in the message that refers to the `JMSTextMessage` is updated to the target character set ID. The text is decoded from the target character set into Unicode by the `getText` method; see the example in [“Sending and receiving a JMSTextmessage” on page 800](#).

A `JMSTextMessage` can be sent in an MQ-style message body, without a JMS MQRFH2 header. The value of the destination attributes, `WMQ_MESSAGE_BODY` and `WMQ_TARGET_DEST` determine the message body style, unless overridden by the application. The application can override the values set on the destination

by calling `destination.setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ)` or `destination.setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ)`.

If you send a `JMSTextMessage` with an MQ style body by sending it to a destination with `WMQ_MESSAGE_BODY` set to `WMQ_MESSAGE_BODY_MQ`, you cannot receive it as a `JMSTextMessage` from the same destination. All messages received from a destination with `WMQ_MESSAGE_BODY` set to `WMQ_MESSAGE_BODY_MQ` are received as a `JMSBytesMessage`. If you try to receive the message as a `JMSTextMessage` it causes an exception, `ClassCastException: com.ibm.jms.JMSBytesMessage cannot be cast to javax.jms.TextMessage`.

Note: Text in a `JMSBytesMessage` is not converted by the JMS client. The client can only receive the text in the message as a byte array. If queue manager conversion is enabled, the text is converted by the queue manager, but the JMS client must still receive it as a byte array in a `JMSBytesMessage`.

It is generally better to use the `WMQ_TARGET_DEST` property to control whether a `JMSTextMessage` is sent with an MQ or JMS body style. You can then receive the message from a destination that has either `WMQ_TARGET_DEST` set to `WMQ_TARGET_DEST_MQ` or `WMQ_TARGET_DEST_JMS`. `WMQ_TARGET_DEST` has no effect on the receiver.

JMSMapMessage and JMSStreamMessage

These two JMS message types are similar. You can read and write primitive types to the messages using methods based on the `DataInputStream` and `DataOutputStream` interfaces; see [“Table of message types and conversion types” on page 803](#). The details are described in [“JMS client message conversion and encoding” on page 804](#). Each primitive is tagged; see [“The JMS message body” on page 789](#).

Numeric data is read and written to the message encoded as XML text. No reference is made to the destination property, `JMS_IBM_ENCODING`. Text data is treated the same way as text in a `JMSTextMessage`. If you were to look at the message contents created by the example in [Figure 138 on page 801](#), all the message data would be in EBCDIC as it was sent with a character set value of 37.

You can send multiple items in a `JMSMapMessage` or `JMSStreamMessage`.

You can retrieve the individual items of data by name from a `JMSMapMessage`, or by position from a `JMSStreamMessage`. Each item is decoded when a `get` or `read` method is called using the `CodedCharacterSetId` value stored in the message. If the method used to retrieve the item returns a different type to the type that was sent, the type is converted. If the type cannot be converted, an exception is thrown. See [Class JMSStreamMessage](#) for details. The example in [“Sending data in a JMSStreamMessage and JMSMapMessage” on page 801](#) illustrates type conversion, and getting the `JMSMapMessage` contents out of sequence.

The `MQRFH2.format` field for the `JMSMapMessage` and `JMSStreamMessage` is set to `"MQSTR "`. If the destination property `WMQ_RECEIVE_CONVERSION` is set to `WMQ_RECEIVE_CONVERSION_QMGR`, the message data is converted by the queue manager before being sent to the JMS client. The `MQRFH2.CodedCharacterSetId` of the message is the `WMQ_RECEIVE_CCSD` of the destination. The `MQRFH2.Encoding` is `Native`. If `WMQ_RECEIVE_CONVERSION` is `WMQ_RECEIVE_CONVERSION_CLIENT_MSG` the `CodedCharacterSetId` and `Encoding` of the `MQRFH2` is the value set by the sender.

A JMS client application can receive a `JMSMapMessage` or `JMSStreamMessage` only in a message that has a JMS-style body, and from a destination that does not specify an MQ style body.

JMSBytesMessage

A `JMSBytesMessage` can contain multiple primitive types. You can read and write primitive types to the messages using methods based on the `DataInputStream` and `DataOutputStream` interfaces; see [“Table of message types and conversion types” on page 803](#). The details are described in [“JMS message types and conversion” on page 798](#).

The encoding of numeric data in the message is controlled by the value of `JMS_IBM_ENCODING` that is set before writing numeric data to the `JMSBytesMessage`. An application can override the default `Native` encoding defined for `JMSBytesMessage` by setting the message property `JMS_IBM_ENCODING`.

Text data can be read and written in UTF-8 using the `readUTF` and `writeUTF`, or in Unicode using the `readChar` and `writeChar` methods. There are no methods that use `CodedCharacterSetId`. Alternatively, the JMS client can encode and decode text into bytes using the `Charset` class. It transfers the bytes between the JVM and message without the WebSphere MQ classes for JMS performing any conversion; see [“Sending and receiving text in a JMSBytesMessage” on page 801](#).

A `JMSBytesMessage` sent to an MQ application is typically sent in an MQ-style message body, without a JMS MQRFH2 header. If it is sent to a JMS application, the message body style is typically JMS. The value of the destination attributes, `WMQ_MESSAGE_BODY` and `WMQ_TARGET_DEST` determine the message body style, unless overridden by the application. The application can override the values set on the destination by calling `destination.setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ)` or `destination.setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ)`.

If you send a `JMSBytesMessage` with an MQ style body, you can receive the message from a destination that defines either an MQ or a JMS message body style. If you send a `JMSBytesMessage` with a JMS style body, then you must receive the message from a destination that defines a JMS message body style. If you do not, the MQRFH2 is treated as part of the user message data, which might not be what you are expecting.

Whether a message has an MQ or a JMS body style, the way it is received is not affected by setting `WMQ_TARGET_DEST`.

The message might be transformed later, by the queue manager, if a `Format` is supplied for the message data, and queue manager data conversion is enabled. Do not use the format field for anything other than specifying the format of the message data, or leave it blank, `MQConstants.MQFMT_NONE`.

You can send multiple items in a `JMSBytesMessage`. Each numeric item is converted when the message is sent using the encoding defined for the message.

You can retrieve the individual items of data from `JMSBytesMessage`. Call read methods in the same order as the write methods were called to create the message. Each numeric item is converted when the message is called using the `Encoding` value stored in the message.

Unlike `JMSMapMessage` and `JMSStreamMessage`, `JMSBytesMessage` contains only data written by the application. No additional data is stored in the message data, such as the XML tags used to define the items in a `JMSMapMessage` and `JMSStreamMessage`. For this reason, use `JMSBytesMessage` to transfer messages formatted for other applications.

Converting between `JMSBytesMessage` and `DataInputStream` and `DataOutputStream` is useful in some applications. Code based on the example, [“Reading and writing messages using DataInputStream and DataOutputStream” on page 802](#), is necessary to use the `com.ibm.mq.header` package with JMS.

Examples

Sending and receiving a JMSObjectMessage

```
ObjectMessage omo = session.createObjectMessage();
omo.setObject(new String("A string"));
producer.send(omo);
...
ObjectMessage omi = (ObjectMessage)consumer.receive();
System.out.println((String)omi.getObject());
...
A string
```

Figure 134. Sending and receiving a JMSObjectMessage

Sending and receiving a JMSTextmessage

A text message cannot contain text in different character sets. The example shows text in different character sets, sent in two different messages.

```
TextMessage tmo = session.createTextMessage();
tmo.setText("Sent in the character set defined for the destination");
producer.send(tmo);
```

Figure 135. Send text message in the character set defined by the destination

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText("Sent in EBCDIC character set 37");
producer.send(tmo);
```

Figure 136. Send text message in ccsid 37

```
TextMessage tmi = (TextMessage)consumer.receive();
System.out.println(tmi.getText());
...
Sent in the character set defined for the destination
```

Figure 137. Receive text message

Sending data in a `JMSStreamMessage` and `JMSMapMessage`

```
StreamMessage smo = session.createStreamMessage();
smo.writeString("256");
smo.writeInt(512);
smo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
producer.send(smo);
...
MapMessage mmo = session.createMapMessage();
mmo.setString("First", "256");
mmo.setInt("Second", 512);
mmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
producer.send(mmo);
...
StreamMessage smi = (StreamMessage)consumer.receive();
System.out.println("Stream: First as float " + smi.readFloat() +
    " Second as String " + smi.readString());
...
Stream: First as float: 256.0, Second as String: 512
...
MapMessage mmi = (MapMessage)consumer.receive();
System.out.println("Map: Second as String " + mmi.getString("Second") +
    " First as double " + mmi.getDouble("First"));
...
Map: Second as String: 512, First as double: 256.0
```

Figure 138. Send data in `JMSStreamMessage` and `JMSMapMessage`

Sending and receiving text in a `JMSBytesMessage`

The code in [Figure 139](#) on page 802 sends a string in a `BytesMessage`. For simplicity, the example sends a single string, for which a `JMSTextMessage` is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called `TEXT_LENGTH` in [Figure 140](#) on page 802. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

```

BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodecpage(((MQDestination) destination)
    .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);

```

Figure 139. Sending a String in a JMSBytesMessage

```

BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);

```

Figure 140. Receiving a String from a JMSBytesMessage

Reading and writing messages using DataInputStream and DataOutputStream

The code in [Figure 141 on page 802](#) creates a JMSBytesMessage using a DataOutputStream.

```

ByteArrayOutputStream bout = new ByteArrayOutputStream();
DataOutputStream dout = new DataOutputStream(bout);
BytesMessage messageOut = prod.session.createBytesMessage();
// messageOut.setIntProperty(WMQConstants.JMS_IBM_ENCODING,
//                             ((MQDestination) (prod.destination)).getIntProperty
//                             (WMQConstants.WMQ_ENCODING));
int ccsidOut = (((MQDestination)prod.destination).getIntProperty(WMQConstants.WMQ_CCSID));
String codePageOut = CCSID.getCodecpage(ccsidOut);
dout.writeInt(ccsidOut);
dout.write(codePageOut.getBytes(codePageOut));
messageOut.writeBytes(bout.toByteArray());
producer.send(messageOut);

```

Figure 141. Send a JMSBytesMessage using a DataOutputStream

The statement that sets the JMS_IBM_ENCODING property is commented out. The statement is valid, if writing directly to a JMSBytesMessage, but has no effect when writing to DataOutputStream. Numbers that are written to the DataOutputStream are encoded in Native encoding. Setting JMS_IBM_ENCODING has no effect.

The code in [Figure 142 on page 802](#) receives a JMSBytesMessage using a DataInputStream.

```

static final int ccsidIn_SIZE = (Integer.SIZE)/8;
...
connection.start();
BytesMessage messageIn = (BytesMessage) consumer.receive();
int messageLength = new Long(messageIn.getBodyLength()).intValue();
byte [] bin = new byte[messageLength];
messageIn.readBytes(bin, messageLength);
DataInputStream din = new DataInputStream(new ByteArrayInputStream(bin));
int ccsidIn = din.readInt();
byte [] codePageByte = new byte[messageLength - ccsidIn_SIZE];
din.read(codePageByte, 0, codePageByte.length);
System.out.println("CCSID " + ccsidIn + " code page " + new String(codePageByte,
    messageIn.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET)));

```

Figure 142. Receive a JMSBytesMessage using a DataInputStream

The code page is printed out using the code page property of the input message data, JMS_IBM_CHARACTER_SET. On input JMS_IBM_CHARACTER_SET is a Java code page and not a numeric coded character set identifier.

Table of message types and conversion types

<i>Table 120. Message types and conversion types</i>				
	Conversion type			
Message type	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

Related concepts

[JMS message conversion approaches](#)

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by WebSphere MQ.

JMS client message conversion and encoding

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

Queue manager data conversion

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

Related tasks

Exchanging a formatted record with a non-JMS application

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using JMSBytesMessage. The exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

JMS client message conversion and encoding

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

Conversion and encoding occur when Java primitives or objects are read or written to and from JMS messages. The conversion is called JMS client data conversion to distinguish it from queue manager data conversion and application data conversion. The conversion takes place strictly when data is read from or written to a JMS message. Text is converted to and from the internal 16 bit Unicode representation ⁶ to the character set used for text in messages. Numeric data is converted to and Java primitive numeric types to the encoding defined for the message. Whether conversion is performed, and what type of conversion is performed, depends on the JMS message type and the read or write operation.

Table 121 on page 804 categorizes the read and write methods for different JMS message types by the type of conversion performed. The conversions types are described in the text following the table.

<i>Table 121. Message types and conversion types</i>				
	Conversion type			
Message type	Text	Numeric	Other	None
JMSObjectMessage				getObject setObject
JMSTextMessage	getText setText			

⁶ Some Unicode representation requires more than 16 bits. See a Java SE reference.

Table 121. Message types and conversion types (continued)

Message type	Conversion type			
	Text	Numeric	Other	None
JMSBytesMessage	readUTF writeUTF	readDouble readFloat readInt readLong readShort readUnsignedShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readUnsignedByte readBytes readChar writeByte writeBytes writeChar
JMSStreamMessage	readString writeString	readDouble readFloat readInt readLong readShort writeDouble writeFloat writeInt writeLong writeShort	readBoolean readObject writeBoolean writeObject	readByte readBytes readChar writeByte writeBytes writeChar
JMSMapMessage	getString setString	getDouble getFloat getInt getLong getShort setDouble setFloat setInt setLong setShort	getBoolean getObject setBoolean setObject	getByte getBytes readChar setByte setBytes setChar

Text

The default CodedCharacterSetId for a destination is 1208, UTF-8. By default, text is converted from Unicode and sent as a UTF-8 text string. On receive, the text is converted from the coded character set in the message received by the client, into Unicode.

The `setText` and `writeString` methods convert text from Unicode into the character set defined for the destination. An application can override the destination character set by setting the message property `JMS_IBM_CHARACTER_SET`. `JMS_IBM_CHARACTER_SET`, when sending a message must be a numeric coded character set identifier⁷.

The code snippets in [“Sending and receiving a JMSTextmessage” on page 808](#) send two messages. One is sent in the character set defined for the destination and the other in character set 37, defined by the application.

⁷ When receiving a message `JMS_IBM_CHARACTER_SET` is a Java Charset code page name.

The `getText` and `readString` methods convert the text in the message from the character set defined in the message into Unicode. The methods use the code page defined in the message property, `JMS_IBM_CHARACTER_SET`. The code page is mapped from `MQRFH2.CodedCharacterSetId` unless the message is an MQ-type message and has no `MQRFH2`. If the message is a MQ-type message, with no `MQRFH2`, the code page is mapped from `MQMD.CodedCharacterSetId`.

The code snippet in [Figure 147 on page 808](#) receives the message that was sent to the destination. The text in the message is converted from code page IBM037 back into Unicode.

Note: A simple way to check that the text is converted to coded character set 37 is to use WebSphere MQ Explorer. Browse the queue and show the properties of the message before it is retrieved.

Contrast the code snippet in [Figure 146 on page 808](#) with the incorrect code snippet in [Figure 143 on page 806](#). In the incorrect snippet the text string is converted twice, once by the application, and again by WebSphere MQ.

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText(new String("Sent in EBCDIC character set 37".getBytes(CCSID.getCodepage(37))));
producer.send(tmo);
```

Figure 143. Incorrect code page conversion

The `writeUTF` method converts text from Unicode to 1208, UTF-8. The text string is prefaced with a 2 byte length. The maximum length of the text string is 65534 bytes. The `readUTF` method reads an item in a message written by the `writeUTF` method. It reads exactly the number of bytes written by the `writeUTF` method.

Numeric

The default numeric encoding for a destination is Native. The Native encoding constant for Java has the value 273, `x'00000111'`, which is the same for all platforms. On receive, the numbers in the message are correctly transformed into numeric Java primitives. The transformation uses the encoding defined in the message and the type returned by the `read` method.

The `send` method converts numbers that are added to a message by the `set` and `write` into the numeric encoding defined for the destination. The destination encoding can be overridden for a message by an application setting the message property, `JMS_IBM_ENCODING`; for example:

```
message.setIntProperty(WMQConstants.JMS_IBM_ENCODING,
    WMQConstants.WMQ_ENCODING_INTEGER_REVERSED);
```

The `get` and `read` numeric methods convert numbers in the message from the numeric encoding defined in the message. They convert the numbers to the type that is specified by the `read` or `get` method; see [The `ENCODING` property](#). The methods use the encoding defined in `JMS_IBM_ENCODING`. The encoding is mapped from `MQRFH2.Encoding` unless the message is an MQ-type message and has no `MQRFH2`. If the message is a MQ-type message, with no `MQRFH2`, then the methods use the encoding defined in `MQMD.Encoding`.

The example in [Figure 148 on page 808](#) shows an application encoding a number in the destination format and sending it in a `JMSStreamMessage`. Compare the example in [Figure 148 on page 808](#) to the example in [Figure 149 on page 808](#). The difference is that `JMS_IBM_ENCODING` must be set in a `JMSBytesMessage`.

Note: A simple way to check that the number is encoded correctly is to use WebSphere MQ Explorer. Browse the queue and show the properties of the message before it is consumed.

Other

The boolean methods encode `true` and `false` as `x'01'` and `x'00'` in a `JMSByteMessage`, `JMSStreamMessage`, and `JMSMapMessage`.

The UTF methods encode and decode Unicode into UTF-8 text strings. The strings are limited to less than 65536 characters, and are preceded by the 2 byte length field.

The Object methods wrap primitive types as objects. Numeric and text types are encoded or converted as if the primitive types had been read or written using the numeric and text methods.

None

The `readByte`, `readBytes`, `readUnsignedByte`, `writeByte`, and `writeBytes` methods get or put single bytes, or arrays of bytes, between the application and the message without conversion. The `readChar` and `writeChar` methods get and put 2 byte Unicode characters between the application and the message without conversion.

Using the `readBytes` and `writeBytes` methods, the application can perform its own code point conversion, as in [“Sending and receiving text in a `JMSBytesMessage`” on page 809](#).

WebSphere MQ does not perform any code page conversion in the client as the message is a `JMSBytesMessage`, and because the `readBytes` and `writeBytes` methods are used. Nonetheless, if the bytes represent text, make sure that code page used by the application matches the coded character set of the destination. The message might be converted again by a queue manager conversion exit. Another possibility is that the receiving JMS client program might follow the convention of converting any byte arrays representing text in the message into strings or characters using the `JMS_IBM_CHARACTER_SET` property in the message.

In this example the client uses the destination coded character set for its conversion:

```
bytes.writeBytes("In the destination code page".getBytes(
    CCSID.getCodepage(((MQDestination) destination)
        .getIntProperty(WMQConstants.WMQ_CCSID))));
```

Alternatively, the client might have chosen a code page and then set the corresponding coded character set in the `JMS_IBM_CHARACTER_SET` property of the message. The WebSphere MQ classes for Java use `JMS_IBM_CHARACTER_SET` to set the `CodedCharacterSetId` field in the JMS properties in the `MQRFH2`, or in the message descriptor, `MQMD`:

```
String codePage = CCSID.getCodepage(37);
message.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, codePage);8
```

If a byte array is written into a `JMSStringMessage` or `JMSMapMessage`, WebSphere MQ classes for JMS does not perform data conversion, as the bytes are typed as hexadecimal data not as text in the `JMSStringMessage` and `JMSMapMessage`.

If the bytes represent characters in your application, you must take into account what code points to read and write to the message. The code in [Figure 144 on page 807](#) follows the convention of using the destination coded character set. If you create the string using the default character set for the JVM, the byte contents depend on the platform. A JVM on Windows typically has a default `Charset` of `windows-1252`, and UNIX, `UTF-8`. Interchange between Windows and UNIX does require that you select an explicit code page for exchanging text as bytes.

```
StreamMessage smo = producer.session.createStreamMessage();
smo.writeBytes("123".getBytes(CCSID.getCodepage(((MQDestination) destination)
    .getIntProperty(WMQConstants.WMQ_CCSID))));
```

Figure 144. Writing bytes representing a string in a `JMSStreamMessage` using the destination character set

Examples

⁸ `SetStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET, codePage)` currently accepts only numeric character set identifiers.

Sending and receiving a JMSTextmessage

A text message cannot contain text in different character sets. The example shows text in different character sets, sent in two different messages.

```
TextMessage tmo = session.createTextMessage();
tmo.setText("Sent in the character set defined for the destination");
producer.send(tmo);
```

Figure 145. Send text message in the character set defined by the destination

```
TextMessage tmo = session.createTextMessage();
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 37);
tmo.setText("Sent in EBCDIC character set 37");
producer.send(tmo);
```

Figure 146. Send text message in ccsid 37

```
TextMessage tmi = (TextMessage)consumer.receive();
System.out.println(tmi.getText());
...
Sent in the character set defined for the destination
```

Figure 147. Receive text message

Encoding examples

Examples showing a number being sent in the encoding defines for a destination. Notice that you must set the JMS_IBM_ENCODING property of a JMSBytesMessage to the value specified for the destination.

```
StreamMessage smo = session.createStreamMessage();
smo.writeInt(256);
producer.send(smo);
...
StreamMessage smi = (StreamMessage)consumer.receive();
System.out.println(smi.readInt());
...
256
```

Figure 148. Sending a number using the destination encoding in a JMSStreamMessage

```
BytesMessage bmo = session.createBytesMessage();
bmo.writeInt(256);
int encoding = ((MQDestination) (destination)).getIntProperty
(WMQConstants.WMQ_ENCODING);
bmo.setIntProperty(WMQConstants.JMS_IBM_ENCODING, encoding);
producer.send(bmo);
...
BytesMessage bmi = (BytesMessage)consumer.receive();
System.out.println(bmi.readInt());
...
256
```

Figure 149. Sending a number using the destination encoding in a JMSBytesMessage

Sending and receiving text in a JMSBytesMessage

The code in [Figure 150](#) on page 809 sends a string in a BytesMessage. For simplicity, the example sends a single string, for which a JMSTextMessage is more appropriate. To receive a text string in bytes message containing a mixture of types, you must know the length of the string in bytes, called `TEXT_LENGTH` in [Figure 151](#) on page 809. Even for a string with a fixed number of characters, the length of the byte representation might be longer.

```
BytesMessage bytes = session.createBytesMessage();
String codePage = CCSID.getCodepage(((MQDestination) destination)
    .getIntProperty(WMQConstants.WMQ_CCSID));
bytes.writeBytes("In the destination code page".getBytes(codePage));
producer.send(bytes);
```

Figure 150. Sending a String in a JMSBytesMessage

```
BytesMessage message = (BytesMessage)consumer.receive();
int TEXT_LENGTH = new Long(message.getBodyLength()).intValue();
byte[] textBytes = new byte[TEXT_LENGTH];
message.readBytes(textBytes, TEXT_LENGTH);
String codePage = message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET);
String textString = new String(textBytes, codePage);
```

Figure 151. Receiving a String from a JMSBytesMessage

Related concepts

JMS message conversion approaches

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by WebSphere MQ.

Queue manager data conversion

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

Related tasks

Exchanging a formatted record with a non-JMS application

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using JMSBytesMessage. The exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

Related reference

JMS message types and conversion

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, JMSObjectMessage, JMSTextMessage, JMSMapMessage, JMSStreamMessage, and JMSBytesMessage.

Queue manager data conversion

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

The queue manager can convert character and numeric data in message data using the values of CodedCharacterSetId, Encoding, and Format set for the message data. For non-JMS applications the conversion capability has always been available by setting the GetMessageOption, GMO_CONVERT. The queue manager conversion capability has not been available to a JMS application receiving a message until V7.0.

You can use queue manager conversion, before V7.0, with a JMS client application that sends a message. The JMS client builds a formatted record, sets the `CodedCharacterSetId`, `Encoding`, and `Format` attributes corresponding to the data placed in the message. A non-JMS receiving application reads the message using `GMO_CONVERT`, and causes a user-written data conversion exit to be called. The data conversion exit is a shared library that has the name set in the `Format` field.

Since V7.0, the queue manager is able to convert messages that are sent to JMS clients. From 7.0.0.0 to 7.0.1.4 inclusive, queue manager conversion is always called for JMS clients. From 7.0.1.5, or from 7.0.1.4 with APAR IC72897 applied, queue manager conversion is controlled by setting the destination property, `WMQ_RECEIVE_CONVERSION`, to `WMQ_RECEIVE_CONVERSION_QMGR`, or `WMQ_RECEIVE_CONVERSION_CLIENT_MSG`. `WMQ_RECEIVE_CONVERSION_CLIENT_MSG` is the default setting, matching the behavior of WebSphere MQ V6.0, which did not support queue manager data conversion for JMS clients. The application can change the destination setting:

```
((MQDestination)destination).setIntProperty(  
    WMQConstants.WMQ_RECEIVE_CONVERSION,  
    WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Or,

```
((MQDestination)destination).setReceiveConversion  
    (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

Figure 152. Enable queue manager data conversion

Queue manager data conversion for a JMS client takes place when the client calls a `consumer.receive` method. Text data is transformed into UTF-8 (1208) by default. Subsequent read and get methods decode text in the received data from UTF-8, creating Java text primitives in their internal Unicode encoding. UTF-8 is not the only target character set from queue manager data conversion. You can choose a different CCSID by setting the `WMQ_RECEIVE_CCSID` destination property.

An application can also change the destination setting, for example setting it to 437, DOS-US:

```
((MQDestination)destination).setIntProperty  
    (WMQConstants.WMQ_RECEIVE_CCSID, 437);
```

Or,

```
((MQDestination)destination).setReceiveCCSID(437);
```

Figure 153. Set target coded character set for queue manager conversion

The reason for changing `WMQ_RECEIVE_CCSID` is specialized; the chosen CCSID makes no difference to the text objects created in the JVM. However, some JVMs, on some platforms, might not be able to handle conversion from the CCSID of text in the message into Unicode. The option gives you a choice of CCSID for any text delivered to the client in the message. Some JMS client platforms have had problems with message text being delivered in UTF-8.

The JMS code is equivalent to the bold text in the C code in [Figure 154 on page 811](#),

```

gmo.Options = MQGMO_WAIT          /* wait for new messages          */
              | MQGMO_NO_SYNCPOINT /* no transaction                */
              | MQGMO_CONVERT;    /* convert if necessary          */

while (CompCode != MQCC_FAILED) {
    buflen = sizeof(buffer) - 1; /* buffer size available for GET */
    memcpy(md.MsgId, MQMI_NONE, sizeof(md.MsgId));
    memcpy(md.CorrelId, MQCI_NONE, sizeof(md.CorrelId));
    md.Encoding = MQENC_NATIVE;
    md.CodedCharSetId = MQCCSI_Q_MGR;

    MQGET(Hcon,          /* connection handle          */
          Hobj,          /* object handle              */
          &md,            /* message descriptor         */
          &gmo,          /* get message options        */
          buflen,        /* buffer length              */
          buffer,        /* message buffer             */
          &messlen,      /* message length             */
          &CompCode,     /* completion code            */
          &Reason);     /* reason code                */
}

```

Figure 154. Code snippet from *amqsgeth0.c*

Note:

Queue manager conversion is only performed on the message data that has a known WebSphere MQ format. MQSTR, or MQCIH are examples of known formats that are predefined. A known format can also be user-defined format, as long as you have supplied a data-conversion exit.

Messages constructed as JMSTextMessage, JMSMapMessage and JMSStreamMessage, have a MQSTR format, and can be converted by the queue manager.

Related concepts

[JMS message conversion approaches](#)

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by WebSphere MQ.

[JMS client message conversion and encoding](#)

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

[“Invoking the data-conversion exit” on page 398](#)

A data-conversion exit is a user-written exit that receives control during the processing of an MQGET call.

Related tasks

[Exchanging a formatted record with a non-JMS application](#)

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using JMSBytesMessage. The exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

Related reference

[JMS message types and conversion](#)

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, JMSObjectMessage, JMSTextMessage, JMSMapMessage, JMSStreamMessage, and JMSBytesMessage.

[Exchanging a formatted record with a non-JMS application](#)

Follow the steps suggested in this task to design and build a data conversion exit, and a JMS client application that can exchange messages with a non-JMS application using JMSBytesMessage. The

exchange of a formatted message with a non-JMS application can take place with or without calling a data conversion exit.

Before you begin

You might be able to design a simpler solution to exchanging messages with a non-JMS application using a `JMSTextMessage`. Eliminate that possibility before following the steps in this task.

About this task

A JMS client is easier to write if it is not involved in the details of formatting JMS messages exchanged with other JMS clients. As long as the message type is `JMSTextMessage`, `JMSMapMessage`, `JMSStreamMessage`, or `JMSObjectMessage`, WebSphere MQ looks after the details of formatting the message. WebSphere MQ deals with differences in code pages and numeric encoding on different platforms.

You can use these message types to exchange messages with non-JMS applications. To do so, you must understand how these messages are constructed by WebSphere MQ classes for JMS. You might be able to modify the non-JMS application to interpret the messages; see [“Mapping JMS messages onto WebSphere MQ messages” on page 775](#).

An advantage of using one of these message types is the JMS client programming does not depend on the type of application that it is exchanging messages with. A disadvantage is that it might require a modification to another program, and you might not be able to change the other program.

An alternative approach is to write a JMS client application that can deal with existing message formats. Often existing messages are fixed format and contain a mixture of unformatted data, text, and numbers. Use the steps in this task, and the example JMS client in [“Writing classes to encapsulate a record layout in a `JMSBytesMessage`” on page 815](#), as a starting point for building a JMS client that can exchange formatted records with non-JMS applications.

Procedure

1. Define the record layout, or use one of the predefined WebSphere MQ header classes.

For handling predefined WebSphere MQ headers, see [Handling WebSphere MQ message headers](#).

[Figure 155 on page 813](#) is an example of a user defined, fixed-length record layout, which can be processed by the data conversion utility.

2. Create the data conversion exit.

Follow the instructions in [Writing a data-conversion exit program](#) to write a data conversion exit.

To try out the example in [“Writing classes to encapsulate a record layout in a `JMSBytesMessage`” on page 815](#), name the data conversion exit MYRECORD.

3. Write Java classes to encapsulate the record layout, and sending and receiving record. Two approaches you might take are:

- Write a class to that reads and writes the `JMSBytesMessage` that contains the record; see [“Writing classes to encapsulate a record layout in a `JMSBytesMessage`” on page 815](#).
- Write a class extending `com.ibm.mq.header.Header` to define the data structure of the record; see [Creating classes for new header types](#).

4. Decide what coded character set to exchange messages in.

See [“Choosing an approach to message conversion: receiver makes good” on page 795](#).

5. Configure the destination to exchange MQ-type messages, without a JMS MQRFH2 header.

Both the sending and receiving destination must be configured to exchange MQ-type messages. You can use the same destination for both sending and receiving.

The application can override the destination message body property:

```
((MQDestination)destination).setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ);
```

The example in “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 815 overrides the destination message body property, ensuring an MQ-style message is sent.

6. Test the solution with JMS and non-JMS applications

Useful tools to test a data conversion exit are:

- The `amqsgetc0.c` sample program is useful to test receiving a message sent by a JMS client. See the suggested modifications to use the example header, `RECORD.h`, in Figure 156 on page 814. With the modifications, `amqsgetc0.c` receives a message sent by the example JMS client, `TryMyRecord.java`; see “Writing classes to encapsulate a record layout in a JMSBytesMessage” on page 815.
- The sample WebSphere MQ browse program, `amqsbcg0.c`, is useful to inspect the contents of the message header, the JMS header, `MQRFH2`, and the message contents.
- The **rfhutil** program, previously available in SupportPac IH03, allows test messages to be captured and stored in files, and then used to drive Message Flows. Output messages can also be read and displayed in a variety of formats. The formats include two types of XML as well as matching against a COBOL copybook. The data can be in EBCDIC or ASCII. An `RFH2` header can be added to the message before the message is sent.

If you try to receive messages using the modified `amqsgetc0.c` sample program, and get an error with reason code 2080, check whether the message has an `MQRFH2`. The modifications assume that the message has been sent to a destination that specifies no `MQRFH2`.

Examples

```
struct RECORD { MQCHAR StrucID[4];
                MQLONG Version;
                MQLONG StructLength;
                MQLONG Encoding;
                MQLONG CodeCharSetId;
                MQCHAR Format[8];
                MQLONG Flags;
                MQCHAR RecordData[32];
            };
```

Figure 155. *RECORD.h*

- Declare the RECORD.h data structure

```
struct tagRECORD {
    MQCHAR4      StrucId;
    MQLONG       Version;
    MQLONG       StrucLength;
    MQLONG       Encoding;
    MQLONG       CCSID;
    MQCHAR8      Format;
    MQLONG       Flags;
    MQCHAR32     RecordData;
};
typedef struct tagRECORD RECORD;
typedef RECORD MQPOINTER PRECORD;
RECORD record;
PRECORD pRecord = &(record);
```

- Modify the MQGET call to use RECORD,

1. Before modification:

```
MQGET(Hcon,          /* connection handle */
      Hobj,          /* object handle */
      &md,            /* message descriptor */
      &gmo,           /* get message options */
      buflen,        /* buffer length */
      buffer,         /* message buffer */
      &messlen,       /* message length */
      &CompCode,      /* completion code */
      &Reason);       /* reason code */
```

2. After modification:

```
MQGET(Hcon,          /* connection handle */
      Hobj,          /* object handle */
      &md,            /* message descriptor */
      &gmo,           /* get message options */
      sizeof(RECORD), /* buffer length */
      pRecord,        /* message buffer */
      &messlen,       /* message length */
      &CompCode,      /* completion code */
      &Reason);       /* reason code */
```

- Change the print statement,

1. From:

```
buffer[messlen] = '\0';          /* add terminator */
printf("message <%s>\n", buffer);
```

2. To:

```
/* buffer[messlen] = '\0';          add terminator */
printf("ccsid <%d>, flags <%d>, message <%32.32s>\n \0",
      md.CodedCharSetId, record.Flags, record.RecordData);
```

Figure 156. Modify amqsget0.c

Related concepts

JMS message conversion approaches

A number of data conversion approaches are open to JMS application designers. These approaches are not exclusive; some applications are likely to use a combination of these approaches. If your application is exchanging only text or is exchanging messages only with other JMS applications, you do not normally consider data conversion. Data conversion is performed automatically for you, by WebSphere MQ.

JMS client message conversion and encoding

The methods you use to do JMS client message conversion and encoding are listed, with code examples of each type of conversion.

Queue manager data conversion

Queue manager data conversion has always been available to non-JMS applications receiving messages from JMS clients. Since V7.0, JMS clients receiving messages also use queue manager data conversion. From 7.0.1.5, or 7.0.1.4 with APAR IC72897, queue manager data conversion is optional.

Related reference

JMS message types and conversion

The choice of message type affects your approach to message conversion. The interaction of message conversion and message type is described for the JMS message types, `JMSObjectMessage`, `JMSTextMessage`, `JMSMapMessage`, `JMSStreamMessage`, and `JMSBytesMessage`.

[Utility for creating conversion-exit code](#)

Writing classes to encapsulate a record layout in a `JMSBytesMessage`

The purpose of this task is to explore, by example, how to combine data conversion and a fixed record layout in a `JMSBytesMessage`. In the task, you create some Java classes to exchange an example record structure in a `JMSBytesMessage`. You can modify the example to write classes to exchange other record structures.

A `JMSBytesMessage` is the best choice of JMS message type to exchange mixed data type records with non-JMS programs. It has no additional data inserted into the message body by the JMS provider. It is therefore the best choice of message type to use if a JMS client program interoperates with an existing IBM WebSphere MQ program. The main challenge in using a `JMSBytesMessage` comes with matching the encoding and character set expected by the other program. A solution is to create a class that encapsulates the record. A class that encapsulates reading and writing a `JMSBytesMessage`, for a specific record type, makes it easier to send and receive fixed-format records in a JMS program. By capturing the generic aspects of the interface in an abstract class, much of the solution can be reused for different record formats. Different record formats can be implemented in classes that extend the abstract generic class.

An alternative approach is to extend the `com.ibm.mq.headers.Header` class. The `Header` class has methods, such as `addMQLONG`, to build a record format in a more declarative way. A disadvantage of using the `Header` class is getting and setting attributes uses a more complicated interpretative interface. Both approaches result in much the same amount of application code.

A `JMSBytesMessage` can encapsulate only a single format, in addition to an `MQRFH2`, in one message, unless each record uses the same format, coded character set, and encoding. The format, encoding, and character set of a `JMSBytesMessage` are properties of all of the message following the `MQRFH2`. The example is written on the assumption that a `JMSBytesMessage` contains only one user record.

Before you begin

1. Your skill level: you must be familiar with Java programming and JMS. No instructions are provided about setting up the Java development environment. It is advantageous to have written a program to exchange a `JMSTextMessage`, `JMSStreamMessage`, or `JMSMapMessage`. You can then see the differences in exchanging a message using a `JMSBytesMessage`.
2. The example requires IBM WebSphere MQ V7.0.
3. The example was created using the Java perspective of the Eclipse workbench. It requires JRE 6.0 or higher. You can use the Java perspective in IBM WebSphere MQ Explorer to develop and run the Java classes. Alternatively use your own Java development environment.
4. Using IBM WebSphere MQ Explorer makes setting up the test environment, and debugging, simpler than using command-line utilities.

About this task

You are guided through creating two classes: `RECORD` and `MyRecord`. Together these two classes encapsulate a fixed-format record. They have methods to get and set attributes. The `get` method reads the record from a `JMSBytesMessage` and the `put` method writes a record to a `JMSBytesMessage`.

The purpose of the task is not to create a production quality class that you can reuse. You might choose to use the examples in the task to get started on your own classes. The purpose of the task

is to provide you with guidance notes, primarily about using character sets, formats, and encoding, when using a `JMSBytesMessage`. Each step in creating the classes is explained, and aspects of using `JMSBytesMessage`, which are sometimes overlooked, are described.

The `RECORD` class is abstract and defines some common fields for a user record. The common fields are modeled on the standard IBM WebSphere MQ header layout of having an eye catcher, a version, and a length field. The encoding, character set, and format fields, found in many IBM WebSphere MQ headers, are omitted. Another header cannot follow a user-defined format. The `MyRecord` class, which extends the `RECORD` class, does so by literally extending the record with additional user fields. A `JMSBytesMessage`, created by the classes, can be processed by the queue manager data conversion exit.

[“Classes used to run example” on page 821](#) includes a full listing of `RECORD` and `MyRecord`. It also includes listings of the extra “scaffolding” classes to test the `RECORD` and `MyRecord`. The extra classes are:

TryMyRecord

The main program to test `RECORD` and `MyRecord`.

EndPoint

An abstract class that encapsulates the JMS connection, destination, and session in a single class. Its interface just meets the needs of testing the `RECORD` and `MyRecord` classes. It is not an established design pattern for writing JMS applications.

Note: The `EndPoint` class includes this line of code after creating a destination:

```
((MQDestination)destination).setReceiveConversion  
    (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
```

In V7.0, from V7.0.1.5, it is necessary to turn on queue manager conversion. It is disabled by default. In V7.0, up to V7.0.1.4 queue manager conversion is enabled by default, and this line of code causes an error.

MyProducer and MyConsumer

Classes that extend `EndPoint`, and create a `MessageConsumer` and `MessageProducer`, connected and ready to accept requests.

Together all the classes make up a complete application you can build and experiment with, to understand how to use data conversion in a `JMSBytesMessage`.

Procedure

1. Create an abstract class to encapsulate the standard fields in an IBM WebSphere MQ header, with a default constructor. Later, you extend the class to tailor the header to your requirements.

```
public abstract class RECORD implements Serializable {  
    private static final long serialVersionUID = -1616617232750561712L;  
    protected final static int UTF8 = 1208;  
    protected final static int MQLONG_LENGTH = 4;  
    protected final static int RECORD_STRUCT_ID_LENGTH = 4;  
    protected final static int RECORD_VERSION_1 = 1;  
    protected final String RECORD_STRUCT_ID = "BLNK";  
    protected final String RECORD_TYPE = "BLANK";  
    private String structID = RECORD_STRUCT_ID;  
    private int version = RECORD_VERSION_1;  
    private int structLength = RECORD_STRUCT_ID_LENGTH + MQLONG_LENGTH * 2;  
    private int headerEncoding = WMQConstants.WMQ_ENCODING_NATIVE;  
    private String headerCharset = "UTF-8";  
    private String headerFormat = RECORD_TYPE;  
  
    public RECORD() {  
        super();  
    }  
}
```

Note:

- a. The attributes, `structID` to `nextFormat`, are listed in the order they are laid out in a standard IBM WebSphere MQ message header.

- b. The attributes, format, messageEncoding, and messageCharset, describe the header itself, and are not part of the header.
 - c. You must decide whether to store the coded character set identifier or character set of the record. Java uses character sets and IBM WebSphere MQ messages use coded character set identifiers. The example code uses character sets.
 - d. int is serialized to MQLONG by IBM WebSphere MQ. MQLONG is 4 bytes.
2. Create the getters and setters for the private attributes.

a) Create or generate the getters:

```
public String getHeaderFormat() { return headerFormat; }
public int getHeaderEncoding() { return headerEncoding; }
public String getMessageCharset() { return headerCharset; }
public int getMessageEncoding() { return headerEncoding; }
public String getStructID() { return structID; }
public int getStructLength() { return structLength; }
public int getVersion() { return version; }
```

b) Create or generate the setters:

```
public void setHeaderCharset(String charset) {
    this.headerCharset = charset; }
public void setHeaderEncoding(int encoding) {
    this.headerEncoding = encoding; }
public void setHeaderFormat(String headerFormat) {
    this.headerFormat = headerFormat; }
public void setStructID(String structID) {
    this.structID = structID; }
public void setStructLength(int structLength) {
    this.structLength = structLength; }
public void setVersion(int version) {
    this.version = version; }
}
```

3. Create a constructor to create a RECORD instance from a JMSBytesMessage.

```
public RECORD(BytesMessage message) throws JMSEException, IOException,
    MQDataException {
    super();
    setHeaderCharset(message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET));
    setHeaderEncoding(message.getIntProperty(WMQConstants.JMS_IBM_ENCODING));
    byte[] structID = new byte[RECORD_STRUCT_ID_LENGTH];
    message.readBytes(structID, RECORD_STRUCT_ID_LENGTH);
    setStructID(new String(structID, getMessageCharset()));
    setVersion(message.readInt());
    setStructLength(message.readInt());
}
```

Note:

- a. The messageCharset and messageEncoding, are captured from the message properties, as they override the values set for the destination. format is not updated. The example does no error checking. If the Record (BytesMessage) constructor is called, it is assumed that the JMSBytesMessage is a RECORD type message. The line "setStructID(new String(structID, getMessageCharset()))" sets the eye catcher.
 - b. The lines of code that complete the method deserialize fields in the message, in order, updating the default values set in the RECORD instance.
4. Create a put method to write the header fields to a JMSBytesMessage.

```
protected BytesMessage put(MyProducer myProducer) throws IOException,
    JMSEException, UnsupportedEncodingException {
    setHeaderEncoding(myProducer.getEncoding());
    setHeaderCharset(myProducer.getCharset());
    myProducer.setMQClient(true);
    BytesMessage bytes = myProducer.session.createBytesMessage();
    bytes.setStringProperty(WMQConstants.JMS_IBM_FORMAT, getHeaderFormat());
    bytes.setIntProperty(WMQConstants.JMS_IBM_ENCODING, getHeaderEncoding());
    bytes.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET,
        myProducer.getCCSID());
    bytes.writeBytes(String.format("%1$-" + RECORD_STRUCT_ID_LENGTH + ".")
```

```

        + RECORD_STRUCT_ID_LENGTH + "s", getStructID())
        .getBytes(getMessageCharset()), 0, RECORD_STRUCT_ID_LENGTH);
    bytes.writeInt(getVersion());
    bytes.writeInt(getStructLength());
    return bytes;
}

```

Note:

- a. MyProducer encapsulates the JMS Connection, Destination, Session, and MessageProducer in a single class. MyConsumer, used later on, encapsulates the JMS Connection, Destination, Session, and MessageConsumer in a single class.
- b. For a JMSBytesMessage, if the encoding is other than Native, the encoding must be set in the message. The destination encoding is copied to the message encoding attribute, JMS_IBM_CHARACTER_SET, and saved as an attribute of the RECORD class.
 - i) "setMessageEncoding(myProducer.getEncoding());" calls "((MQDestination) destination).getIntProperty(WMQConstants.WMQ_ENCODING);" to get the destination encoding.
 - ii) "Bytes.setIntProperty(WMQConstants.JMS_IBM_ENCODING, getMessageEncoding());" sets the message encoding.
- c. The character set used to transform text into bytes is obtained from the destination, and saved as an attribute of the RECORD class. It is not set in the message, because it is not used by the IBM WebSphere MQ classes for JMS when writing a JMSBytesMessage.

"messageCharset = myProducer.getCharset();" calls

```

    public String getCharset() throws UnsupportedEncodingException,
        JMSEException {
        return CCSID.getCodepage(getCCSID());
    }

```

It gets the Java character set from a coded character set identifier.

"CCSID.getCodepage(ccsid)" is in the package com.ibm.mq.headers. The ccsid is obtained from another method in MyProducer, which queries the destination:

```

    public int getCCSID() throws JMSEException {
        return (((MQDestination) destination)
            .getIntProperty(WMQConstants.WMQ_CCSID));
    }

```

- d. "myProducer.setMQClient(true);" overrides the destination setting for the client type, forcing it to an IBM WebSphere MQ MQI client. You might prefer to omit this line of code, as it obscures an administrative configuration error.

"myProducer.setMQClient(true);" calls:

```

    ((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ); }
    if (!getMQDest()) setMQBody();

```

The code has the side-effect of setting the IBM WebSphere MQ body style to unspecified, if it must override a setting of JMS.

Note:

The IBM WebSphere MQ classes for JMS write the format, encoding, and character set identifier of the message into the message descriptor, MQMD, or into the JMS header, MQRFH2. It depends on whether the message has an IBM WebSphere MQ style body. Do not set the MQMD fields manually.

A method exists to set the message descriptor properties manually. It uses the JMS_IBM_MQMD_* properties. You must set the destination property, WMQ_MQMD_WRITE_ENABLED to set the JMS_IBM_MQMD_* properties:

```

    ((MQDestination)destination).setMQMDWriteEnabled(true);

```

You must set the destination property, `WMQ_MQMD_READ_ENABLED`, to read the properties.

Use the `JMS_IBM_MQMD_*` only if you take full control over the whole message payload. Unlike the `JMS_IBM_*` properties, the `JMS_IBM_MQMD_*` properties do not control how IBM WebSphere MQ classes for JMS construct a JMS message. It is possible to create message descriptor properties that conflict with the properties of the JMS message.

- e. The lines of code that completes the method `serialize` in class `asFields` in the message.

The string attributes are padded with blanks. The strings are converted to bytes using the character set defined for the record, and truncated to the length of the message fields.

5. Complete the class by adding the imports.

```
package com.ibm.mq.id;
import java.io.IOException;
import java.io.Serializable;
import java.io.UnsupportedEncodingException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.constants.MQConstants;
import com.ibm.mq.headers.MQDataException;
import com.ibm.msg.client.wmq.WMQConstants;
```

6. Create a class to extend the `RECORD` class to include additional fields. Include a default constructor.

```
public class MyRecord extends RECORD {
    private static final long serialVersionUID = -370551723162299429L;
    private final static int FLAGS = 1;
    private final static String STRUCT_ID = "MYRD";
    private final static int DATA_LENGTH = 32;
    private final static String FORMAT = "MYRECORD";
    private int flags = FLAGS;
    private String recordData = "ABCDEFGHIIJKLMNOPQRSTUVWXYZ012345";

    public MyRecord() {
        super();
        super.setStructID(STRUCT_ID);
        super.setHeaderFormat(FORMAT);
        super.setStructLength(super.getStructLength() + MQLONG_LENGTH
                               + DATA_LENGTH);
    }
}
```

Note:

- a. The `RECORD` subclass, `MyRecord`, customizes the eye catcher, format, and length of the header.

7. Create or generate the getters and setters.

- a) Create the getters:

```
public int getFlags() { return flags; }
public String getRecordData() { return recordData; } .
```

- b) Create the setters:

```
public void setFlags(int flags) {
    this.flags = flags; }
public void setRecordData(String recordData) {
    this.recordData = recordData; }
}
```

8. Create a constructor to create a `MyRecord` instance from a `JMSBytesMessage`.

```
public MyRecord(BytesMessage message) throws JMSEException, IOException,
    MQDataException {
    super(message);
    setFlags(message.readInt());
    byte[] recordData = new byte[DATA_LENGTH];
    message.readBytes(recordData, DATA_LENGTH);
    setRecordData(new String(recordData, super.getMessageCharset()));
}
```

Note:

- a. The fields that make up the standard message template are read first by the RECORD class.
 - b. The recordData text is converted to String using the character set property of the message.
9. Create a static method to get a message from a consumer and create a new MyRecord instance.

```
public static MyRecord get(MyConsumer myConsumer) throws JMSEException,
    MQDataException, IOException {
    BytesMessage message = (BytesMessage) myConsumer.receive();
    return new MyRecord(message);
}
```

Note:

- a. In the example, for brevity, the MyRecord(BytesMessage) constructor is called from the static get method. Typically, you might separate receiving the message from creating a new MyRecord instance.
10. Create a put method to append the customer fields to a JMSBytesMessage containing a message header.

```
public BytesMessage put(MyProducer myProducer) throws JMSEException,
    IOException {
    BytesMessage bytes = super.put(myProducer);
    bytes.writeInt(getFlags());
    bytes.writeBytes(String.format("%1$-" + DATA_LENGTH + "."
        + DATA_LENGTH + "s", getRecordData())
        .getBytes(super.getMessageCharset()), 0, DATA_LENGTH);
    myProducer.send(bytes);
    return bytes;
}
```

Note:

- a. The method calls in the code serialize the attributes in the MyRecord class as fields in the message.
 - The recordData String attribute is padded with blanks, converted to bytes using the character set defined for the record, and truncated to the length of the RecordData fields.
11. Complete the class by adding the include statements.

```
package com.ibm.mq.id;
import java.io.IOException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.headers.MQDataException;
```

Results

Results:

- The results from running the TryMyRecord class:
 - Sending message in coded character set 37, and using a queue manager conversion exit:

```
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true
In flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 273 CCSID UTF-8
```

- Sending message in coded character set 37, and *not* using a queue manager conversion exit:

```
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true
Out flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID 37 MQ true
In flags 1 text ABCDEFGHIJKLMNOPQRSTUVWXYZ012345 Encoding 546 CCSID IBM037
```

- The results from modifying the TryMyRecord class not to receive the message, and instead receiving it using the modified amqsget0.c sample. The modified sample accepts a formatted record; see [Figure 156 on page 814](#) in “Exchanging a formatted record with a non-JMS application” on page 811.

- Sending message in coded character set 37, and using a queue manager conversion exit:

```
Sample AMQSGETO start
ccsid <850>, flags <1>, message <ABCDEFGHIJKLMNOPQRSTUVWXYZ012345>
no more messages
Sample AMQSGETO end
```

- Sending message in coded character set 37, and *not* using a queue manager conversion exit:

```
Sample AMQSGETO start
MQGET ended with reason code 2110
ccsid <37>, flags <1>, message <---+~ãÄ++ÐÊËËiÐÎÐ+ÔðööµþÚ-±=¾¶§>
no more messages
Sample AMQSGETO end
```

To try out the example and experiment with different code pages and a data conversion exit. Create the Java classes, configure IBM WebSphere MQ, and run the main program, TryMyRecord; see [Figure 157 on page 822](#).

1. Configure IBM WebSphere MQ and JMS to run the example. The instructions are for running the example on Windows.

1. Create a queue manager

```
crtmqm -sa -u SYSTEM.DEAD.LETTER.QUEUE QM1
strmqm QM1
```

2. Create a queue

```
echo DEFINE QL('Q1') REPLACE | runmqsc QM1
```

3. Create a JNDI directory

```
cd c:\
md JNDI-Directory
```

4. Switch to the JMS bin directory

The JMS Administration program must be run from here. The path is `MQ_INSTALLATION_PATH\java\bin`.

5. Create the following JMS definitions in a file called JMSQM1Q1.txt

```
DEF CF(QM1) PROVIDERVERSION(7) QMANAGER(QM1)
DEF Q(Q1) CCSID(37) ENCODING(RRR) MSGBODY(MQ) QMANAGER(QM1) QUEUE(Q1) TARGCLIENT(MQ)
VERSION(7)
END
```

6. Run the JMSAdmin program to create the JMS resources

```
JMSAdmin < JMSQM1Q1.txt
```

2. You can create, alter, and browse the definitions you have created using IBM WebSphere MQ Explorer.
3. Run TryMyRecord.

Classes used to run example

The classes listed in [Figure 157 on page 822](#) to [Figure 162 on page 826](#) are also available in a compressed file; download [jm25529_.zip](#) or [jm25529_.tar.gz](#).

```
package com.ibm.mq.id;
public class TryMyRecord {
    public static void main(String[] args) throws Exception {
        MyProducer producer = new MyProducer();
        MyRecord outrec = new MyRecord();
        System.out.println("Out flags " + outrec.getFlags() + " text "
            + outrec.getRecordData() + " Encoding "
            + producer.getEncoding() + " CCSID " + producer.getCCSID()
            + " MQ " + producer.getMQDest());
        outrec.put(producer);
        System.out.println("Out flags " + outrec.getFlags() + " text "
            + outrec.getRecordData() + " Encoding "
            + producer.getEncoding() + " CCSID " + producer.getCCSID()
            + " MQ " + producer.getMQDest());
        MyRecord inrec = MyRecord.get(new MyConsumer());
        System.out.println("In flags " + inrec.getFlags() + " text "
            + inrec.getRecordData() + " Encoding "
            + inrec.getMessageEncoding() + " CCSID "
            + inrec.getMessageCharset());
    }
}
```

Figure 157. TryMyRecord

```

package com.ibm.mq.id;
import java.io.IOException;
import java.io.Serializable;
import java.io.UnsupportedEncodingException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.constants.MQConstants;
import com.ibm.mq.headers.MQDataException;
import com.ibm.msg.client.wmq.WMQConstants;

public abstract class RECORD implements Serializable {
    private static final long serialVersionUID = -1616617232750561712L;
    protected final static int UTF8 = 1208;
    protected final static int MQLONG_LENGTH = 4;
    protected final static int RECORD_STRUCT_ID_LENGTH = 4;
    protected final static int RECORD_VERSION_1 = 1;
    protected final String RECORD_STRUCT_ID = "BLNK";
    protected final String RECORD_TYPE = "BLANK ";
    private String structID = RECORD_STRUCT_ID;
    private int version = RECORD_VERSION_1;
    private int structLength = RECORD_STRUCT_ID_LENGTH + MQLONG_LENGTH * 2;
    private int headerEncoding = WMQConstants.WMQ_ENCODING_NATIVE;
    private String headerCharset = "UTF-8";
    private String headerFormat = RECORD_TYPE;

    public RECORD() {
        super();
    }

    public RECORD(BytesMessage message) throws JMSEException, IOException,
        MQDataException {
        super();
        setHeaderCharset(message.getStringProperty(WMQConstants.JMS_IBM_CHARACTER_SET));
        setHeaderEncoding(message.getIntProperty(WMQConstants.JMS_IBM_ENCODING));
        byte[] structID = new byte[RECORD_STRUCT_ID_LENGTH];
        message.readBytes(structID, RECORD_STRUCT_ID_LENGTH);
        setStructID(new String(structID, getMessageCharset()));
        setVersion(message.readInt());
        setStructLength(message.readInt());
    }

    public String getHeaderFormat() { return headerFormat; }
    public int getHeaderEncoding() { return headerEncoding; }
    public String getMessageCharset() { return headerCharset; }
    public int getMessageEncoding() { return headerEncoding; }
    public String getStructID() { return structID; }
    public int getStructLength() { return structLength; }
    public int getVersion() { return version; }

    protected BytesMessage put(MyProducer myProducer) throws IOException,
        JMSEException, UnsupportedEncodingException {
        setHeaderEncoding(myProducer.getEncoding());
        setHeaderCharset(myProducer.getCharset());
        myProducer.setMQClient(true);
        BytesMessage bytes = myProducer.session.createBytesMessage();
        bytes.setStringProperty(WMQConstants.JMS_IBM_FORMAT, getHeaderFormat());
        bytes.setIntProperty(WMQConstants.JMS_IBM_ENCODING, getHeaderEncoding());
        bytes.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET,
            myProducer.getCCSID());
        bytes.writeBytes(String.format("%1$-" + RECORD_STRUCT_ID_LENGTH + " ."
            + RECORD_STRUCT_ID_LENGTH + "s", getStructID())
            .getBytes(getMessageCharset()), 0, RECORD_STRUCT_ID_LENGTH);
        bytes.writeInt(getVersion());
        bytes.writeInt(getStructLength());
        return bytes;
    }

    public void setHeaderCharset(String charset) {
        this.headerCharset = charset;
    }
    public void setHeaderEncoding(int encoding) {
        this.headerEncoding = encoding;
    }
    public void setHeaderFormat(String headerFormat) {
        this.headerFormat = headerFormat;
    }
    public void setStructID(String structID) {
        this.structID = structID;
    }
    public void setStructLength(int structLength) {
        this.structLength = structLength;
    }
    public void setVersion(int version) {
        this.version = version;
    }
}

```

Figure 158. RECORD

```

package com.ibm.mq.id;
import java.io.IOException;
import javax.jms.BytesMessage;
import javax.jms.JMSEException;
import com.ibm.mq.headers.MQDataException;

public class MyRecord extends RECORD {
    private static final long serialVersionUID = -370551723162299429L;
    private final static int FLAGS = 1;
    private final static String STRUCT_ID = "MYRD";
    private final static int DATA_LENGTH = 32;
    private final static String FORMAT = "MYRECORD";
    private int flags = FLAGS;
    private String recordData = "ABCDEFGHIJKLMNOPQRSTUVWXYZ012345";

    public MyRecord() {
        super();
        super.setStructID(STRUCT_ID);
        super.setHeaderFormat(FORMAT);
        super.setStructLength(super.getStructLength() + MQLONG_LENGTH
            + DATA_LENGTH);
    }

    public MyRecord(BytesMessage message) throws JMSEException, IOException,
        MQDataException {
        super(message);
        setFlags(message.readInt());
        byte[] recordData = new byte[DATA_LENGTH];
        message.readBytes(recordData, DATA_LENGTH);
        setRecordData(new String(recordData, super.getMessageCharset()));
    }

    public static MyRecord get(MyConsumer myConsumer) throws JMSEException,
        MQDataException, IOException {
        BytesMessage message = (BytesMessage) myConsumer.receive();
        return new MyRecord(message);
    }

    public int getFlags() { return flags; }
    public String getRecordData() { return recordData; }

    public BytesMessage put(MyProducer myProducer) throws JMSEException,
        IOException {
        BytesMessage bytes = super.put(myProducer);
        bytes.writeInt(getFlags());
        bytes.writeBytes(String.format("%1$-" + DATA_LENGTH + " "
            + DATA_LENGTH + "s", getRecordData())
            .getBytes(super.getMessageCharset()), 0, DATA_LENGTH);
        myProducer.send(bytes);
        return bytes;
    }

    public void setFlags(int flags) {
        this.flags = flags;
    }
    public void setRecordData(String recordData) {
        this.recordData = recordData;
    }
}

```

Figure 159. MyRecord


```

package com.ibm.mq.id;
import java.io.UnsupportedEncodingException;
import javax.jms.Connection;
import javax.jms.ConnectionFactory;
import javax.jms.Destination;
import javax.jms.JMSException;
import javax.jms.Session;
import javax.naming.Context;
import javax.naming.InitialContext;
import javax.naming.NamingException;
import com.ibm.mq.headers.CCSID;
import com.ibm.mq.jms.MQDestination;
import com.ibm.msg.client.wmq.WMQConstants;
public abstract class EndPoint {
    public Context ctx;
    public ConnectionFactory cf;
    public Connection connection;
    public Destination destination;
    public Session session;
    protected EndPoint() throws NamingException, JMSException {
        System.setProperty("java.naming.provider.url", "file:/C:/JNDI-Directory");
        System.setProperty("java.naming.factory.initial",
            "com.sun.jndi.fscontext.RefFSContextFactory");

        ctx = new InitialContext();
        cf = (ConnectionFactory) ctx.lookup("QM1");
        connection = cf.createConnection();
        destination = (Destination) ctx.lookup("Q1");
        ((MQDestination)destination).setReceiveConversion
            (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
        session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE); }
    protected EndPoint(String cFactory, String dest) throws NamingException,
        JMSException {
        System.setProperty("java.naming.provider.url", "file:/C:/JNDI-Directory");
        System.setProperty("java.naming.factory.initial",
            "com.sun.jndi.fscontext.RefFSContextFactory");

        ctx = new InitialContext();
        cf = (ConnectionFactory) ctx.lookup(cFactory);
        connection = cf.createConnection();
        destination = (Destination) ctx.lookup(dest);
        ((MQDestination)destination).setReceiveConversion
            (WMQConstants.WMQ_RECEIVE_CONVERSION_QMGR);
        session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE); }
    public int getCCSID() throws JMSException {
        return (((MQDestination) destination)
            .getIntProperty(WMQConstants.WMQ_CCSID)); }
    public String getCharSet() throws UnsupportedEncodingException,
        JMSException {
        return CCSID.getCodepage(getCCSID()); }
    public int getEncoding() throws JMSException {
        return (((MQDestination) destination)
            .getIntProperty(WMQConstants.WMQ_ENCODING)); }
    public boolean getMQDest() throws JMSException {
        if (((MQDestination) destination).getMessageBodyStyle()
            == WMQConstants.WMQ_MESSAGE_BODY_MQ)
            || (((MQDestination) destination).getMessageBodyStyle()
            == WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED)
            && (((MQDestination) destination).getTargetClient()
            == WMQConstants.WMQ_TARGET_DEST_MQ))
            return true;
        else
            return false; }
    public void setCCSID(int ccsid) throws JMSException {
        ((MQDestination) destination).setIntProperty(WMQConstants.WMQ_CCSID,
            ccsid); }
    public void setEncoding(int encoding) throws JMSException {
        ((MQDestination) destination).setIntProperty(WMQConstants.WMQ_ENCODING,
            encoding); }
    public void setMQBody() throws JMSException {
        ((MQDestination) destination)
            .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED); }
    public void setMQBody(boolean mqbody) throws JMSException {
        if (mqbody) ((MQDestination) destination)
            .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_MQ);
        else
            ((MQDestination) destination)
            .setMessageBodyStyle(WMQConstants.WMQ_MESSAGE_BODY_JMS); }
    public void setMQClient(boolean mqclient) throws JMSException {
        if (mqclient){
            ((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ);
            if (!getMQDest()) setMQBody();
        }
        else
            ((MQDestination) destination).setTargetClient(WMQConstants.WMQ_TARGET_DEST_JMS); }
}

```

Figure 160. EndPoint

```

package com.ibm.mq.id;
import javax.jms.JMSException;
import javax.jms.Message;
import javax.jms.MessageProducer;
import javax.naming.NamingException;
public class MyProducer extends Endpoint {
    public MessageProducer producer;
    public MyProducer() throws NamingException, JMSException {
        super();
        producer = session.createProducer(destination); }
    public MyProducer(String cFactory, String dest) throws NamingException,
        JMSException {
        super(cFactory, dest);
        producer = session.createProducer(destination); }
    public void send(Message message) throws JMSException {
        producer.send(message); }
}

```

Figure 161. MyProducer

```

package com.ibm.mq.id;
import javax.jms.JMSException;
import javax.jms.Message;
import javax.jms.MessageConsumer;
import javax.naming.NamingException;
public class MyConsumer extends Endpoint {
    public MessageConsumer consumer;
    public MyConsumer() throws NamingException, JMSException {
        super();
        consumer = session.createConsumer(destination);
        connection.start(); }
    public MyConsumer(String cFactory, String dest) throws NamingException,
        JMSException {
        super(cFactory, dest);
        consumer = session.createConsumer(destination);
        connection.start(); }
    public Message receive() throws JMSException {
        return consumer.receive(); }
}

```

Figure 162. MyConsumer

Creating and configuring connection factories and destinations in a WebSphere MQ classes for JMS application

A WebSphere MQ classes for JMS application can create connection factories and destinations by retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace, by using the IBM JMS extensions, or by using the WebSphere MQ JMS extensions. An application can also use the IBM JMS extensions or WebSphere MQ JMS extensions to set the properties of connection factories and destinations.

Connection factories and destinations are starting points in the flow of logic of a JMS application. An application uses a ConnectionFactory object to create a connection to a messaging server, and uses a Queue or Topic object as a target to send messages to or a source from which to receive messages. An application therefore needs to create at least one connection factory and one or more destinations. Having created a connection factory or destination, the application might then need to configure the object by setting one or more of its properties.

In summary, an application can create and configure connection factories and destinations in the following ways:

Using JNDI to retrieve administered objects

An administrator can use the WebSphere MQ JMS administration tool or WebSphere MQ Explorer to create and configure connection factories and destinations as administered objects in a JNDI namespace. An application can then retrieve the administered objects from the JNDI namespace. Having retrieved an administered object, the application can, if required, set or change one or more of its properties by using either the IBM JMS extensions or the WebSphere MQ JMS extensions.

Using the IBM JMS extensions

An application can use the IBM JMS extensions to create connection factories and destinations dynamically at run time. The application first creates a JmsFactoryFactory object, and then uses

methods of this object to create connection factories and destinations. Having created a connection factory or destination, the application can use methods inherited from the `JmsPropertyContext` interface to set its properties. Alternatively, the application can use a uniform resource identifier (URI) to specify one or more properties of a destination when it creates the destination.

Using the WebSphere MQ JMS extensions

An application can also use the WebSphere MQ JMS extensions to create connection factories and destinations dynamically at run time. The application uses the supplied constructors to create connection factories and destinations. Having created a connection factory or destination, the application can use methods of the object to set its properties. Alternatively, the application can use a URI to specify one or more properties of a destination when it creates the destination.

Using JNDI to retrieve administered objects in a JMS application

To retrieve administered objects from a Java Naming and Directory Interface (JNDI) namespace, a JMS application must create an initial context and then use the `lookup()` method to retrieve the objects.

Before an application can retrieve administered objects from a JNDI namespace, an administrator must first create the administered objects. The administrator can use the WebSphere MQ JMS administration tool or WebSphere MQ Explorer to create and maintain administered objects in a JNDI namespace. For information about how to use the WebSphere MQ JMS administration tool, see [“Using the WebSphere MQ JMS administration tool”](#) on page 893. For information about how to use WebSphere MQ Explorer, see the help provided with WebSphere MQ Explorer. An application server, however, typically provides its own repository for administered objects and its own tools for creating and maintaining the objects.

To retrieve administered objects from a JNDI namespace, an application must first create an initial context, as shown in the following example:

```
import javax.jms.*;
import javax.naming.*;
import javax.naming.directory.*;
.
.
String url = "ldap://server.company.com/o=company_us,c=us";
String icf = "com.sun.jndi.ldap.LdapCtxFactory";
.
java.util.Hashtable environment = new java.util.Hashtable();
environment.put(Context.PROVIDER_URL, url);
environment.put(Context.INITIAL_CONTEXT_FACTORY, icf);
Context ctx = new InitialDirContext(environment);
```

In this code, the String variables `url` and `icf` have the following meanings:

url

The uniform resource locator (URL) of the directory service. The URL can have one of the following formats:

- `ldap://hostname/contextName`, for a directory service based on an LDAP server
- `file://directoryPath`, for a directory service based on the local file system

icf

The class name of the initial context factory, which can be one of the following values:

- `com.sun.jndi.ldap.LdapCtxFactory`, for a directory service based on an LDAP server
- `com.sun.jndi.fscontext.RefFSContextFactory`, for a directory service based on the local file system

Note that some combinations of a JNDI package and a Lightweight Directory Access Protocol (LDAP) service provider can cause LDAP error 84 to occur. To resolve this problem, insert the following line of code before the call to `InitialDirContext()`:

```
environment.put(Context.REFERRAL, "throw");
```

After an initial context is obtained, the application can retrieve administered objects from the JNDI namespace by using the `lookup()` method, as shown in the following example:

```

ConnectionFactory factory;
Queue queue;
Topic topic;
.
.
.
factory = (ConnectionFactory)ctx.lookup("cn=myCF");
queue = (Queue)ctx.lookup("cn=myQ");
topic = (Topic)ctx.lookup("cn=myT");

```

This code retrieves the following objects from an LDAP based namespace:

- A ConnectionFactory object bound with the name myCF
- A Queue object bound with the name myQ
- A Topic object bound with the name myT

Using the IBM JMS extensions

WebSphere MQ classes for JMS contains a set of extensions to the JMS API called the IBM JMS extensions. An application can use these extensions to create connection factories and destinations dynamically at run time, and to set the properties of WebSphere MQ classes for JMS objects. The extensions can be used with any messaging provider.

The IBM JMS extensions are a set of interfaces and classes in the following packages:

- com.ibm.msg.client.jms
- com.ibm.msg.client.services

The packages can be found in com.ibm.mqjms.jar which is located in <MQ_Install_Dir>/java/lib.

These extensions provide the following function:

- A factory-based mechanism for creating connection factories and destinations dynamically at run time, instead of retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace
- A set of methods for setting the properties of WebSphere MQ classes for JMS objects
- A set of exception classes with methods for obtaining detailed information about a problem
- A set of methods for controlling tracing
- A set of methods for obtaining version information about WebSphere MQ classes for JMS

With regard to creating connection factories and destinations dynamically at run time, and setting and getting their properties, the IBM JMS extensions provide an alternative set of interfaces to the WebSphere MQ JMS extensions. However, whereas the WebSphere MQ JMS extensions are specific to the WebSphere MQ messaging provider, the IBM JMS extensions are not specific to WebSphere MQ and can be used with any messaging provider within the layered architecture described in [“A layered architecture” on page 763](#).

The interface com.ibm.msg.client.wmq.WMQConstants contains the definitions of constants, which an application can use when setting the properties of WebSphere MQ classes for JMS objects using the IBM JMS extensions. The interface contains constants for the WebSphere MQ messaging provider and JMS constants that are independent of any messaging provider.

The examples of code that follow assume that the following import statements have been run:

```

import com.ibm.msg.client.jms.*;
import com.ibm.msg.client.services.*;
import com.ibm.msg.client.wmq.WMQConstants;

```

Creating connection factories and destinations

Before an application can create connection factories and destinations using the IBM JMS extensions, it must first create a `JmsFactoryFactory` object. To create a `JmsFactoryFactory` object, the application calls the `getInstance()` method of the `JmsFactoryFactory` class, as shown in the following example:

```
JmsFactoryFactory ff = JmsFactoryFactory.getInstance(JmsConstants.WMQ_PROVIDER);
```

The parameter on the `getInstance()` call is a constant that identifies the WebSphere MQ messaging provider as the chosen messaging provider. The application can then use the `JmsFactoryFactory` object to create connection factories and destinations.

To create a connection factory, the application calls the `createConnectionFactory()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsConnectionFactory factory = ff.createConnectionFactory();
```

This statement creates a `JmsConnectionFactory` object with the default values for all its properties, which means that the application connects to the default queue manager in bindings mode. If you want an application to connect in client mode, or connect to a queue manager other than the default queue manager, the application must set the appropriate properties of the `JmsConnectionFactory` object before creating the connection. For information about how to do this, see [“Setting the properties of WebSphere MQ classes for JMS objects”](#) on page 830.

The `JmsFactoryFactory` class also contains methods to create connection factories of the following types:

- `JmsQueueConnectionFactory`
- `JmsTopicConnectionFactory`
- `JmsXAConnectionFactory`
- `JmsXAQueueConnectionFactory`
- `JmsXATopicConnectionFactory`

To create a Queue object, the application calls the `createQueue()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsQueue q1 = ff.createQueue("Q1");
```

This statement creates an `JmsQueue` object with the default values for all its properties. The object represents a WebSphere MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

The `createQueue()` method can also accept a queue uniform resource identifier (URI) as a parameter. A queue URI is a string that specifies the name of a WebSphere MQ queue and, optionally, the name of the queue manager that owns the queue, and one or more properties of the `JmsQueue` object. The following statement contains an example of a queue URI:

```
JmsQueue q2 = ff.createQueue("queue://QM2/Q2?persistence=2&priority=5");
```

The `JmsQueue` object created by this statement represents a WebSphere MQ queue called Q2 that is owned by queue manager QM2, and all messages sent to this destination are persistent and have a priority of 5. For more information about queue URIs, see [“Uniform resource identifiers \(URIs\)”](#) on page 841. For an alternative way of setting the properties of a `JmsQueue` object, see [“Setting the properties of WebSphere MQ classes for JMS objects”](#) on page 830.

To create a Topic object, an application can use the `createTopic()` method of the `JmsFactoryFactory` object, as shown in the following example:

```
JmsTopic t1 = ff.createTopic("Sport/Football/Results");
```

This statement creates a `JmsTopic` object with the default values for all its properties. The object represents a topic called Sport/Football/Results.

The `createTopic()` method can also accept a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the `JmsTopic` object. The following statements contain an example of a topic URI:

```
String s1 = "topic://Sport/Tennis/Results?persistence=1&priority=0";
JmsTopic t2 = ff.createTopic(s1);
```

The `JmsTopic` object created by these statements represents a topic called `Sport/Tennis/Results`, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see [“Uniform resource identifiers \(URIs\)”](#) on page 841. For an alternative way of setting the properties of a `JmsTopic` object, see [“Setting the properties of WebSphere MQ classes for JMS objects”](#) on page 830.

After an application has created a connection factory or destination, that object can be used only with the selected messaging provider.

Setting the properties of WebSphere MQ classes for JMS objects

To set the properties of WebSphere MQ classes for JMS objects using the IBM JMS extensions, an application uses the methods of the `com.ibm.msg.client.JmsPropertyContext` interface.

For each Java data type, the `JmsPropertyContext` interface contains a method to set the value of a property with that data type, and a method to get the value of a property with that data type. For example, an application calls the `setIntProperty()` method to set a property with an integer value, and calls the `getIntProperty()` method to get a property with an integer value.

Instances of classes in the `com.ibm.mq.jms` package also inherit the methods of the `JmsPropertyContext` interface. An application can therefore use these methods to set the properties of `MQConnectionFactory`, `MQQueue`, and `MQTopic` objects.

When an application creates a WebSphere MQ classes for JMS object, any properties with default values are set automatically. When an application sets a property, the new value replaces any previous value the property had. After a property has been set, it cannot be deleted, but its value can be changed.

If an application attempts to set a property to a value that is not valid value for the property, WebSphere MQ classes for JMS throws a `JMSEException` exception. If an application attempts to get a property that has not been set, the behavior is as described in the JMS specification. WebSphere MQ classes for JMS throws a `NumberFormatException` exception for primitive data types and returns null for referenced data types.

In addition to the predefined properties of a WebSphere MQ classes for JMS object, an application can set its own properties. These application defined properties are ignored by WebSphere MQ classes for JMS.

For more information about the properties of WebSphere MQ classes for JMS objects, see [Properties of IBM WebSphere MQ classes for JMS objects](#).

The following code is an example of how to set properties using the IBM JMS extensions. The code sets five properties of a connection factory.

```
factory.setIntProperty(WMQConstants.WMQ_CONNECTION_MODE,
    WMQConstants.WMQ_CM_CLIENT);
factory.setStringProperty(WMQConstants.WMQ_QUEUE_MANAGER, "QM1");
factory.setStringProperty(WMQConstants.WMQ_HOST_NAME, "HOST1");
factory.setIntProperty(WMQConstants.WMQ_PORT, 1415);
factory.setStringProperty(WMQConstants.WMQ_CHANNEL, "QM1.SVR");
factory.setStringProperty(WMQConstants.WMQ_APPLICATIONNAME, "My Application");
```

The effect of setting these properties is that the application connects to queue manager QM1 in client mode, using an MQI channel called QM1.SVR. The queue manager is running on a system with host name HOST1, and the listener for the queue manager is listening in port number 1415. This connection and other queue manager connections associated with sessions under it, have the application name "My Application" associated with them.

Note: Queue managers running on z/OS platforms do not support setting application names, and this setting is therefore ignored.

The `JmsPropertyContext` interface also contains the `setObjectProperty()` method, which an application can use to set properties. The second parameter of the method is an object that encapsulates the value of the property. For example, the following code creates an `Integer` object that encapsulates the integer 1415, and then calls `setObjectProperty()` to set the `PORT` property of a connection factory to the value 1415:

```
Integer port = new Integer(1415);
factory.setObjectProperty(WMQConstants.WMQ_PORT, port);
```

This code is therefore equivalent to the following statement:

```
factory.setIntProperty(WMQConstants.WMQ_PORT, 1415);
```

Conversely, the `getObjectProperty()` method returns an object that encapsulates the value of a property.

Implicit conversion of a property value from one data type to another

When an application uses a method of the `JmsPropertyContext` interface to set or get the property of a WebSphere MQ classes for JMS object, the value of the property can be implicitly converted from one data type to another.

For example, the following statement sets the `PRIORITY` property of the `JmsQueue` object `q1`:

```
q1.setStringProperty(WMQConstants.WMQ_PRIORITY, "5");
```

The `PRIORITY` property has an integer value, and so the `setStringProperty()` call implicitly converts the string "5" (the source value) to the integer 5 (the target value), which then becomes the value of the `PRIORITY` property.

Conversely, the following statement gets the `PRIORITY` property of the `JmsQueue` object `q1`:

```
String s1 = q1.getStringProperty(WMQConstants.WMQ_PRIORITY);
```

The integer 5 (the source value), which is the value of the `PRIORITY` property, is implicitly converted to the string "5" (the target value) by the `getStringProperty()` call.

The conversions supported by WebSphere MQ classes for JMS are shown in [Table 122 on page 831](#).

Table 122. Supported conversions from one data type to another	
Source data type	Supported target data types
boolean	String
byte	int, long, short, String
char	String
double	String
float	double, String
int	long, String
long	String
short	int, long, String
String	boolean, byte, double, float, int, long, short

The general rules governing the supported conversions are as follows:

- Numeric values can be converted from one data type to another provided no data is lost during the conversion. For example, a value with data type `int` can be converted into a value with data type `long`, but cannot be converted into a value with data type `short`.
- A value of any data type can be converted into a string.

- A string can be converted to a value of any other data type (except `char`) provided the string is in the correct format for the conversion. If an application attempts to convert a string that is not in the correct format, WebSphere MQ classes for JMS throws a `NumberFormatException` exception.
- If an application attempts a conversion that is not supported, WebSphere MQ classes for JMS throws a `MessageFormatException` exception.

The specific rules for converting a value from one data type to another are as follows:

- When converting a boolean value to a string, the value `true` is converted to the string "true", and the value `false` is converted to the string "false".
- When converting a string to a boolean value, the string "true" (not case-sensitive) is converted to `true`, and the string "false" (not case-sensitive) is converted to `false`. Any other string is converted to `false`.
- When converting a string to a value with data type `byte`, `int`, `long`, or `short`, the string must have the following format:

[blanks][sign]digits

The meanings of the components of the string are as follows:

blanks

Optional leading blank characters.

sign

An optional plus sign (+) or minus sign (-).

digits

A contiguous sequence of digits (0-9). At least one digit must be present.

After the sequence of digits, the string can contain other characters that are not digits, but the conversion stops as soon as the first of these characters is reached. The string is assumed to represent a decimal integer.

If the string is not in the correct format, WebSphere MQ classes for JMS throws a `NumberFormatException` exception.

- When converting a string to a value with data type `double` or `float`, the string must have the following format:

[blanks][sign]digits[e_char[e_sign]e_digits]

The meanings of the components of the string are as follows:

blanks

Optional leading blank characters.

sign

An optional plus sign (+) or minus sign (-).

digits

A contiguous sequence of digits (0-9). At least one digit must be present.

e_char

An exponent character, which is either *E* or *e*.

e_sign

An optional plus sign (+) or minus sign (-) for the exponent.

e_digits

A contiguous sequence of digits (0-9) for the exponent. At least one digit must be present if the string contains an exponent character.

After the sequence of digits, or the optional characters representing an exponent, the string can contain other characters that are not digits, but the conversion stops as soon as the first of these characters is reached. The string is assumed to represent a decimal floating point number with an exponent that is a power of 10.

If the string is not in the correct format, WebSphere MQ classes for JMS throws a `NumberFormatException` exception.

- When converting a numeric value (including a value with data type byte) to a string, the value is converted to the string representation of the value as a decimal number, not the string containing the ASCII character for that value. For example, the integer 65 is converted to the string "65", not the string "A".

Setting more than one property in a single call

The `JmsPropertyContext` interface also contains the `setBatchProperties()` method, which an application can use to set more than one property in a single call. The parameter of the method is a `Map` object that encapsulates a set of property name-value pairs.

For example, the following code uses the `setBatchProperties()` method to set the same five properties of a connection factory as shown in [“Setting the properties of WebSphere MQ classes for JMS objects” on page 830](#). The code creates an instance of the `HashMap` class, which implements the `Map` interface.

```
HashMap batchProperties = new HashMap();
batchProperties.put(WMQConstants.WMQ_CONNECTION_MODE,
    new Integer(WMQConstants.WMQ_CM_CLIENT));
batchProperties.put(WMQConstants.WMQ_QUEUE_MANAGER, "QM1");
batchProperties.put(WMQConstants.WMQ_WMQ_HOST_NAME, "HOST1");
batchProperties.put(WMQConstants.WMQ_PORT, "1414");
batchProperties.put(WMQConstants.WMQ_CHANNEL, "QM1.SVR");
factory.setBatchProperties(batchProperties);
```

Note that the second parameter of the `Map.put()` method must be an object. Therefore a property value with a primitive data type must be encapsulated within an object or represented by a string, as shown in the example.

The `setBatchProperties()` method validates each property. If the `setBatchProperties()` method cannot set a property because, for example, its value is not valid, none of the specified properties are set.

Property names and values

If an application uses the methods of the `JmsPropertyContext` interface to set and get the properties of WebSphere MQ classes for JMS objects, the application can specify the names and values of properties in any of the following ways. Each of the accompanying examples shows how to set the `PRIORITY` property of the `JmsQueue` object `q1` so that a message sent to the queue has the priority specified on the `send()` call.

Using the property names and values that are defined as constants in the `com.ibm.msg.client.wmq.WMQConstants` interface

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setIntProperty(WMQConstants.WMQ_PRIORITY, WMQConstants.WMQ_PRI_APP);
```

Using the property names and values that can be used in queue and topic uniform resource identifiers (URIs)

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setIntProperty("priority", -2);
```

Only the names and values of properties of destinations can be specified in this way.

Using the property names and values that are recognized by the WebSphere MQ JMS administration tool

The following statement is an example of how to specify the names and values of properties in this way:

```
q1.setStringProperty("PRIORITY", "APP");
```

The short form of the property name is also acceptable, as shown in the following statement:

```
q1.setStringProperty("PRI", "APP");
```

When an application gets a property, the value returned depends on the way in which the application specifies the name of the property. For example, if an application specifies the constant `WMQConstants.WMQ_PRIORITY` as the property name, the value returned is the integer `-2`:

```
int n1 = getIntProperty(WMQConstants.WMQ_PRIORITY);
```

The same value is returned if the application specifies the string `"priority"` as the property name:

```
int n2 = getIntProperty("priority");
```

However, if the application specifies the string `"PRIORITY"` or `"PRI"` as the property name, the value returned is the string `"APP"`:

```
String s1 = getStringProperty("PRI");
```

Internally, WebSphere MQ classes for JMS stores property names and values as the literal values defined in the `com.ibm.msg.client.wmq.WMQConstants` interface. This is the defined canonical format for property names and values. As a general rule, if an application sets properties using one of the other two ways of specifying property names and values, WebSphere MQ classes for JMS has to convert the names and values from the specified input format into the canonical format. Similarly, if an application gets properties using one of the other two ways of specifying property names and values, WebSphere MQ classes for JMS must convert the names from the specified input format into the canonical format, and convert the values from the canonical format into the required output format. Having to perform these conversions might have implications for performance.

Property names and values returned by exceptions, in trace files, or in the WebSphere MQ classes for JMS log are always in the canonical format.

Using the Map interface

The `JmsPropertyContext` interface extends the `java.util.Map` interface. An application can therefore use the methods of the `Map` interface to access the properties of a WebSphere MQ classes for JMS object.

For example, the following code prints out the names and values of all the properties of a connection factory. The code uses only the methods of the `Map` interface to get the names and values of the properties.

```
// Get the names of all the properties
Set propNames = factory.keySet();

// Loop round all the property names and get the property values
Iterator iterator = propNames.iterator();
while (iterator.hasNext()){
    String pName = (String)iterator.next();
    System.out.println(pName+"="+factory.get(pName));
}
```

Using the methods of the `Map` interface does not bypass any property validations or conversions.

Using the WebSphere MQ JMS extensions

WebSphere MQ classes for JMS contains a set of extensions to the JMS API called the WebSphere MQ JMS extensions. An application can use these extensions to create connection factories and destinations dynamically at run time, and to set the properties of connection factories and destinations.

WebSphere MQ classes for JMS contains a set of classes in the packages `com.ibm.jms` and `com.ibm.mq.jms`. These classes implement the JMS interfaces and contain the WebSphere MQ JMS extensions. The examples of code that follow assume that these packages have been imported by the following statements:

```
import com.ibm.jms.*;
import com.ibm.mq.jms.*;
```

An application can use the WebSphere MQ JMS extensions to perform the following functions:

- Create connection factories and destinations dynamically at run time, instead of retrieving them as administered objects from a Java Naming and Directory Interface (JNDI) namespace
- Set the properties of connection factories and destinations

Creating connection factories

To create a connection factory, an application can use the `MQConnectionFactory` constructor, as shown in the following example:

```
MQConnectionFactory factory = new MQConnectionFactory();
```

This statement creates an `MQConnectionFactory` object with the default values for all its properties, which means that the application connects to the default queue manager in bindings mode. If you want an application to connect in client mode, or connect to a queue manager other than the default queue manager, the application must set the appropriate properties of the `MQConnectionFactory` object before creating the connection. For information about how to do this, see [“Setting the properties of connection factories” on page 835](#).

An application can create connection factories of the following types in a similar way:

- `MQQueueConnectionFactory`
- `MQTopicConnectionFactory`
- `MQXAConnectionFactory`
- `MQXAQueueConnectionFactory`
- `MQXATopicConnectionFactory`

Setting the properties of connection factories

An application can set the properties of a connection factory by calling the appropriate methods of the connection factory. The connection factory can either be an administered object or an object created dynamically at run time.

Consider the following code, for example:

```
MQConnectionFactory factory = new MQConnectionFactory();
factory.setTransportType(WMQConstants.WMQ_CM_CLIENT);
factory.setQueueManager("QM1");
factory.setHostName("HOST1");
factory.setPort(1415);
factory.setChannel("QM1.SVR");
```

This code creates an `MQConnectionFactory` object and then sets five properties of the object. The effect of setting these properties is that the application connects to queue manager QM1 in client mode using an MQI channel called QM1.SVR. The queue manager is running on a system with host name HOST1, and the listener for the queue manager is listening in port number 1415.

For a real-time connection to a broker, an application can use the following code:

```
MQConnectionFactory factory = new MQConnectionFactory();
factory.setTransportType(WMQConstants.WMQ_CM_DIRECT);
factory.setHostName("HOST2");
factory.setPort(1507);
```

This code assumes that the broker is running on a system with host name HOST2 and listening on port number 1507.

An application that uses a real-time connection to a broker can use only the publish/subscribe style of messaging. It cannot use the point-to-point style of messaging.

Only certain combinations of properties of a connection factory are valid. For information about which combinations are valid, see [Dependencies between properties of WebSphere MQ classes for JMS objects](#).

For more information about the properties of a connection factory, and the methods used to set its properties, see [Properties of IBM WebSphere MQ classes for JMS objects](#).

Creating destinations

To create a Queue object, an application can use the `MQQueue` constructor, as shown in the following example:

```
MQQueue q1 = new MQQueue("Q1");
```

This statement creates an `MQQueue` object with the default values for all its properties. The object represents a WebSphere MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

An alternative form of the `MQQueue` constructor has two parameters, as shown in the following example:

```
MQQueue q2 = new MQQueue("QM2", "Q2");
```

The `MQQueue` object created by this statement represents a WebSphere MQ queue called Q2 that is owned by queue manager QM2. The queue manager identified in this way can be the local queue manager or a remote queue manager. If it is a remote queue manager, WebSphere MQ must be configured so that, when the application sends a message to this destination, WebSphere MQ can route the message from the local queue manager to the remote queue manager.

The `MQQueue` constructor can also accept a queue uniform resource identifier (URI) as a single parameter. A queue URI is a string that specifies the name of a WebSphere MQ queue and, optionally, the name of the queue manager that owns the queue, and one or more properties of the `MQQueue` object. The following statement contains an example of a queue URI:

```
MQQueue q3 = new MQQueue("queue://QM3/Q3?persistence=2&priority=5");
```

The `MQQueue` object created by this statement represents a WebSphere MQ queue called Q3 that is owned by queue manager QM3, and all messages sent to this destination are persistent and have a priority of 5. For more information about queue URIs, see [“Uniform resource identifiers \(URIs\)” on page 841](#). For an alternative way of setting the properties of an `MQQueue` object, see [“Setting the properties of destinations” on page 837](#).

To create a Topic object, an application can use the `MQTopic` constructor, as shown in the following example:

```
MQTopic t1 = new MQTopic("Sport/Football/Results");
```

This statement creates an `MQTopic` object with the default values for all its properties. The object represents a topic called Sport/Football/Results.

The `MQTopic` constructor can also accept a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the `MQTopic` object. The following statement contains an example of a topic URI:

```
MQTopic t2 = new MQTopic("topic://Sport/Tennis/Results?persistence=1&priority=0");
```

The `MQTopic` object created by this statement represents a topic called Sport/Tennis/Results, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see [“Uniform resource identifiers \(URIs\)” on page 841](#). For an alternative way of setting the properties of an `MQTopic` object, see [“Setting the properties of destinations” on page 837](#).

Setting the properties of destinations

An application can set the properties of a destination by calling the appropriate methods of the destination. The destination can either be an administered object or an object created dynamically at run time.

Consider the following code, for example:

```
MQQueue q1 = new MQQueue("Q1");
.
q1.setPersistence(WMQConstants.WMQ_PER_PER);
q1.setPriority(5);
```

This code creates an MQQueue object and then sets two properties of the object. The effect of setting these properties is that all messages sent to the destination are persistent and have a priority of 5.

An application can set the properties of MQTopic object in a similar way, as shown in the following example:

```
MQTopic t1 = new MQTopic("Sport/Football/Results");
.
t1.setPersistence(WMQConstants.WMQ_PER_NON);
t1.setPriority(0);
```

This code creates an MQTopic object and then sets two properties of the object. The effect of setting these properties is that all messages sent to the destination are nonpersistent and have a priority of 0.

For more information about the properties of a destination, and the methods used to set its properties, see [Properties of IBM WebSphere MQ classes for JMS objects](#).

Building a connection in a JMS application

To build a connection, a JMS application uses a ConnectionFactory object to create a Connection object and then starts the connection.

To create a Connection object, an application uses the createConnection() method of a ConnectionFactory object, as shown in the following example:

```
ConnectionFactory factory;
Connection connection;
.
.
.
connection = factory.createConnection();
```

When a JMS connection is created, the IBM WebSphere MQ classes for JMS creates a connection handle (Hconn) and starts a conversation with the queue manager.

The QueueConnectionFactory interface and the TopicConnectionFactory interface each inherits the createConnection() method from the ConnectionFactory interface. You can therefore use the createConnection() method to create a domain specific object, as shown in the following example:

```
QueueConnectionFactory qcf;
Connection connection;
.
.
.
connection = qcf.createConnection();
```

This fragment of code creates a QueueConnection object. An application can now perform a domain independent operation on this object, or an operation that is applicable only to the point-to-point domain. However, if the application attempts to perform an operation that is applicable only to the publish/subscribe domain, an IllegalStateException exception is thrown with the following message:

```
JMSMQ1112: Operation for a domain specific object was not valid.
Operation createProducer() is not valid for type com.ibm.mq.jms.MQTopic
```

This is because the connection was created from a domain specific connection factory.

Note: Note that the application process ID is used as the default user identity to be passed to the queue manager. If the application is running in client transport mode then this process ID must exist, with the relevant authorizations, on the server. If you want a different identity to be used, then use the `createConnection(username, password)` method.

The JMS specification states that a connection is created in the stopped state. Until a connection starts, a message consumer that is associated with the connection cannot receive any messages. To start a connection, an application uses the `start()` method of a `Connection` object, as shown in the following example:

```
connection.start();
```

Creating a session in a JMS application

To create a session, a JMS application uses the `createSession()` method of a `Connection` object.

The `createSession()` method has two parameters:

1. A parameter that specifies whether the session is transacted or not transacted
2. A parameter that specifies the acknowledgment mode for the session

For example, the following code creates a session that is not transacted and has an acknowledgment mode of `AUTO_ACKNOWLEDGE`:

```
Session session;  
  
boolean transacted = false;  
session = connection.createSession(transacted, Session.AUTO_ACKNOWLEDGE);
```

When a JMS session is created, the IBM WebSphere MQ classes for JMS creates a connection handle (Hconn) and starts a conversation with the queue manager.

A `Session` object, and any `MessageProducer` or `MessageConsumer` object created from it, cannot be used concurrently by different threads of a multithreaded application. The simplest way of ensuring that these objects are not used concurrently is to create a separate `Session` object for each thread.

Transacted sessions in JMS applications

JMS applications can run local transactions by first creating a transacted session. An application can commit or roll back a transaction.

JMS applications can run local transactions. A local transaction is a transaction that involves changes only to the resources of the queue manager to which the application is connected. To run local transactions, an application must first create a transacted session by calling the `createSession()` method of a `Connection` object, specifying as a parameter that the session is transacted. Subsequently, all messages sent and received within the session are grouped into a sequence of transactions. A transaction ends when the application commits or rolls back the messages it has sent and received since the transaction began.

To commit a transaction, an application calls the `commit()` method of the `Session` object. When a transaction is committed, all messages sent within the transaction become available for delivery to other applications, and all messages received within the transaction are acknowledged so that the messaging server does not attempt to deliver them to the application again. In the point-to-point domain, the messaging server also removes the received messages from their queues.

To roll back a transaction, an application calls the `rollback()` method of the `Session` object. When a transaction is rolled back, all messages sent within the transaction are discarded by the messaging server, and all messages received within the transaction become available for delivery again. In the point-to-point domain, the messages that were received are put back on their queues and become visible to other applications again.

A new transaction starts automatically when an application creates a transacted session or calls the `commit()` or `rollback()` method. Therefore, a transacted session always has an active transaction.

When an application closes a transacted session, an implicit rollback occurs. When an application closes a connection, an implicit rollback occurs for all the connection's transacted sessions.

If an application ends without closing a connection, an implicit rollback also occurs for all the connection's transacted sessions.

A transaction is wholly contained within a transacted session. A transaction cannot span sessions. This means that it is not possible for an application to send and receive messages in two or more transacted sessions and then commit or roll back all these actions as a single transaction.

Acknowledgment modes of JMS sessions

Every session that is not transacted has an acknowledgment mode that determines how messages received by the application are acknowledged. Three acknowledgment modes are available, and the choice of acknowledgment mode affects the design of the application.

If a session is not transacted, the way that messages received by the application are acknowledged is determined by the acknowledgment mode of the session. The three acknowledgment modes are described in the following paragraphs:

AUTO_ACKNOWLEDGE

The session automatically acknowledges each message received by the application.

If messages are delivered synchronously to the application, the session acknowledges receipt of a message every time a Receive call completes successfully. If messages are delivered asynchronously, the session acknowledges receipt of a message every time a call to the `onMessage()` method of a message listener completes successfully.

If the application receives a message successfully, but a failure prevents acknowledgment from occurring, the message becomes available for delivery again. The application must therefore be able to handle a message that is re-delivered.

DUPS_OK_ACKNOWLEDGE

The session acknowledges the messages received by the application at times it selects.

Using this acknowledgment mode reduces the amount of work the session must do, but a failure that prevents message acknowledgment might result in more than one message becoming available for delivery again. The application must therefore be able to handle messages that are re-delivered.

Restriction: In `AUTO_ACKNOWLEDGE` and `DUPS_OK_ACKNOWLEDGE` modes, JMS does not support an application throwing an unhandled exception in a message listener. This means that messages are always acknowledged when the message listener returns, regardless of whether it was processed successfully (provided any failures are non-fatal and do not prevent the application from continuing). If you require finer control of message acknowledgment, use the `CLIENT_ACKNOWLEDGE` or transacted modes, which give the application full control of the acknowledgment functions.

CLIENT_ACKNOWLEDGE

The application acknowledges the messages it receives by calling the `Acknowledge` method of the `Message` class.

The application can acknowledge the receipt of each message individually, or it can receive a batch of messages and call the `Acknowledge` method only for the last message it receives. When the `Acknowledge` method is called all messages received since the last time the method was called are acknowledged.

In conjunction with any of these acknowledgment modes, an application can stop and restart the delivery of messages in a session by calling the `Recover` method of the `Session` class. Messages received but previously unacknowledged are re-delivered. However, they might not be delivered in the same sequence in which they were previously delivered. In the meantime, higher priority messages might have arrived, and some of the original messages might have expired. In the point-to-point domain, some of the original messages might have been consumed by another application.

An application can determine whether a message is being re-delivered by examining the contents of the `JMSRedelivered` header field of the message. The application does this by calling the `getJMSRedelivered()` method of the `Message` class.

Creating destinations in a JMS application

Instead of retrieving destinations as administered objects from a Java Naming and Directory Interface (JNDI) namespace, a JMS application can use a session to create destinations dynamically at run time. An application can use a uniform resource identifier (URI) to identify a WebSphere MQ queue or a topic and, optionally, to specify one or more properties of a Queue or Topic object.

Using a session to create Queue objects

To create a Queue object, an application can use the `createQueue()` method of a Session object, as shown in the following example:

```
Session session;  
.  
Queue q1 = session.createQueue("Q1");
```

This code creates a Queue object with the default values for all its properties. The object represents a WebSphere MQ queue called Q1 that belongs to the local queue manager. This queue can be a local queue, an alias queue, or a remote queue definition.

The `createQueue()` method also accepts a queue URI as a parameter. A queue URI is a string that specifies the name of a WebSphere MQ queue and, optionally, the name of the queue manager that owns the queue and one or more properties of the Queue object. The following statement contains an example of a queue URI:

```
Queue q2 = session.createQueue("queue://QM2/Q2?persistence=2&priority=5");
```

The Queue object created by this statement represents a WebSphere MQ queue called Q2 that is owned by a queue manager called QM2, and all messages sent to this destination are persistent and have a priority of 5. The queue manager identified in this way can be the local queue manager or a remote queue manager. If it is a remote queue manager, WebSphere MQ must be configured so that, when the application sends a message to this destination, WebSphere MQ can route the message from the local queue manager to queue manager QM2. For more information about URIs, see [“Uniform resource identifiers \(URIs\)” on page 841](#).

Note that the parameter on the `createQueue()` method contains provider specific information. Therefore, using the `createQueue()` method to create a Queue object, instead of retrieving a Queue object as an administered object from a JNDI namespace, might make your application less portable.

An application can create a TemporaryQueue object by using the `createTemporaryQueue()` method of a Session object, as shown in the following example:

```
TemporaryQueue q3 = session.createTemporaryQueue();
```

Although a session is used to create a temporary queue, the scope of a temporary queue is the connection that was used to create the session. Any of the connection's sessions can create message producers and message consumers for the temporary queue. The temporary queue remains until the connection ends or the application explicitly deletes the temporary queue by using the `TemporaryQueue.delete()` method, whichever is the sooner.

When an application creates a temporary queue, WebSphere MQ classes for JMS creates a dynamic queue in the queue manager to which the application is connected. The `TEMPMODEL` property of the connection factory specifies the name of the model queue that is used to create the dynamic queue, and the `TEMPQPREFIX` property of the connection factory specifies the prefix that is used to form the name of the dynamic queue.

Using a session to create Topic objects

To create a Topic object, an application can use the `createTopic()` method of a Session object, as shown in the following example:


```
Session session;
Topic t1 = session.createTopic("Sport/Football/Results");
```

This code creates an Topic object with the default values for all its properties. The object represents a topic called Sport/Football/Results.

The createTopic() method also accepts a topic URI as a parameter. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the Topic object. The following code contains an example of a topic URI:

```
String uri = "topic://Sport/Tennis/Results?persistence=1&priority=0";
Topic t2 = session.createTopic(uri);
```

The Topic object created by this code represents a topic called Sport/Tennis/Results, and all messages sent to this destination are nonpersistent and have a priority of 0. For more information about topic URIs, see [“Uniform resource identifiers \(URIs\)”](#) on page 841.

Note that the parameter on the createTopic() method contains provider specific information. Therefore, using the createTopic() method to create a Topic object, instead of retrieving a Topic object as an administered object from a JNDI namespace, might make your application less portable.

An application can create a TemporaryTopic object by using the createTemporaryTopic() method of a Session object, as shown in the following example:

```
TemporaryTopic t3 = session.createTemporaryTopic();
```

Although a session is used to create a temporary topic, the scope of a temporary topic is the connection that was used to create the session. Any of the connection's sessions can create message producers and message consumers for the temporary topic. The temporary topic remains until the connection ends or the application explicitly deletes the temporary topic by using the TemporaryTopic.delete() method, whichever is the sooner.

When an application creates a temporary topic, WebSphere MQ classes for JMS creates a topic with a name that commences with the characters TEMP/tempTopicPrefix, where tempTopicPrefix is the value of the TEMPTOPICPREFIX property of the connection factory.

Uniform resource identifiers (URIs)

A queue URI is a string that specifies the name of a WebSphere MQ queue and, optionally, the name of the queue manager that owns the queue and one or more properties of the Queue object created by the application. A topic URI is a string that specifies the name of a topic and, optionally, one or more properties of the Topic object created by the application.

A queue URI has the following format:

```
queue://[qMgrName]/qName[?propertyName1=propertyValue1
                        &propertyName2=propertyValue2
                        &...]
```

A topic URI has the following format:

```
topic://topicName[?propertyName1=propertyValue1
                  &propertyName2=propertyValue2
                  &...]
```

The variables in these formats have the following meanings:

qMgrName

The name of the queue manager that owns the queue identified by the URI.

The queue manager can be the local queue manager or a remote queue manager. If it is a remote queue manager, WebSphere MQ must be configured so that, when an application sends a message to the queue, WebSphere MQ can route the message from the local queue manager to the remote queue manager.

If no name is specified, the local queue manager is assumed.

qName

The name of the WebSphere MQ queue.

The queue can be a local queue, an alias queue, or a remote queue definition.

For the rules for creating queue names, see [Rules for naming IBM WebSphere MQ objects](#).

topicName

The name of the topic.

For the rules for creating topic names, see [Rules for naming IBM WebSphere MQ objects](#). Avoid the use of the wildcard characters +, #, *, and ? in topic names. Topic names containing these characters can cause unexpected results when you subscribe to them. See [Using topic strings](#).

propertyName1, propertyName2, ...

The names of the properties of the Queue or Topic object created by the application. [Table 123 on page 842](#) lists the valid property names that can be used in a URI.

If no properties are specified, the Queue or Topic object has the default values for all its properties.

propertyValue1, propertyValue2, ...

The values of the properties of the Queue or Topic object created by the application. [Table 123 on page 842](#) lists the valid property values that can be used in a URI.

Brackets ([]) denotes an optional component, and the ellipsis (...) means that the list of property name-value pairs, if present, can contain one or more name-value pairs.

[Table 123 on page 842](#) lists the valid property names and valid values that can be used in queue and topic URIs. Although the WebSphere MQ JMS administration tool uses symbolic constants for the values of properties, URIs cannot contain symbolic constants.

Table 123. Property names and valid values for use in queue and topic URIs		
Property name	Description	Valid values
CCSID	How the character data in the body of a message is represented when WebSphere MQ classes for JMS forwards the message to the destination	<ul style="list-style-type: none">Any coded character set identifier supported by WebSphere MQ.
encoding	How the numeric data in the body of a message is represented when WebSphere MQ classes for JMS forwards the message to the destination	<ul style="list-style-type: none">Any valid value for the <i>Encoding</i> field in a WebSphere MQ message descriptor.
expiry	The time to live for messages sent to the destination	<ul style="list-style-type: none">-2 - As specified on the send() call or, if not specified on the send() call, the default time to live of the message producer.0 - A message sent to the destination never expires.A positive integer specifying the time to live in milliseconds.

Table 123. Property names and valid values for use in queue and topic URIs (continued)

Property name	Description	Valid values
multicast	The multicast setting for a topic when using a real-time connection to a broker	<p>The following list contains the valid values. Associated with each value is the corresponding value of the MULTICAST property as used in the WebSphere MQ JMS administration tool. For a description of the MULTICAST property and its valid values, see Properties of IBM WebSphere MQ classes for JMS objects.</p> <ul style="list-style-type: none"> • -1 - ASCF • 0 - DISABLED • 3 - NOTR • 5 - RELIABLE • 7 - ENABLED
persistence	The persistence of messages sent to the destination	<ul style="list-style-type: none"> • -2 - As specified on the send() call or, if not specified on the send() call, the default persistence of the message producer. • -1 - As specified by the <i>DefPersistence</i> attribute of the WebSphere MQ queue or topic. • 1 - Nonpersistent. • 2 - Persistent. • 3 - Equivalent to the value HIGH for the PERSISTENCE property as used in the WebSphere MQ JMS administration tool. For an explanation of this value, see “JMS persistent messages” on page 866.
priority	The priority of messages sent to the destination	<ul style="list-style-type: none"> • -2 - As specified on the send() call or, if not specified on the send() call, the default priority of the message producer. • -1 - As specified by the <i>DefPriority</i> attribute of the WebSphere MQ queue or topic. • An integer in the range 0-9 specifying the priority of messages sent to the destination.
targetClient	Whether messages sent to the destination contain an MQRFH2 header	<ul style="list-style-type: none"> • 0 - Messages contain an MQRFH2 header. • 1 - Messages do not contain an MQRFH2 header.

For example, the following URI identifies a WebSphere MQ queue called Q1 that is owned by the local queue manager. A Queue object created using this URI has the default values for all its properties.

```
queue:///Q1
```

The following URI identifies a WebSphere MQ queue called Q2 that is owned by a queue manager called QM2. All messages sent to this destination have a priority of 6. The remaining properties of the Queue object created using this URI have their default values.

```
queue://QM2/Q2?priority=6
```

The following URI identifies a topic called Sport/Athletics/Results. All messages sent to this destination are nonpersistent and have a priority of 0. The remaining properties of the Topic object created using this URI have their default values.

```
topic://Sport/Athletics/Results?persistence=1&priority=0
```

Sending messages in a JMS application

Before a JMS application can send messages to a destination, it must first create a `MessageProducer` object for the destination. To send a message to the destination, the application creates a `Message` object and then calls the `send()` method of the `MessageProducer` object.

An application uses a `MessageProducer` object to send messages. An application normally creates a `MessageProducer` object for a specific destination, which can be a queue or a topic, so that all messages sent using the message producer are sent to the same destination. Therefore, before an application can create a `MessageProducer` object, it must first create a `Queue` or `Topic` object. For information about how to create a `Queue` or `Topic` object, see the following topics:

- [“Using JNDI to retrieve administered objects in a JMS application” on page 827](#)
- [“Using the IBM JMS extensions” on page 828](#)
- [“Using the WebSphere MQ JMS extensions” on page 834](#)
- [“Creating destinations in a JMS application” on page 840](#)

To create a `MessageProducer` object, an application uses the `createProducer()` method of a `Session` object, as shown in the following example:

```
MessageProducer producer = session.createProducer(destination);
```

The parameter `destination` is a `Queue` or `Topic` object that the application has created previously.

Before an application can send a message, it must create a `Message` object. The body of a message contains the application data, and JMS defines five types of message body:

- Bytes
- Map
- Object
- Stream
- Text

Each type of message body has its own JMS interface, which is a sub-interface of the `Message` interface, and a method in the `Session` interface for creating a message with that type of body. For example, the interface for a text message is called `TextMessage`, and an application uses the `createTextMessage()` method of a `Session` object to create a text message, as shown in the following statement:

```
TextMessage outMessage = session.createTextMessage(outString);
```

For more information about messages and message bodies, see [“JMS messages” on page 772](#).

To send a message, an application uses the `send()` method of a `MessageProducer` object, as shown in the following example:

```
producer.send(outMessage);
```

An application can use the `send()` method to send messages in either messaging domain. The nature of the destination determines which messaging domain is used. However, `TopicPublisher`, the sub-interface of `MessageProducer` that is specific to the publish/subscribe domain, also has a `publish()` method, which can be used instead of the `send()` method. The two methods are functionally the same.

An application can create a `MessageProducer` object with no specified destination. In this case, the application must specify the destination when calling the `send()` method.

If an application sends a message within a transaction, the message is not delivered to its destination until the transaction is committed. This means that an application cannot send a message and receive a reply to the message within the same transaction.

A destination can be configured so that when an application sends messages to it, WebSphere MQ classes for JMS forwards the message and returns control back to the application without determining whether the queue manager has received the message safely. This is sometimes referred to as *asynchronous put*. For more information, see [“Putting messages asynchronously in IBM WebSphere MQ classes for JMS” on page 880](#).

Receiving messages in a JMS application

An application uses a message consumer to receive messages. A durable topic subscriber is a message consumer that receives all messages sent to a destination, including those sent while the consumer is inactive. An application can select which messages it wants to receive by using a message selector, and can receive messages asynchronously by using a message listener.

An application uses a `MessageConsumer` object to receive messages. An application creates a `MessageConsumer` object for a specific destination, which can be a queue or a topic, so that all messages received using the message consumer are received from the same destination. Therefore, before an application can create a `MessageConsumer` object, it must first create a `Queue` or `Topic` object. For information about how to create a `Queue` or `Topic` object, see the following topics:

- [“Using JNDI to retrieve administered objects in a JMS application” on page 827](#)
- [“Using the IBM JMS extensions” on page 828](#)
- [“Using the WebSphere MQ JMS extensions” on page 834](#)
- [“Creating destinations in a JMS application” on page 840](#)

To create a `MessageConsumer` object, an application uses the `createConsumer()` method of a `Session` object, as shown in the following example:

```
MessageConsumer consumer = session.createConsumer(destination);
```

The parameter `destination` is a `Queue` or `Topic` object that the application has created previously.

The application then uses the `receive()` method of the `MessageConsumer` object to receive a message from the destination, as shown in the following example:

```
Message inMessage = consumer.receive(1000);
```

The parameter on the `receive()` call specifies how long in milliseconds the method waits for a suitable message to arrive if no message is available immediately. If you omit this parameter, the call blocks indefinitely until a suitable message arrives. If you do not want the application to wait for a message, use the `receiveNoWait()` method instead.

The `receive()` method returns a message of a specific type. For example, when an application receives a text message, the object returned by the `receive()` call is a `TextMessage` object.

However, the declared type of object returned by a `receive()` call is a `Message` object. Therefore, in order to extract the data from the body of a message that has just been received, the application must cast from the `Message` class to the more specific subclass, such as `TextMessage`. If the type of the message is not known, the application can use the `instanceof` operator to determine the type. It is always

good practice for an application to determine the type of a message before casting so that errors can be handled gracefully.

The following code uses the `instanceof` operator and shows how to extract the data from the body of a text message:

```
if (inMessage instanceof TextMessage) {
    String replyString = ((TextMessage) inMessage).getText();
    .
    .
} else {
    // Print error message if Message was not a TextMessage.
    System.out.println("Reply message was not a TextMessage");
}
```

If an application sends a message within a transaction, the message is not delivered to its destination until the transaction is committed. This means that an application cannot send a message and receive a reply to the message within the same transaction.

If a message consumer receives messages from a destination that is configured for read ahead, any nonpersistent messages that are in the read ahead buffer when the application ends are discarded.

In the publish/subscribe domain, JMS identifies two types of message consumer, nondurable topic subscriber and durable topic subscriber, which are described in the following two sections.

Nondurable topic subscribers

A nondurable topic subscriber receives only those messages that are published while the subscriber is active. A nondurable subscription starts when an application creates a nondurable topic subscriber and ends when the application closes the subscriber, or when the subscriber falls out of scope. As an extension in WebSphere MQ classes for JMS, a nondurable topic subscriber also receives retained publications, but not when using a real-time connection to a broker.

To create a nondurable topic subscriber, an application can use the domain independent `createConsumer()` method, specifying a `Topic` object as the destination. Alternatively, an application can use the domain specific `createSubscriber()` method, as shown in the following example:

```
TopicSubscriber subscriber = session.createSubscriber(topic);
```

The parameter `topic` is a `Topic` object that the application has created previously.

Durable topic subscribers

Restriction: An application cannot create durable topic subscribers when using a real-time connection to a broker.

A durable topic subscriber receives all messages that are published during the life of a durable subscription. These messages include all those that are published while the subscriber is not active. As an extension in WebSphere MQ classes for JMS, a durable topic subscriber also receives retained publications.

To create a durable topic subscriber, an application uses the `createDurableSubscriber()` method of a `Session` object, as shown in the following example:

```
TopicSubscriber subscriber = session.createDurableSubscriber(topic, "D_SUB_000001");
```

On the `createDurableSubscriber()` call, the first parameter is a `Topic` object that the application has created previously, and the second parameter is a name that is used to identify the durable subscription.

The session used to create a durable topic subscriber must have an associated client identifier. The client identifier associated with a session is the same as the client identifier for the connection that is used to create the session. The client identifier can be specified by setting the `CLIENTID` property of the `ConnectionFactory` object. Alternatively, an application can specify the client identifier by calling the `setClientID()` method of the `Connection` object.

The name that is used to identify a durable subscription must be unique only within the client identifier, and therefore the client identifier forms part of the full, unique identifier of a durable subscription. To continue using a durable subscription that was created previously, an application must create a durable topic subscriber using a session with the same client identifier as that associated with the durable subscription, and using the same subscription name.

A durable subscription starts when an application creates a durable topic subscriber using a client identifier and subscription name for which no durable subscription currently exists. However, a durable subscription does not end when the application closes the durable topic subscriber. To end a durable subscription, an application must call the `unsubscribe()` method of a `Session` object that has the same client identifier as that associated with the durable subscription. The parameter on the `unsubscribe()` call is the subscription name, as shown in the following example:

```
session.unsubscribe("D_SUB_000001");
```

The scope of a durable subscription is a queue manager. If a durable subscription exists on one queue manager, and an application connected to another queue manager creates a durable subscription with the same client identifier and subscription name, the two durable subscriptions are completely independent.

Message selectors

An application can specify that only those messages that satisfy certain criteria are returned by successive `receive()` calls. When creating a `MessageConsumer` object, the application can specify a Structured Query Language (SQL) expression that determines which messages are retrieved. This SQL expression is called a *message selector*. The message selector can contain the names of JMS message header fields and message properties. For information about how to construct a message selector, see [“Message selectors in JMS” on page 773](#).

The following example shows how an application can select messages based on a user defined property called `myProp`:

```
MessageConsumer consumer;  
.  
consumer = session.createConsumer(destination, "myProp = 'blue'");
```

The JMS specification does not allow an application to change the message selector of a message consumer. After an application creates a message consumer with a message selector, the message selector remains for the life of that consumer. If an application requires more than one message selector, the application must create a message consumer for each message selector.

Note that, when an application is connected to a Version 7 queue manager, the `MSGSELECTION` property of the connection factory has no effect. To optimize performance, all message selection is done by the queue manager.

Suppressing local publications

An application can create a message consumer that ignores publications published on the consumer's own connection. The application does this by setting the third parameter on a `createConsumer()` call to `true`, as shown in the following example:

```
MessageConsumer consumer = session.createConsumer(topic, null, true);
```

On a `createDurableSubscriber()` call, the application does this by setting the fourth parameter to `true`, as shown in the following example

```
String selector = "company = 'IBM'";  
TopicSubscriber subscriber = session.createDurableSubscriber(topic, "D_SUB_000001",  
                                                             selector, true);
```

Asynchronous delivery of messages

An application can receive messages asynchronously by registering a message listener with a message consumer. The message listener has a method called `onMessage`, which is called asynchronously when a suitable message is available and whose purpose is to process the message. The following code illustrates the mechanism:

```
import javax.jms.*;

public class MyClass implements MessageListener
{
    // The method that is called asynchronously when a suitable message is available
    public void onMessage(Message message)
    {
        System.out.println("Message is "+message);

        // The code to process the message
        .
        .
        .
    }
}

// Main program (possibly in another class)
// Creating the message listener
MyClass listener = new MyClass();

// Registering the message listener with a message consumer
consumer.setMessageListener(listener);

// The main program now continues with other processing
```

An application can use a session either for receiving messages synchronously using `receive()` calls, or for receiving messages asynchronously using message listeners, but not for both. If an application needs to receive messages synchronously and asynchronously, it must create separate sessions.

Once a session is set up to receive messages asynchronously, the following methods cannot be called on that session or on objects created from that session:

- `MessageConsumer.receive()`
- `MessageConsumer.receive(long)`
- `MessageConsumer.receiveNoWait()`
- `Session.acknowledge()`
- `MessageProducer.send(Destination, Message)`
- `MessageProducer.send(Destination, Message, int, int, long)`
- `MessageProducer.send(Message)`
- `MessageProducer.send(Message, int, int, long)`
- `Session.commit()`
- `Session.createBrowser(Queue)`
- `Session.createBrowser(Queue, String)`
- `Session.createBytesMessage()`
- `Session.createConsumer(Destination)`
- `Session.createConsumer(Destination, String, boolean)`
- `Session.createDurableSubscriber(Topic, String)`
- `Session.createDurableSubscriber(Topic, String, String, boolean)`
- `Session.createMapMessage()`
- `Session.createMessage()`
- `Session.createObjectMessage()`

- Session.createObjectMessage(Serializable)
- Session.createProducer(Destination)
- Session.createQueue(String)
- Session.createStreamMessage()
- Session.createTemporaryQueue()
- Session.createTemporaryTopic()
- Session.createTextMessage()
- Session.createTextMessage(String)
- Session.createTopic()
- Session.getAcknowledgeMode()
- Session.getMessageListener()
- Session.getTransacted()
- Session.rollback()
- Session.unsubscribe(String)

If any of these methods are called, a `JMSEException` containing the message:

```
JMSCC0033: A synchronous method call is not permitted when a session is being used
asynchronously: 'method name'
```

is thrown.

Receiving poison messages

An application can receive a message that cannot be processed. There can be several reasons why the message cannot be processed, for example the message might have an incorrect format. Such messages are described as poison messages and require special handling to prevent the message being recursively processed.

For details on how to handle poison messages, see [“Handling poison messages in IBM WebSphere MQ classes for JMS” on page 850](#).

V7.5.0.8 Retrieval of subscription user data

If the messages that an IBM WebSphere MQ classes for JMS application is consuming from a queue are put by an administratively defined durable subscription, the application needs to access the user data information that is associated with the subscription. This information is added to the message as a property.

From Version 7.5.0, Fix Pack 8, when a message is consumed from a queue that contains an RFH2 header with the MQPS folder, the value that is associated with the `Sud` key, if it exists, is added as a String property to the JMS Message object returned to the IBM WebSphere MQ classes for JMS application. To enable the retrieval of this property from the message, the constant `JMS_IBM_SUBSCRIPTION_USER_DATA` in the `JmsConstants` interface can be used with the method `javax.jms.Message.getStringProperty(java.lang.String)` to get the subscription user data.

In the following example, an administrative durable subscription is defined by using the MQSC command **DEFINE SUB:**

```
DEFINE SUB('MY.SUBSCRIPTION') TOPICSTR('PUBLIC') DEST('MY.SUBSCRIPTION.Q')
USERDATA('Administrative durable subscription to put message to the queue MY.SUBSCRIPTION.Q')
```

Copies of messages that are published to the topic string `PUBLIC` are put to the queue, `MY.SUBSCRIPTION.Q`. The user data that is associated with the durable subscription is then added as a property to the message, which is stored in the MQPS folder of the RFH2 header with the key `Sud`.

The IBM WebSphere MQ classes for JMS application can call:

```
javax.jms.Message.getStringProperty(JmsConstants.JMS_IBM_SUBSCRIPTION_USER_DATA);
```

The following String is then returned:

```
Administrative durable subscription to put message to the queue MY.SUBSCRIPTION.Q
```

Related concepts

[“The MQRFH2 header and JMS” on page 776](#)

Related tasks

[Defining an administrative subscription](#)

Related reference

[DEFINE SUB](#)

[Interface JmsConstants](#)

Closing down a WebSphere MQ classes for JMS application

It is important for a WebSphere MQ classes for JMS application to close certain JMS objects explicitly before stopping. Finalizers might not be called, so do not rely on them to free resources. Do not allow an application to terminate with compressed trace active.

Garbage collection alone cannot release all WebSphere MQ classes for JMS and WebSphere MQ resources in a timely manner, especially if an application creates many short lived JMS objects at the session level or lower. It is therefore important for an application to close a Connection, Session, MessageConsumer, or MessageProducer object when it is no longer required.

If an application ends without closing a Connection, an implicit rollback occurs for all the connection's transacted sessions. To ensure any changes made by the application are committed, close the Connection explicitly before closing the application.

Do not use finalizers in an application to close JMS objects. Because finalizers might not be called, resources might not be freed. When a Connection is closed it closes all the Sessions that were created from it. Similarly, the MessageConsumers and MessageProducers created from a Session are closed when the Session is closed. However, consider closing Sessions, MessageConsumers, and MessageProducers explicitly to ensure resources are freed in a timely manner.

If trace compression is activated, System.Halt() shutdowns and abnormal, uncontrolled JVM terminations are likely to result in a corrupt trace file. Where possible, turn off the trace facility when you have collected the trace information you need. If you are tracing an application up to an abnormal end, use uncompressed trace output.

Handling poison messages in IBM WebSphere MQ classes for JMS

A poison message is one which cannot be processed by a receiving MDB application. If a poison message is encountered, the JMS MessageConsumer and ConnectionConsumer objects can requeue it according to two queue properties, BOQNAME, and BOTHRESH.

Sometimes, a badly formatted message arrives on a queue. In this context, badly formatted means that the receiving application cannot process the message correctly. Such a message can cause the receiving application to fail and to back out this badly formatted message. The message can then be repeatedly delivered to the input queue and repeatedly backed out by the application. These messages are known as *poison messages*. The JMS MessageConsumer object detects poison messages and reroutes them to an alternative destination.

The IBM WebSphere MQ queue manager keeps a record of the number of times that each message has been backed out. When this number reaches a configurable threshold value, the message consumer requeues the message to a named backout queue. If this requeuing fails for any reason, the message is removed from the input queue and either requeued to the dead-letter queue, or discarded. See [“Removing messages from the queue in ASF” on page 888](#) for more details.

There is a difference between the way in which poison messages are requeued by MessageConsumers and ConnectionConsumers. ConnectionConsumers are able to requeue poison messages without affecting message delivery. The requeue process takes place outside of any unit of work associated with actual message delivery to application code. This is possible because of the multi-threaded nature of ConnectionConsumer operation.

MessageConsumers, however, are single threaded below the Session level, and any requeuing of poison messages takes place within the current unit of work. This does not affect the operation of the application, however when poison messages are requeued under a transacted or Client_acknowledge Session, the requeue action itself is not committed until the current unit of work is committed by the application code or, if appropriate, the application container code.

JMS ConnectionConsumer objects handle poison messages in the same way and using the same queue properties. If multiple connection consumers are monitoring the same queue, it is possible that the poison message may be delivered to an application more times than the threshold value before the requeue occurs. This behavior is due to the way individual connection consumers monitor queues and requeue poison messages.

The threshold value and the name of the back out queue are attributes of an IBM WebSphere MQ queue. The names of the attributes are BackoutThreshold and BackoutRequeueQName. The queue they apply to is as follows:

- For point-to-point messaging, this is the underlying local queue. This is important when message consumers and connection consumers use queue aliases.
- For publish/subscribe messaging in IBM WebSphere MQ messaging provider normal mode, the Topic's managed queue is created from the model queue.
- For publish/subscribe messaging in IBM WebSphere MQ messaging provider migration mode, it is the CCSUB queue defined on the TopicConnectionFactory object, or the CCDSUB queue defined on the Topic object.

IBM WebSphere MQ classes for JMS queries the BackoutThreshold and BackoutRequeueQName of the queue. You must therefore grant inquire access on the queue to the user running the application.

V7.5.0.9 If the target queue is a cluster queue, then the authorities needed depend on the version of the IBM WebSphere MQ classes for JMS being used:

- When using the IBM WebSphere MQ classes for JMS for Version 7.5.0, Fix Pack 9 plus an interim fix for APAR IT26482, inquire access is required.
- For all other versions, grant inquire, browse and get access.

To set the BackoutThreshold and BackoutRequeueQName attributes, issue the following MQSC command:

```
ALTER QLOCAL(your.queue.name) BOTHRESH(threshold value) BOQNAME(your.backout.queue.name)
```

If the BackoutThreshold attribute is set to a value other than zero, to avoid unexpected behavior set the BackoutRequeueQName attribute to a valid queue name.

For publish/subscribe messaging, if your system creates a dynamic queue for each subscription, these attribute values are obtained from the IBM WebSphere MQ classes for JMS model queue, SYSTEM.JMS.MODEL.QUEUE. To alter these settings, use:

```
ALTER QMODEL(SYSTEM.JMS.MODEL.QUEUE) BOTHRESH(threshold value) BOQNAME(your.backout.queue.name)
```

If the backout threshold value is zero, poison message handling is disabled, and poison messages remain on the input queue. Otherwise, when the backout count reaches the threshold value, the message is sent to the named backout queue. If the backout count reaches the threshold value, but the message cannot go to the backout queue, the message is sent to the dead-letter queue or it is discarded. This situation occurs if the backout queue is not defined, or if the MessageConsumer object cannot send the message to the backout queue. See [“Removing messages from the queue in ASF” on page 888](#) for further details.

When a message is requeued to the backout requeue queue, some of the field values in the Message Descriptor (MQMD) of the message change. See [MQMD - Message Descriptor](#) for details about the format of the MQMD.

The following MQMD fields change value when the message goes onto the backout queue.

- PutDate is updated to the date it goes onto the backout requeue queue.
- PutTime is updated to the time it goes onto the backout requeue queue.
- Backout count is reset to zero.
- Expiry of the message is updated to reflect the remaining expiry at the time the original message was received by the JMS application.

The values in the following fields stay the same when the message goes onto the backout queue:

- StructId
- Version
- Report
- MessageType
- Feedback
- Encoding
- CodedCharSetId
- MsgId
- CorrelId
- ReplyToQ
- ReplyToQMgr
- Format
- Persistence
- Priority

Exceptions in IBM WebSphere MQ classes for JMS

A IBM WebSphere MQ classes for JMS application must be able to handle exceptions that are thrown by a JMS API calls or delivered to an exception handler.

IBM WebSphere MQ classes for JMS reports runtime problems by throwing exceptions. JMSEException is the root class for exceptions thrown by JMS methods, and catching JMSEException exceptions provides a generic way of handling all JMS related exceptions.

Every JMSEException exception encapsulates the following information:

- A provider specific exception message, which an application obtains by calling the Throwable.getMessage() method.
- A provider specific error code, which an application obtains by calling the JMSEException.getErrorCode() method.
- A linked exception. An exception thrown by a JMS API call is often the result of a lower level problem, which is reported by another exception linked to this exception. An application obtains a linked exception by calling the JMSEException.getLinkedException() or the Throwable.getCause() method.

Most exceptions thrown by IBM WebSphere MQ classes for JMS are instances of subclasses of JMSEException. These subclasses implement the com.ibm.msg.client.jms.JmsExceptionDetail interface, which provides the following additional information:

- An explanation of the exception message, which an application obtains by calling the JmsExceptionDetail.getExplanation() method.
- A recommended user response to the exception, which an application obtains by calling the JmsExceptionDetail.getUserAction() method.

- The keys for the message inserts in the exception message. An application obtains an iterator for all the keys by calling the `JmsExceptionDetail.getKeys()` method.
- The message inserts in the exception message. For example, a message insert might be the name of the queue that caused the exception, and it might be useful for an application to be able to access that name. An application obtains the message insert corresponding to a specified key by calling the `JmsExceptionDetail.getValue()` method.

All the methods in the `JmsExceptionDetail` interface might return null if no details are available.

For example, if an application tries to create a message producer for an IBM WebSphere MQ queue that does not exist, an exception is thrown with the following information:

```
Message : JMSWMQ2008: Failed to open MQ queue 'Q_test'.
Class : class com.ibm.msg.client.jms.DetailedInvalidDestinationException
Error Code : JMSWMQ2008
Explanation : JMS attempted to perform an MQOPEN, but WebSphere MQ reported an
              error.
User Action : Use the linked exception to determine the cause of this error. Check
              that the specified queue and queue manager are defined correctly.
```

The exception thrown, `com.ibm.msg.client.jms.DetailedInvalidDestinationException`, is a subclass of `javax.jms.InvalidDestinationException` and implements the `com.ibm.msg.client.jms.JmsExceptionDetail` interface.

Linked exceptions

A linked exception provides further information about a runtime problem. Therefore, for each `JMSEException` exception that is thrown, an application should check the linked exception. The linked exception itself might have another linked exception, and so the linked exceptions form a chain leading back to the original underlying problem. A linked exception is implemented by using the chained exception mechanism of the `java.lang.Throwable` class, and an application obtains a linked exception by calling the `Throwable.getCause()` method. For a `JMSEException` exception, the `getLinkedException()` method actually delegates to the `Throwable.getCause()` method.

For example, if an application specifies an incorrect port number when connecting to a queue manager, the exceptions form the following chain:

```
com.ibm.msg.client.jms.DetailIllegalStateException
|
+--->com.ibm.mq.MQException
      |
      +--->com.ibm.mq.jmqi.JmqiException
            |
            +--->java.net.ConnectionException
```

Typically, each exception in a chain is thrown from a different layer in the code. For example, the exceptions in the preceding chain are thrown by the following layers:

- The first exception, an instance of a subclass of `JMSEException`, is thrown by the common layer in IBM WebSphere MQ classes for JMS.
- The next exception, an instance of `com.ibm.mq.MQException`, is thrown by the IBM WebSphere MQ messaging provider.
- The next exception, an instance of `com.ibm.mq.jmqi.JmqiException`, is thrown by the common Java interface to the MQI.
- The final exception, an instance of `java.net.ConnectionException`, is thrown by the Java class library.

For more information about the layered architecture of IBM WebSphere MQ classes for JMS, see [“IBM WebSphere MQ classes for JMS architecture” on page 763](#).

Using code similar to the following code, an application can iterate through this chain to extract all the appropriate information:

```
import com.ibm.msg.client.jms.JmsExceptionDetail;
import com.ibm.mq.MQException;
import com.ibm.mq.jmqi.JmqiException;
```

```

import javax.jms.JMSException;
.
.
.
catch (JMSException je) {
    System.err.println("Caught JMSException");

    // Check for linked exceptions in JMSException
    Throwable t = je;
    while (t != null) {
        // Write out the message that is applicable to all exceptions
        System.err.println("Exception Msg: " + t.getMessage());
        // Write out the exception stack trace
        t.printStackTrace(System.err);

        // Add on specific information depending on the type of exception
        if (t instanceof JMSException) {
            JMSException je1 = (JMSException) t;
            System.err.println("JMS Error code: " + je1.getErrorCode());

            if (t instanceof JmsExceptionDetail){
                JmsExceptionDetail jed = (JmsExceptionDetail)je1;
                System.err.println("JMS Explanation: " + jed.getExplanation());
                System.err.println("JMS Explanation: " + jed.getUserAction());
            }
        } else if (t instanceof MQException) {
            MQException mqe = (MQException) t;
            System.err.println("WMQ Completion code: " + mqe.getCompCode());
            System.err.println("WMQ Reason code: " + mqe.getReason());
        } else if (t instanceof JmqiException){
            JmqiException jmqie = (JmqiException)t;
            System.err.println("WMQ Log Message: " + jmqie.getWmqLogMessage());
            System.err.println("WMQ Explanation: " + jmqie.getWmqMsgExplanation());
            System.err.println("WMQ Msg Summary: " + jmqie.getWmqMsgSummary());
            System.err.println("WMQ Msg User Response: "
                + jmqie.getWmqMsgUserResponse());
            System.err.println("WMQ Msg Severity: " + jmqie.getWmqMsgSeverity());
        }

        // Get the next cause
        t = t.getCause();
    }
}

```

Note that an application should always check the type of each exception in a chain because the type of exception can vary and exceptions of different types encapsulate different information.

Obtaining IBM WebSphere MQ specific information about a problem

Instances of `com.ibm.mq.MQException` and `com.ibm.mq.jmqi.JmqiException` encapsulate IBM WebSphere MQ specific information about a problem.

An `MQException` exception encapsulates the following information:

- A completion code, which an application obtains by calling the `getCompCode()` method
- A reason code, which an application obtains by calling the `getReason()` method

A `JmqiException` exception also encapsulates a completion code and a reason code. Additionally, however, a `JmqiException` exception encapsulates the information in an `AMQnnnn` or `CSQnnnn` message, if one is associated with the exception. By calling the appropriate methods of the exception, an application can obtain the various components of this message, such as the severity, explanation, and user response.

For examples of how to use of the methods mentioned in this section, see the sample code in [“Linked exceptions”](#) on page 853.

Upgrading from previous versions of IBM WebSphere MQ classes for JMS

Compared to previous versions of IBM WebSphere MQ classes for JMS, most error codes and exception messages have changed in Version 7. The reason for these changes is that IBM WebSphere MQ classes for JMS now has a layered architecture and exceptions are thrown from different layers in the code.

For example, if an application tries to connect to a queue manager that does not exist, a previous version of IBM WebSphere MQ classes for JMS threw a `JMSEException` exception with the following information:

```
MQJMS2005: Failed to create MQQueueManager for 'localhost:QM_test'.
```

This exception contained a linked `MQException` exception with the following information:

```
MQJE001: Completion Code 2, Reason 2058
```

By comparison in the same circumstances, Version 7 of IBM WebSphere MQ classes for JMS throws a `JMSEException` exception with the following information:

```
Message : JMSWMQ0018: Failed to connect to queue manager 'QM_test' with
           connection mode 'Client' and host name 'localhost'.
Class : class com.ibm.msg.client.jms.DetailedJMSEException
Error Code : JMSWMQ0018
Explanation : null
User Action : 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.
```

This exception contains a linked `MQException` exception with the following information:

```
Message : JMSCMQ0001: WebSphere MQ call failed with compcode '2' ('MQCC_FAILED')
           reason '2058' ('MQRC_Q_MGR_NAME_ERROR').
Class : class com.ibm.mq.MQException
Completion Code : 2
Reason Code : 2058
```

If your application parses or tests exception messages returned by the `Throwable.getMessage()` method, or error codes returned by the `JMSEException.getErrorCode()` method, and you are upgrading from a release before Version 7, your application probably needs to be modified in order to use Version 7 of IBM WebSphere MQ classes for JMS.

Exception listeners

An application can register an exception listener with a `Connection` object. Subsequently, if a problem occurs that makes the connection unusable, IBM WebSphere MQ classes for JMS delivers an exception to the exception listener by calling its `onException()` method. The application then has the opportunity to reestablish the connection.

V7.5.0.8 APAR IT14820, included from IBM WebSphere MQ Version 7.5.0, Fix Pack 8, fixed a defect where an application's JMS `ExceptionListener` would not be invoked for non-connection broken exceptions (for example `MQRC_GET_INHIBITED`) even though the `ASYNC_EXCEPTIONS` property on the JMS `ConnectionFactory` used by the application, was set to `ASYNC_EXCEPTIONS_ALL`. This was the default value prior to Version 7.5.0, Fix Pack 8.

V7.5.0.8 To maintain behavior for current JMS applications that configure a JMS `MessageListener` and a JMS `ExceptionListener`, and to ensure that the IBM WebSphere MQ classes for JMS are consistent with the JMS specification, the default value for the `ASYNC_EXCEPTIONS` JMS `ConnectionFactory` property has been changed to `ASYNC_EXCEPTIONS_CONNECTIONBROKEN` for the IBM WebSphere MQ classes for JMS. As a result, by default, only exceptions corresponding to broken connection error codes are delivered to an application's JMS `ExceptionListener`.

V7.5.0.8 From Version 7.5.0, Fix Pack 8, the IBM WebSphere MQ classes for JMS have also been updated such that `JMSEExceptions` relating to non-connection broken errors, which occur during message delivery to asynchronous message consumers, are still delivered to a registered `ExceptionListener` when the JMS `ConnectionFactory` used by the application has the `ASYNC_EXCEPTIONS` property set to the value `ASYNC_EXCEPTIONS_ALL`.

V7.5.0.8 For more information about what has changed for exception listeners for Version 7.5.0, Fix Pack 8 and why the changes have been made from earlier releases, see [JMS: Exception listener changes in Version 7.5](#).

For any other type of problem, a `JMSEException` exception is thrown by the current JMS API call.

If an application does not register an exception listener with a `Connection` object, any exceptions that would have been delivered to the exception listener are written to the IBM WebSphere MQ classes for JMS log.

Related reference

[ASYNCEXCEPTION](#)

Logging errors in WebSphere MQ classes for JMS

Information about runtime problems that might require corrective action by the user is written to the WebSphere 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, WebSphere 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 WebSphere MQ classes for JMS configuration file. For information about the WebSphere MQ classes for JMS configuration file, see ["The IBM WebSphere MQ classes for JMS configuration file"](#) on page 695, and for more detail about valid values for the `com.ibm.msg.client.commonservices.log.outputName` property, see ["Logging and IBM WebSphere MQ classes for JMS"](#) on page 761.

First failure support technology (FFST) in WebSphere MQ classes for JMS

If a serious internal error occurs within WebSphere MQ classes for JMS, first failure support technology (FFST) information is generated.

The FFST information is written to a file called `JMSCnnnn.FDC`, where `nnnn` is a four digit number. This file is in a directory called `FFDC`, which is a subdirectory of the directory to which trace output is written. By default, trace output is written to the current working directory, but you can redirect trace output to a different directory by setting the `com.ibm.msg.client.commonservices.trace.outputName` property in the WebSphere MQ classes for JMS configuration file. For information about the WebSphere MQ classes for JMS configuration file, see ["The IBM WebSphere MQ classes for JMS configuration file"](#) on page 695.

If tracing is enabled when FFST information is generated, the FFST information is also be written to the trace file. For more information about tracing JMS programs, see [Tracing IBM WebSphere MQ classes for JMS applications](#).

To suppress the production of FFDC files, set the property `com.ibm.msg.client.commonservices.ffst.suppress`, as follows:

0

Output all FFDC files (default).

-1

Output only the first FFDC files of a particular type.

integer

Suppress all FFDC files except those that are a multiple of this number.

Accessing WebSphere MQ features from a WebSphere MQ classes for JMS application

WebSphere MQ classes for JMS provides facilities to exploit a number of features of WebSphere MQ.



Attention: These features are outside the JMS specification or, in certain cases, violate the JMS specification. If you use them, your application is unlikely to be compatible with other JMS providers. Those features which do not comply with the JMS specification are labeled with an Attention notice.

Reading and writing the message descriptor from a WebSphere MQ classes for JMS application

You control the ability to access the message descriptor (MQMD) by setting properties on a Destination and a Message.

Some WebSphere MQ applications require specific values to be set in the MQMD of messages sent to them. WebSphere MQ classes for JMS provides message attributes that allow JMS applications to set MQMD fields and so enable JMS applications to "drive" WebSphere MQ applications.

You must set the Destination object property `WMQ_MQMD_WRITE_ENABLED` to true for the setting of MQMD properties to have any effect. You can then use the property setting methods of the message (for example `setStringProperty`) to assign values to the MQMD fields. All MQMD fields are exposed except `StrucId` and `Version`; `BackoutCount` can be read but not written to.

This example results in a message being put to a queue or topic with `MQMD.UserIdIdentifier` set to "JoeBloggs".

```
// Create a ConnectionFactory, connection, session, producer, message
// ...

// Create a destination
// ...

// Enable MQMD write
dest.setBooleanProperty(WMQConstants.WMQ_MQMD_WRITE_ENABLED, true);

// Optionally, set a message context if applicable for this MD field
dest.setIntProperty(WMQConstants.WMQ_MQMD_MESSAGE_CONTEXT,
    WMQConstants.WMQ_MDCTX_SET_IDENTITY_CONTEXT);

// On the message, set property to provide custom UserId
msg.setStringProperty("JMS_IBM_MQMD_UserIdIdentifier", "JoeBloggs");

// Send the message
// ...
```

It is necessary to set `WMQ_MQMD_MESSAGE_CONTEXT` before setting `JMS_IBM_MQMD_UserIdIdentifier`. For more information about the use of `WMQ_MQMD_MESSAGE_CONTEXT`, see ["JMS message object properties"](#) on page 860.

Similarly, you can extract the contents of the MQMD fields by setting `WMQ_MQMD_READ_ENABLED` to true before receiving a message and then using the get methods of the message, such as `getStringProperty`. Any properties received are read-only.

This example results in the *value* field holding the value of the `MQMD.ApplIdentityData` field of a message got from a queue or a topic.

```
// Create a ConnectionFactory, connection, session, consumer
// ...

// Create a destination
// ...

// Enable MQMD read
dest.setBooleanProperty(WMQConstants.WMQ_MQMD_READ_ENABLED, true);

// Receive a message
// ...

// Get MQMD field value using a property
String value = rcvMsg.getStringProperty("JMS_IBM_MQMD_ApplIdentityData");
```

JMS destination object properties

Two properties of the Destination object control access to the MQMD from JMS, and a third controls message context.

Table 124. Property names and descriptions		
Property	Short form	Description
WMQ_MQMD_WRITE_ENABLED	MDW	Whether a JMS application can set the values of MQMD fields
WMQ_MQMD_READ_ENABLED	MDR	Whether a JMS application can extract the values of MQMD fields
WMQ_MQMD_MESSAGE_CONTEXT	MDCTX	What level of message context is to be set by the JMS application. The application must be running with appropriate context authority for this property to take effect

Table 125. Property names, values, and set methods			
Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
WMQ_MQMD_WRITE_ENABLED	<ul style="list-style-type: none">• NO All JMS_IBM_MQMD* properties are ignored and their values are not copied into the underlying MQMD structure.• YES JMS_IBM_MQMD* properties are processed. Their values are copied into the underlying MQMD structure.	<ul style="list-style-type: none">• False• True	setMQMDWriteEnabled

Table 125. Property names, values, and set methods (continued)

Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
WMQ_MQMD_READ_ENABLED	<ul style="list-style-type: none"> • NO When sending messages, the JMS_IBM_MQMD* properties on a sent message are not updated to reflect the updated field values in the MQMD. When receiving messages, none of the JMS_IBM_MQMD* properties are available on a received message, even if the sender had set some or all of them. • YES When sending messages, all of the JMS_IBM_MQMD* properties on a sent message are updated to reflect the updated field values in the MQMD, including those that the sender did not set explicitly. When receiving messages, all of the JMS_IBM_MQMD* properties are available on a received message, including those that the sender did not set explicitly. 	<ul style="list-style-type: none"> • False • True 	setMQMDReadEnabled
WMQ_MQMD_MESSAGE_CONTEXT	<ul style="list-style-type: none"> • DEFAULT The MQOPEN API call and the MQPMO structure specify no explicit message context options • SET_IDENTITY_CONTEXT The MQOPEN API call specifies the message context option MQOO_SET_IDENTITY_CONTEXT and the MQPMO structure specifies MQPMO_SET_IDENTITY_CONTEXT • SET_ALL_CONTEXT The MQOPEN API call specifies the message context option MQOO_SET_ALL_CONTEXT and the MQPMO structure specifies MQPMO_SET_ALL_CONTEXT 	<ul style="list-style-type: none"> • WMQ_MD CTX_DEFAULT • WMQ_MD CTX_SET_IDENTITY_CONTEXT • WMQ_MD CTX_SET_ALL_CONTEXT 	setMQMDMessageContext

JMS message object properties

Message object properties prefixed JMS_IBM_MQMD allow you to set or read the corresponding MQMD field.

Sending messages

All MQMD fields except StrucId and Version are represented. These properties refer only to the MQMD fields; where a property occurs both in the MQMD and in the MQRFH2 header, the version in the MQRFH2 is not set or extracted.

Any of these properties can be set, except JMS_IBM_MQMD_BackoutCount. Any value set for JMS_IBM_MQMD_BackoutCount is ignored.

If a property has a maximum length and you supply a value that is too long, the value is truncated.

For certain properties, you must also set the WMQ_MQMD_MESSAGE_CONTEXT property on the Destination object. The application must be running with appropriate context authority for this property to take effect. If you do not set WMQ_MQMD_MESSAGE_CONTEXT to an appropriate value, the property value is ignored. If you set WMQ_MQMD_MESSAGE_CONTEXT to an appropriate value but you do not have sufficient context authority for the queue manager, a JMSException is issued. Properties requiring specific values of WMQ_MQMD_MESSAGE_CONTEXT are as follows.

The following properties require WMQ_MQMD_MESSAGE_CONTEXT to be set to WMQ_MDCTX_SET_IDENTITY_CONTEXT or WMQ_MDCTX_SET_ALL_CONTEXT:

- JMS_IBM_MQMD_UserIdentifier
- JMS_IBM_MQMD_AccountingToken
- JMS_IBM_MQMD_ApplIdentityData

The following properties require WMQ_MQMD_MESSAGE_CONTEXT to be set to WMQ_MDCTX_SET_ALL_CONTEXT :

- JMS_IBM_MQMD_PutApplType
- JMS_IBM_MQMD_PutApplName
- JMS_IBM_MQMD_PutDate
- JMS_IBM_MQMD_PutTime
- JMS_IBM_MQMD_ApplOriginData

Receiving messages

All these properties are available on a received message if WMQ_MQMD_READ_ENABLED property is set to true, irrespective of the actual properties the producing application has set. An application cannot modify the properties of a received message unless all properties are cleared first, according to the JMS specification. The received message can be forwarded without modifying the properties.



Attention: If your application receives a message from a destination with WMQ_MQMD_READ_ENABLED property set to true, and forwards it to a destination with WMQ_MQMD_WRITE_ENABLED set to true, this results in all the MQMD field values of the received message being copied into the forwarded message.





Table of properties

This table lists the properties of the Message object representing the MQMD fields. See the links for full descriptions of the fields and their allowable values.

Table 126. Property names, descriptions, and types			
Property	Description	Java Type	Link to full description
JMS_IBM_MQMD_Report	Options for report messages	Integer	Report

Table 126. Property names, descriptions, and types (continued)

Property	Description	Java Type	Link to full description
JMS_IBM_MQMD_MsgType	Message type	Integer	MsgType
JMS_IBM_MQMD_Expiry	Message lifetime	Integer	Expiry
JMS_IBM_MQMD_Feedback	Feedback or reason code	Integer	Feedback
JMS_IBM_MQMD_Encoding	Numeric encoding of message data	Integer	Encoding
JMS_IBM_MQMD_CodedCharSetId	Character set identifier of message data	Integer	CodedCharSetId
JMS_IBM_MQMD_Format	Format name of message data	String	Format
JMS_IBM_MQMD_Priority ¹	Message priority	Integer	Priority
JMS_IBM_MQMD_Persistence	Message persistence	Integer	Persistence
JMS_IBM_MQMD_MsgId ²	Message identifier	Object (byte[]) ⁴	MsgId
JMS_IBM_MQMD_CorrelId ³	Correlation identifier	Object (byte[]) ⁴	CorrelId
JMS_IBM_MQMD_BackoutCount	Backout counter	Integer	BackoutCount
JMS_IBM_MQMD_ReplyToQ	Name of reply queue	String	ReplyToQ
JMS_IBM_MQMD_ReplyToQMgr	Name of reply queue manager	String	ReplyToQMgr
JMS_IBM_MQMD_UserIdentifier	User identifier	String	UserIdentifier
JMS_IBM_MQMD_AccountingToken	Accounting token	Object (byte[]) ⁴	AccountingToken
JMS_IBM_MQMD_ApplIdentityData	Application data relating to identity	String	ApplIdentityData
JMS_IBM_MQMD_PutApplType	Type of application that put the message	Integer	PutApplType
JMS_IBM_MQMD_PutApplName	Name of application that put the message	String	PutApplName
JMS_IBM_MQMD_PutDate	Date when message was put	String	PutDate
JMS_IBM_MQMD_PutTime	Time when message was put	String	PutTime
JMS_IBM_MQMD_ApplOriginData	Application data relating to origin	String	ApplOriginData
JMS_IBM_MQMD_GroupId	Group identifier	Object (byte[]) ⁴	GroupId
JMS_IBM_MQMD_MsgSeqNumber	Sequence number of logical message within group	Integer	MsgSeqNumber

Table 126. Property names, descriptions, and types (continued)			
Property	Description	Java Type	Link to full description
JMS_IBM_MQMD_Offset	Offset of data in physical message from start of logical message	Integer	Offset
JMS_IBM_MQMD_MsgFlags	Message flags	Integer	MsgFlags
JMS_IBM_MQMD_OriginalLength	Length of original message	Integer	OriginalLength
<ol style="list-style-type: none">  Attention: If you assign a value to JMS_IBM_MQMD_Priority that is not within the range 0-9, this violates the JMS specification.  Attention: The JMS specification states that the message ID must be set by the JMS provider and that it must either be unique or null. If you assign a value to JMS_IBM_MQMD_MsgId, this value is copied to the JMSMessageID. Thus it is not set by the JMS provider and might not be unique: this violates the JMS specification.  Attention: If you assign a value to JMS_IBM_MQMD_CorrelId that starts with the string 'ID:', this violates the JMS specification.  Attention: The use of byte array properties on a message violates the JMS specification. 			

Accessing IBM WebSphere MQ Message data from an application using WebSphere MQ classes for JMS

You can access the complete WebSphere MQ message data within an application using IBM WebSphere MQ classes for JMS. To access all the data, the message must be a JMSBytesMessage. The body of the JMSBytesMessage includes any MQRFH2 header, any other IBM WebSphere MQ headers, and the following message data.

Set the WMQ_MESSAGE_BODY property of the destination to WMQ_MESSAGE_BODY_MQ, to receive all the message body data in the JMSBytesMessage.

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_JMS or WMQ_MESSAGE_BODY_UNSPECIFIED, the message body is returned without the JMS MQRFH2 header, and the properties of the JMSBytesMessage reflect the properties set in the RFH2.

Some applications cannot use the functions described in this topic. If an application is connected to a WebSphere MQ V6 queue manager, or if it has set PROVIDERVERSION to 6, the functions are not available.

Sending a message

When sending messages the destination property, WMQ_MESSAGE_BODY, takes precedence over WMQ_TARGET_CLIENT.

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_JMS, WebSphere MQ classes for JMS automatically generates an MQRFH2 header based on the settings of the JMSMessage properties and header fields.

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_MQ, no additional header is added to the message body

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_UNSPECIFIED, WebSphere MQ classes for JMS sends an MQRFH2 header, unless WMQ_TARGET_CLIENT is set to WMQ_TARGET_DEST_MQ. On receive, setting WMQ_TARGET_CLIENT to WMQ_TARGET_DEST_MQ results in any MQRFH2 being removed from the message body.

Note: JMSBytesMessage and JMSTextMessage do not require an MQRFH2, whereas JMSStreamMessage, JMSMapMessage, and JMSObjectMessage do.

WMQ_MESSAGE_BODY_UNSPECIFIED is the default setting for WMQ_MESSAGE_BODY, and WMQ_TARGET_DEST_JMS is the default setting for WMQ_TARGET_CLIENT.

If you send a JMSBytesMessage, you can override the default settings for the JMS message body when the WebSphere MQ message is constructed. Use the following properties:

- JMS_IBM_Format or JMS_IBM_MQMD_Format: This property specifies the format of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding Websphere MQ header.
- JMS_IBM_Character_Set or JMS_IBM_MQMD_CodedCharSetId: This property specifies the CCSID of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding Websphere MQ header.
- JMS_IBM_Encoding or JMS_IBM_MQMD_Encoding: This property specifies the encoding of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding Websphere MQ header.

If both types of property are specified, the JMS_IBM_MQMD_* properties override the corresponding JMS_IBM_* properties, as long as the destination property WMQ_MQMD_WRITE_ENABLED is set to true.

The differences in effect between setting message properties using JMS_IBM_MQMD_* and JMS_IBM_* are significant:

1. The JMS_IBM_MQMD_* properties are specific to the IBM WebSphere MQ JMS provider.
2. The JMS_IBM_MQMD_* properties are only set in the MQMD. JMS_IBM_* properties are set in the MQMD only if the message does not have an MQRFH2 JMS header. Otherwise they are set in the JMS RFH2 header.
3. The JMS_IBM_MQMD_* properties have no affect on the encoding of text and numbers written into a JMSMessage.

A receiving application is likely to assume the values of MQMD.Encoding and MQMD.CodedCharSetId correspond to the encoding and character set of numbers and text in the message body. If JMS_IBM_MQMD_* properties are used, it is the responsibility of the sending application to make it so. The encoding and character set of numbers and text in the message body are set by the JMS_IBM_* properties.

The badly coded snippet in [Figure 163 on page 863](#) sends a message encoded in character set 1208, with MQMD.CodedCharSetId set to 37.

a. Send wrongly encoded message

```
TextMessage tmo = session.createTextMessage();
((MQDestination) destination).setMessageBodyStyle
    (WMQConstants.WMQ_MESSAGE_BODY_MQ);
((MQDestination) destination).setMQMDWriteEnabled(true);
tmo.setIntProperty(WMQConstants.JMS_IBM_MQMD_CODEDCHARSETID, 37);
tmo.setIntProperty(WMQConstants.JMS_IBM_CHARACTER_SET, 1208);
tmo.setText("String one");
producer.send(tmo);
```

b. Receiving the message, relying on the value of JMS_IBM_CHARACTER_SET set by the value of MQMD.CodedCharSetId:

```
TextMessage tmi = (TextMessage) cons.receive();
System.out.println("Message is \"" + tmi.getText() + "\"");
```

c. Resulting output:

```
Message is "ëËË'>...??>?"
```

Figure 163. Inconsistently coded MQMD and message data

Either of the snippets of code in [Figure 164](#) on [page 864](#) results in a message being put to a queue or topic, with its body containing the application payload without an automatically generated MQRFH2 header being added.

1. Setting WMQ_MESSAGE_BODY_MQ:

```
((MQDestination) destination).setMessageBodyStyle  
    (WMQConstants.WMQ_MESSAGE_BODY_MQ);
```

2. Setting WMQ_TARGET_DEST_MQ:

```
((MQDestination) destination).setMessageBodyStyle  
    (WMQConstants.WMQ_MESSAGE_BODY_UNSPECIFIED);  
((MQDestination) destination).  
    setTargetClient(WMQConstants.WMQ_TARGET_DEST_MQ);
```

Figure 164. Send a message with an MQ message body.

Receiving a message

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_JMS, the inbound JMS message type and body are determined by the contents of the received WebSphere MQ message. The message type and body are determined by fields in the MQRFH2 header, or in the MQMD, if there is no MQRFH2.

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_MQ, the inbound JMS message type is JMSBytesMessage. The JMS message body is the message data returned by the underlying MQGET API call. The length of message body is the length returned by the MQGET call. The character set and encoding of the data in the message body is determined by the CodedCharSetId and Encoding fields of the MQMD. The format of the data in the message body is determined by the Format field of the MQMD.

If WMQ_MESSAGE_BODY is set to WMQ_MESSAGE_BODY_UNSPECIFIED, the default value, IBM WebSphere MQ classes for JMS sets it to WMQ_MESSAGE_BODY_JMS.

When you receive a JMSBytesMessage, you can decode it by reference to the following properties:

- JMS_IBM_Format or JMS_IBM_MQMD_Format: This property specifies the format of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.
- JMS_IBM_Character_Set or JMS_IBM_MQMD_CodedCharSetId: This property specifies the CCSID of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.
- JMS_IBM_Encoding or JMS_IBM_MQMD_Encoding: This property specifies the encoding of the WebSphere MQ header or application payload that starts the JMS message body if there is no preceding WebSphere MQ header.

The following code snippet results in a received message that is a JMSBytesMessage. Irrespective of the content of the received message and of the format field of the received MQMD, the message is a JMSBytesMessage.

```
((MQDestination)destination).setMessageBodyStyle  
    (WMQConstants.WMQ_MESSAGE_BODY_MQ);
```


Destination property *WMQ_MESSAGE_BODY*

WMQ_MESSAGE_BODY determines whether a JMS application processes the MQRFH2 of a WebSphere MQ message as part of the message payload (that is, as part of the JMS message body).

Table 127. Property names and descriptions		
Property	Short form	Description
WMQ_MESSAGE_BODY	MBODY	Whether a JMS application processes the MQRFH2 of a WebSphere MQ message as part of the message payload (that is, as part of the JMS message body).

Table 128. Property names, values, and set methods			
Property	Valid values in administration tool (defaults in bold)	Valid values in programs	Set method
WMQ_MESSAGE_BODY	<ul style="list-style-type: none"> • UNSPECIFIED When sending, WebSphere MQ classes for JMS does or does not generate and include an MQRFH2 header, depending on the value of WMQ_TARGET_CLIENT. When receiving, acts as value JMS. • JMS When sending, WebSphere MQ classes for JMS automatically generates an MQRFH2 header and includes it in the WebSphere MQ message. When receiving, WebSphere MQ classes for JMS set the JMS message properties according to values in the MQRFH2 (if present); it does not present the MQRFH2 as part of the JMS message body. • MQ When sending, WebSphere MQ classes for JMS does not generate an MQRFH2. When receiving, WebSphere MQ classes for JMS presents the MQRFH2 as part of the JMS message body. 	<ul style="list-style-type: none"> • WMQ_MESSAGE_BODY_UNSPECIFIED • WMQ_MESSAGE_BODY_JMS • WMQ_MESSAGE_BODY_MQ 	setMessageBodyStyle

JMS persistent messages

WebSphere MQ classes for JMS applications can use the **NonPersistentMessageClass** queue attribute to provide better performance for JMS persistent messages, at the expense of some reliability.

A WebSphere MQ queue has an attribute called **NonPersistentMessageClass**. The value of this attribute determines whether nonpersistent messages on the queue are discarded when the queue manager restarts.

You can set the attribute for a local queue by using the WebSphere MQ Script (MQSC) command, **DEFINE QLOCAL**, with either of the following parameters:

NPMCLASS(NORMAL)

Nonpersistent messages on the queue are discarded when the queue manager restarts. This is the default value.

NPMCLASS(HIGH)

Nonpersistent messages on the queue are not discarded when the queue manager restarts following a quiesced or immediate shutdown. Nonpersistent messages might be discarded, however, following a preemptive shutdown or a failure.

This topic describes how WebSphere MQ classes for JMS applications can use this queue attribute to provide better performance for JMS persistent messages.

The **PERSISTENCE** property of a Queue or Topic object can have the value **HIGH**. You can use the WebSphere MQ JMS administration tool to set this value, or an application can call the `Destination.setPersistence()` method passing the value `WMQConstants.WMQ_PER_NPHIGH` as a parameter.

If an application sends a JMS persistent message or a JMS nonpersistent message to a destination where the **PERSISTENCE** property has the value **HIGH**, and the underlying WebSphere MQ queue is set to **NPMCLASS(HIGH)**, the message is put on the queue as a WebSphere MQ nonpersistent message. If the **PERSISTENCE** property of the destination does not have the value **HIGH**, or if the underlying queue is set to **NPMCLASS(NORMAL)**, a JMS persistent message is put on the queue as a WebSphere MQ persistent message, and a JMS nonpersistent message is put on the queue as a WebSphere MQ nonpersistent message.

If a JMS persistent message is put on a queue as a WebSphere MQ nonpersistent message, and you want to ensure that the message is not discarded following a quiesced or immediate shutdown of a queue manager, all queues through which the message might be routed must be set to **NPMCLASS(HIGH)**. In the publish/subscribe domain, these queues include subscriber queues. As an aid to enforcing this configuration, WebSphere MQ classes for JMS throws an `InvalidDestinationException` if an application tries to create a message consumer for a destination where the **PERSISTENCE** property has the value **HIGH** and the underlying WebSphere MQ queue is set to **NPMCLASS(NORMAL)**.

Setting the **PERSISTENCE** property of a destination to **HIGH** does not affect how a message is received from that destination. A message sent as a JMS persistent message is received as a JMS persistent message, and a message sent as a JMS nonpersistent message is received as a JMS nonpersistent message.

When an application sends the first message to a destination where the **PERSISTENCE** property has the value **HIGH**, or when an application creates the first message consumer for a destination where the **PERSISTENCE** property has the value **HIGH**, WebSphere MQ classes for JMS issues an `MQINQ` call to determine whether **NPMCLASS(HIGH)** is set on the underlying WebSphere MQ queue. The application must therefore have the authority to inquire on the queue. In addition, WebSphere MQ classes for JMS preserves the result of the `MQINQ` call until the destination is deleted, and does not issue more `MQINQ` calls. Therefore, if you change the **NPMCLASS** setting on the underlying queue while the application is still using the destination, WebSphere MQ classes for JMS does not notice the new setting.

By allowing JMS persistent messages to be put on WebSphere MQ queues as WebSphere MQ nonpersistent messages, you are gaining performance at the expense of some reliability. If you require maximum reliability for JMS persistent messages, do not send the messages to a destination where the **PERSISTENCE** property has the value **HIGH**.

The JMS Layer can use `SYSTEM.JMS.TEMPQ.MODEL`, instead of `SYSTEM.DEFAULT.MODEL.QUEUE`. `SYSTEM.JMS.TEMPQ.MODEL` creates permanent dynamic queues that accept persistent messages, because `SYSTEM.DEFAULT.MODEL.QUEUE` cannot accept persistent messages. If you wish to use temporary queues to accept persistent messages, you must therefore use `SYSTEM.JMS.TEMPQ.MODEL`, or change the model queue to an alternative queue of your choosing.

Using Secure Sockets Layer (SSL) with WebSphere MQ classes for JMS

WebSphere MQ classes for JMS applications can use SSL encryption. To do this they require a JSSE provider.

WebSphere MQ classes for JMS connections using `TRANSPORT(CLIENT)` support Secure Sockets Layer (SSL) encryption. SSL provides communication encryption, authentication, and message integrity. It is typically used to secure communications between any two peers on the Internet or within an intranet.

WebSphere MQ classes for JMS uses Java Secure Socket Extension (JSSE) to handle SSL encryption, and therefore requires a JSSE provider. JSE v1.4 JVMs have a JSSE provider built-in. Details of how to manage and store certificates can vary from provider to provider. For information about this, see your JSSE provider's documentation.

This section assumes that your JSSE provider is correctly installed and configured, and that suitable certificates have been installed and made available to your JSSE provider. You can now use JMSAdmin to set a number of administrative properties.

If your WebSphere MQ classes for JMS application uses a client channel definition table (CCDT) to connect to a queue manager, see [“Using a client channel definition table with IBM WebSphere MQ classes for JMS”](#) on page 875.

SSLCIPHERSUITE object property

Set `SSLCIPHERSUITE` to enable SSL encryption on a `ConnectionFactory` object.

To enable SSL encryption on a `ConnectionFactory` object, use JMSAdmin to set the `SSLCIPHERSUITE` property to a CipherSuite supported by your JSSE provider. This must match the CipherSpec set on the target channel. However, CipherSuites are distinct from CipherSpecs and therefore have different names. [“SSL CipherSpecs and CipherSuites in JMS”](#) on page 870 contains a table mapping the CipherSpecs supported by WebSphere MQ to their equivalent CipherSuites as known to JSSE. For more information about CipherSpecs and CipherSuites with WebSphere MQ, see [Security](#).

For example, to set up a `ConnectionFactory` object that can be used to create a connection over an SSL enabled MQI channel with a CipherSpec of `RC4_MD5_EXPORT`, issue the following command to JMSAdmin:

```
ALTER CF(my.cf) SSLCIPHERSUITE(SSL_RSA_EXPORT_WITH_RC4_40_MD5)
```

This can also be set from an application, using the `setSSLCipherSuite()` method on an `MQConnectionFactory` object.

For convenience, if a CipherSpec is specified on the `SSLCIPHERSUITE` property, JMSAdmin attempts to map the CipherSpec to an appropriate CipherSuite and issues a warning. This attempt to map is not made if the property is specified by an application.

Alternatively, use the Client Channel Definition Table (CCDT). For more information, see [“Using a client channel definition table with IBM WebSphere MQ classes for JMS”](#) on page 875.

SSLFIPSREQUIRED object property

If you require a connection to use a CipherSuite that is supported by the IBM Java JSSE FIPS provider (`IBMJSSEFIPS`), set the `SSLFIPSREQUIRED` property of the connection factory to `YES`.

The default value of this property is `NO`, which means that a connection can use any CipherSuite that is supported by WebSphere MQ.

If an application uses more than one connection, the value of `SSLFIPSREQUIRED` that is used when the application creates the first connection determines the value that is used when the application

creates any subsequent connection. This means that the value of the SSLFIPSREQUIRED property of the connection factory that is used to create a subsequent connection is ignored. You must restart the application if you want to use a different value of SSLFIPSREQUIRED.

An application can set this property by calling the `setSSLFipsRequired()` method of a `ConnectionFactory` object. The property is ignored if no `CipherSuite` is set.

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](#)

SSLPEERNAME object property

Use `SSLPEERNAME` to specify a distinguished name pattern, to ensure that your JMS application connects to the correct queue manager.

A JMS application can ensure that it connects to the correct queue manager by specifying a distinguished name (DN) pattern. The connection succeeds only if the queue manager presents a DN that matches the pattern. For more details of the format of this pattern, see the related topics.

The DN is set using the `SSLPEERNAME` property of a `ConnectionFactory` object. For example, the following `JMSAdmin` command sets a `ConnectionFactory` object to expect the queue manager to identify itself with a Common Name beginning with the characters `QMGR.`, and with at least two Organizational Unit names, the first of which must be `IBM` and the second `WEBSphere`:

```
ALTER CF(my.cf) SSLPEERNAME(CN=QMGR.*, OU=IBM, OU=WEBSphere)
```

Checking is not case sensitive and semicolons can be used in place of commas. `SSLPEERNAME` can also be set from an application using the `setSSLPeerName()` method on an `MQConnectionFactory` object. If this property is not set, no checking is performed on the Distinguished Name supplied by the queue manager. This property is ignored if no `CipherSuite` is set.

SSLCERTSTORES object property

Use `SSLCERTSTORES` to specify a list of LDAP servers to use for certificate revocation list (CRL) checking.

It is common to use a certificate revocation list (CRL) to identify certificates that are no longer trusted. CRLs are typically hosted on LDAP servers. JMS allows an LDAP server to be specified for CRL checking under Java 2 v1.4 or later. The following `JMSAdmin` example directs JMS to use a CRL hosted on an LDAP server named `cr11.ibm.com`:

```
ALTER CF(my.cf) SSLCRL(ldap://cr11.ibm.com)
```

Note: To use a `CertStore` successfully with a CRL hosted on an LDAP server, make sure that your Java Software Development Kit (SDK) is compatible with the CRL. Some SDKs require that the CRL conforms to RFC 2587, which defines a schema for LDAP v2. Most LDAP v3 servers use RFC 2256 instead.

If your LDAP server is not running on the default port of 389, you can specify the port by appending a colon (:) and the port number to the host name. If the certificate presented by the queue manager is present in the CRL hosted on `cr11.ibm.com`, the connection is not completed. To avoid a single point of failure, JMS allows multiple LDAP servers to be supplied by supplying a list of LDAP servers delimited by the space character. Here is an example:

```
ALTER CF(my.cf) SSLCRL(ldap://cr11.ibm.com ldap://cr12.ibm.com)
```

When multiple LDAP servers are specified, JMS tries each one in turn until it finds a server with which it can successfully verify the queue manager's certificate. Each server must contain identical information.

A string in this format can be supplied by an application on the `MQConnectionFactory.setSSLCertStores()` method. Alternatively, the application can create one or more `java.security.cert.CertStore` objects, place these in a suitable `Collection` object, and supply this `Collection` object to the `setSSLCertStores()` method.

In this way, the application can customize CRL checking. See your JSSE documentation for details on constructing and using CertStore objects.

The certificate presented by the queue manager when a connection is being set up is validated as follows:

1. The first CertStore object in the Collection identified by sslCertStores is used to identify a CRL server.
2. An attempt is made to contact the CRL server.
3. If the attempt is successful, the server is searched for a match for the certificate.
 - a. If the certificate is found to be revoked, the search process is over and the connection request fails with reason code MQRC_SSL_CERTIFICATE_REVOKED.
 - b. If the certificate is not found, the search process is over and the connection is allowed to proceed.
4. If the attempt to contact the server is unsuccessful, the next CertStore object is used to identify a CRL server and the process repeats from step 2.

If this was the last CertStore in the Collection, or if the Collection contains no CertStore objects, the search process has failed and the connection request fails with reason code MQRC_SSL_CERT_STORE_ERROR.

The Collection object determines the order in which CertStores are used.

If your application uses setSSLCertStores() to set a Collection of CertStore objects, the MQConnectionFactory can no longer be bound into a JNDI namespace. Attempting to do so causes an exception. If the sslCertStores property is not set, no revocation checking is performed on the certificate provided by the queue manager. This property is ignored if no CipherSuite is set.

SSLRESETCOUNT object property

This property represents the total number of bytes sent and received by a connection before the secret key that is used for encryption is renegotiated.

The number of bytes sent is the number before encryption, and the number of bytes received is the number after decryption. The number of bytes also includes control information sent and received by WebSphere MQ classes for JMS.

For example, to configure a ConnectionFactory object that can be used to create a connection over an SSL enabled MQI channel with a secret key that is renegotiated after 4 MB of data have flowed, issue the following command to JMSAdmin:

```
ALTER CF(my.cf) SSLRESETCOUNT(4194304)
```

An application can set this property by calling the setSSLResetCount() method of a ConnectionFactory object.

If the value of this property is zero, which is the default value, the secret key is never renegotiated. The property is ignored if no CipherSuite is set.

SSLSocketFactory object property

To customize other aspects of the SSL connection for an application, create an SSLSocketFactory and configure JMS to use it.

You might want to customize other aspects of the SSL connection for an application. For example, you might want to initialize cryptographic hardware or change the keystore and truststore in use. To do this, the application must first create a javax.net.ssl.SSLSocketFactory object that is customized accordingly. See your JSSE documentation for information about how to do this, because the customizable features vary from provider to provider. After a suitable SSLSocketFactory object is obtained, use the MQConnectionFactory.setSSLSocketFactory() method to configure JMS to use the customized SSLSocketFactory object.

If your application uses the setSSLSocketFactory() method to set a customized SSLSocketFactory object, the MQConnectionFactory object can no longer be bound into a JNDI namespace. Attempting to do so causes an exception. If this property is not set, the default SSLSocketFactory object is used. See your

JSSE documentation for details of the behavior of the default SSLSocketFactory object. This property is ignored if no CipherSuite is set.

Important: Do not assume that the use of the SSL properties ensures security when a ConnectionFactory object is retrieved from a JNDI namespace that is not itself secure. Specifically, the standard LDAP implementation of JNDI is not secure. An attacker can imitate the LDAP server, misleading a JMS application into connecting to the wrong server without noticing. With suitable security arrangements in place, other implementations of JNDI (such as the fscontext implementation) are secure.

Making changes to the JSSE keystore or truststore

If you make changes to the keystore or truststore, you must take certain actions for the changes to be picked up.

If you change the contents of the JSSE keystore or truststore, or change the location of the keystore or truststore file, WebSphere MQ classes for JMS applications that are running at the time do not automatically pick up the changes. For the changes to take effect, the following actions must be performed:

- The applications must close all their connections, and destroy any unused connections in connection pools.
- If your JSSE provider caches information from the keystore and truststore, this information must be refreshed.

After these actions have been performed, the applications can then re-create their connections.

Depending on how you design your applications, and on the function provided by your JSSE provider, it might be possible to perform these actions without stopping and restarting your applications. However, stopping and restarting the applications might be the simplest solution.

SSL CipherSpecs and CipherSuites in JMS

CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites.

Table 129 on page 870 lists the CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites. If the ConnectionFactory property SSLFIPSREQUIRED is set to NO, a WebSphere MQ classes for JMS application can connect to a queue manager if any supported CipherSpec is specified at the server end of the MQI channel and the equivalent CipherSuite is specified at the client end. If SSLFIPSREQUIRED is set to YES, the combination of CipherSpec and CipherSuite determines whether the application can connect to the queue manager.

At the server end of an MQI channel, the name of a CipherSpec can be specified as the value of the SSLCIPH parameter on a DEFINE CHANNEL CHLTYPE(SVRCONN) command. At the client end of an MQI channel, the name of a CipherSuite can be specified in the following ways:

- An application can call the setSSLCipherSuite() method of a ConnectionFactory object.
- Using the WebSphere MQ JMS administration tool, you can set the SSLCIPHERSUITE property of a ConnectionFactory object.

Table 129. CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites		
CipherSpec	Equivalent CipherSuite	Connection possible if SFIPS ¹ is set to YES?
NULL_MD5	SSL_RSA_WITH_NULL_MD5	No
NULL_SHA	SSL_RSA_WITH_NULL_SHA	No
RC4_MD5_EXPORT	SSL_RSA_EXPORT_WITH_RC4_40_MD5	No
RC4_MD5_US	SSL_RSA_WITH_RC4_128_MD5	No
RC4_SHA_US	SSL_RSA_WITH_RC4_128_SHA	No

Table 129. CipherSpecs supported by WebSphere MQ and their equivalent CipherSuites (continued)

CipherSpec	Equivalent CipherSuite	Connection possible if SFIPS ¹ is set to YES?
RC2_MD5_EXPORT	SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5	No
DES_SHA_EXPORT	SSL_RSA_WITH_DES_CBC_SHA	No
RC4_56_SHA_EXPORT1024	SSL_RSA_EXPORT1024_WITH_RC4_56_SHA	No
DES_SHA_EXPORT1024	SSL_RSA_EXPORT1024_WITH_DES_CBC_SHA	No
TRIPLE_DES_SHA_US	SSL_RSA_WITH_3DES_EDE_CBC_SHA	No
TLS_RSA_WITH_NULL_SHA256	SSL_RSA_WITH_NULL_SHA256	No ⁷
TLS_RSA_WITH_AES_128_CBC_SHA	SSL_RSA_WITH_AES_128_CBC_SHA	Yes ^{5 7}
TLS_RSA_WITH_AES_256_CBC_SHA	SSL_RSA_WITH_AES_256_CBC_SHA	Yes ^{5 7}
TLS_RSA_WITH_AES_256_CBC_SHA256	SSL_RSA_WITH_AES_256_CBC_SHA256	Yes ^{5 7}
AES_SHA_US ²		
TLS_RSA_WITH_DES_CBC_SHA ^{8 9}	SSL_RSA_WITH_DES_CBC_SHA	No ³
TLS_RSA_WITH_3DES_EDE_CBC_SHA ⁸	SSL_RSA_WITH_3DES_EDE_CBC_SHA	Yes
FIPS_WITH_DES_CBC_SHA	SSL_RSA_FIPS_WITH_DES_CBC_SHA	No ⁴
FIPS_WITH_3DES_EDE_CBC_SHA	SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA	No ⁶

Notes:

1. When using the WebSphere MQ JMS administration tool, SFIPS is the short name of the ConnectionFactory property SSLFIPSREQUIRED.
2. This CipherSpec has no equivalent CipherSuite.
3. This CipherSpec was FIPS 140-2 certified before 19th May 2007.
4. This CipherSpec was FIPS 140-2 certified before 19th May 2007. The name FIPS_WITH_DES_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.
5. These CipherSpecs (TLS_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA256) cannot be used to secure a connection from the WebSphere MQ Explorer to a queue manager unless the appropriate unrestricted policy files are applied to the JRE used by the Explorer.
See [Security information](#) for further information on policy files.
6. The name FIPS_WITH_3DES_EDE_CBC_SHA is historical and reflects the fact that this CipherSpec was previously (but is no longer) FIPS-compliant. This CipherSpec is deprecated and its use is not recommended.
7. These CipherSpecs (TLS_RSA_WITH_NULL_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA256) require IBM JREs 6.0 SR13 FP2, 7.0 SR4 FP2 or later.
8. These CipherSpecs (TLS_RSA_WITH_3DES_EDE_CBC_SHA, TLS_RSA_WITH_DES_CBC_SHA, TLS_RSA_WITH_RC4_128_SHA256) can use either SSLv3 or TLS. By default, when FIPS is not enabled, SSLv3 is used. To use TLS, set the Java System Property **com.ibm.mq.cfg.preferTLS** to true.

9. This CipherSpec TLS_RSA_WITH_3DES_EDE_CBC_SHA is deprecated. However, it can still be used to transfer up to 32 GB of data before the connection is terminated with error AMQ9288. To avoid this error, you need to either avoid using triple DES, or enable secret key reset when using this CipherSpec.

Related information

Specifying that only FIPS-certified CipherSpecs are used at run time on the MQI client
[Federal Information Processing Standards \(FIPS\) for UNIX, Linux and Windows](#)

Writing channel exits in Java for WebSphere MQ classes for JMS

You create channel exits by defining Java classes that implement specified interfaces.

Three interfaces are defined in the com.ibm.mq.exits package:

- WMQSendExit, for a send exit
- WMQReceiveExit, for a receive exit
- WMQSecurityExit, for a security exit

The following sample code defines a class that implements all three interfaces:

```
public class MyMQExits implements
WMQSendExit, WMQReceiveExit, WMQSecurityExit {
    // Default constructor
    public MyMQExits(){
    }
    // This method implements the send exit interface
    public ByteBuffer channelSendExit(
MQCXP channelExitParms,
                                MQCD channelDefinition,
                                ByteBuffer agentBuffer)
    {
        // Complete the body of the send exit here
    }
    // This method implements the receive exit interface
    public ByteBuffer channelReceiveExit(
MQCXP channelExitParms,
                                MQCD channelDefinition,
                                ByteBuffer agentBuffer)
    {
        // Complete the body of the receive exit here
    }
    // This method implements the security exit interface
    public ByteBuffer channelSecurityExit(
MQCXP channelExitParms,
                                MQCD channelDefinition,
                                ByteBuffer agentBuffer)
    {
        // Complete the body of the security exit here
    }
}
```

Each exit receives as parameters an MQCXP object and an MQCD object. These objects represent the MQCXP and MQCD structures defined in the procedural interface.

When a send exit is called, the agentBuffer parameter contains the data that is about to be sent to the server queue manager. A length parameter is not required because the expression agentBuffer.limit() provides the length of the data. The send exit returns as its value the data to be sent to the server queue manager. However, if the send exit is not the last send exit in a sequence of send exits, the data returned is passed instead to the next send exit in the sequence. A send exit can return a modified version of the data that it receives in the agentBuffer parameter, or it can return the data unchanged. The simplest possible exit body is therefore:

```
{ return agentBuffer; }
```

When a receive exit is called, the agentBuffer parameter contains the data that has been received from the server queue manager. The receive exit returns as its value the data to be passed to the application

by WebSphere MQ classes for JMS. However, if the receive exit is not the last receive exit in a sequence of receive exits, the data returned is passed instead to the next receive exit in the sequence.

When a security exit is called, the `agentBuffer` parameter contains the data that has been received in a security flow from the security exit at the server end of the connection. The security exit returns as its value the data to be sent in a security flow to the server security exit.

Channel exits are called with a buffer that has a backing array. For best performance, the exit should return a buffer with a backing array.

Up to 32 characters of user data can be passed to a channel exit when it is called. The exit accesses the user data by calling the `getExitData()` method of the MQCXP object. Although the exit can change the user data by calling the `setExitData()` method, the user data is refreshed every time the exit is called. Any changes made to the user data are therefore lost. However, the exit can pass data from one call to the next by using the exit user area of the MQCXP object. The exit accesses the exit user area by reference by calling the `getExitUserArea()` method.

Every exit class must have a constructor. The constructor can be either the default constructor, as shown in the previous example, or a constructor with a string parameter. The constructor is called to create an instance of the exit class for each exit defined in the class. Therefore, in the previous example, an instance of the `MyMQExits` class is created for the send exit, another instance is created for the receive exit, and a third instance is created for the security exit. When a constructor with a string parameter is called, the parameter contains the same user data that is passed to the channel exit for which the instance is being created. If an exit class has both a default constructor and a single parameter constructor, the single parameter constructor takes precedence.

Do not close the connection from within a channel exit.

When data is sent to the server end of a connection, SSL encryption is performed *after* any channel exits are called. Similarly, when data is received from the server end of a connection, SSL decryption is performed *before* any channel exits are called.

In versions of WebSphere MQ classes for JMS earlier than Version 7.0, channel exits were implemented using the interfaces `MQSendExit`, `MQReceiveExit`, and `MQSecurityExit`. You can still use these interfaces, but the new interfaces are preferred for improved function and performance.

Configuring IBM WebSphere MQ classes for JMS to use channel exits

An IBM WebSphere MQ classes for JMS application can use channel security, send, and receive exits on the MQI channel that starts when the application connects to a queue manager. The application can use exits written in Java, C, or C++. The application can also use a sequence of send or receive exits that are run in succession.

The following properties are used specify a send exit, or a sequence of send exits, used by a JMS connection:

- The **SENDEXIT** property of an `MQConnectionFactory` object.
- The **sendexit** property on an activation specification used by the IBM WebSphere MQ resource adapter for inbound communication,
- The **sendexit** property on a `ConnectionFactory` object used by the IBM WebSphere MQ resource adapter for output communication.

The value of the property is a string that comprises one or more items separated by commas. Each item identifies a send exit in one of the following ways:

- The name of a class that implements the `WMQSendExit` interface for a send exit written in Java.
- A string in the format *libraryName (entryPointName)* for a send exit written in C or C++.

In a similar way, the following properties specify the receive exit, or sequence of receive exits, used by a connection:

- The **RECEXIT** property of an `MQConnectionFactory` object.
- The **receiveexit** property on an activation specification used by the IBM WebSphere MQ resource adapter for inbound communication,

- The **receiveexit** property on a ConnectionFactory object used by the IBM WebSphere MQ resource adapter for output communication.

The following properties specify the security exit used by a connection:

- The **SECEXIT** property of an MQConnectionFactory object.
- The **securityexit** property on an activation specification used by the IBM WebSphere MQ resource adapter for inbound communication,
- The **securityexit** property on a ConnectionFactory object used by the IBM WebSphere MQ resource adapter for output communication.

For MQConnectionFactory, you can set the **SENDEXIT**, **RECEXIT** and **SECEXIT** properties by using the IBM WebSphere MQ JMS administration tool or IBM WebSphere MQ Explorer. Alternatively, an application can set the properties by calling the `setSendExit()`, `setReceiveExit()`, and `setSecurityExit()` methods.

Channel exits are loaded by their own class loader. To find a channel exit, the class loader searches the following locations in the specified order.

1. The class path specified by the property **com.ibm.mq.cfg.ClientExitPath.JavaExitsClasspath** or by the **JavaExitsClassPath** attribute in the Channels stanza of the IBM WebSphere MQ client configuration file.
2. The class path specified by the Java system property **com.ibm.mq.exitClasspath**. Note that this property is now deprecated.
3. The IBM WebSphere MQ exits directory, as shown in [Table 130 on page 874](#). The class loader first searches the directory for class files that are not packaged in Java archive (JAR) files. If the channel exit is not found, the class loader then searches the JAR files in the directory.

Table 130. The IBM WebSphere MQ exits directory	
Platform	Directory
UNIX and Linux	/var/mqm/exits (32-bit channel exits) /var/mqm/exits64 (64-bit channel exits)
Windows	<i>install_data_dir</i> \exits where <i>install_data_dir</i> is the directory that you chose for the IBM WebSphere MQ data files during installation. The default directory is C:\Program Files\IBM\WebSphere MQ.

Note: If a channel exit exists in more than one location, the IBM WebSphere MQ classes for JMS loads the first instance that it finds.

The parent of the class loader is the class loader that is used to load IBM WebSphere MQ classes for JMS. It is therefore possible for the parent class loader to load a channel exit if it cannot be found in any of the preceding locations. However, when you are using the IBM WebSphere MQ classes for JMS in an environment such as a JEE application server, you are not likely to be able to influence the choice of the parent class loader and so the class loader should be configured by setting the Java system property **com.ibm.mq.cfg.ClientExitPath.JavaExitsClasspath** on the application server.

If your application is being run with the Java Security Manager enabled, then the policy configuration file used by the Java runtime environment that the application is running in must have the permissions to load a channel exit class. For information on how to do this, see [Running IBM MQ classes for JMS applications under the Java Security Manager](#).

The MQSendExit, MQReceiveExit, and MQSecurityExit interfaces supplied with versions of IBM WebSphere MQ earlier than Version 7.0 are still supported. If you use channel exits that implement these interfaces, `com.ibm.mq.jar` must be present in the class path.

For information about how to write channel exits in C, see [“Channel-exit programs for messaging channels”](#) on page 378. You must store channel exit programs written in C or C++ in the directory shown in Table 130 on page 874.

If your application uses a client channel definition table (CCDT) to connect to a queue manager, see [“Using a client channel definition table with IBM WebSphere MQ classes for JMS”](#) on page 875.

Specifying the user data to be passed to channel exits when using WebSphere MQ classes for JMS

Up to 32 characters of user data can be passed to a channel exit when it is called.

The `SENDEXITINIT` property of an `MQConnectionFactory` object specifies the user data that is passed to each send exit when it is called. The value of the property is a string that comprises one or more items of user data separated by commas. The position of each item of user data within the string determines which send exit, in a sequence of send exits, the user data is passed to. For example, the first item of user data in the string is passed to the first send exit in a sequence of send exits.

You can set the `SENDEXITINIT` property by using the WebSphere MQ JMS administration tool or WebSphere MQ Explorer. Alternatively, an application can set the property by calling the `setSendExitInit()` method.

In a similar way, the `RECEXITINIT` property of a `ConnectionFactory` object specifies the user data that is passed to each receive exit, and the `SECXITINIT` property specifies the user data passed to a security exit. You can set these properties by using the WebSphere MQ JMS administration tool or WebSphere MQ Explorer. Alternatively, an application can set the properties by calling the `setReceiveExitInit()` and `setSecurityExitInit()` methods.

Note the following rules when specifying user data that is passed to channel exits:

- If the number of items of user data in a string is more than the number of exits in a sequence, the excess items of user data are ignored.
- If the number of items of user data in a string is less than the number of exits in a sequence, each unspecified item of user data is set to an empty string. Two commas in succession within a string, or a comma at the beginning of a string, also denotes an unspecified item of user data.

If an application uses a client channel definition table (CCDT) to connect to a queue manager, any user data specified in a client connection channel definition is passed to channel exits when they are called. For more information about using a client channel definition table, see [“Using a client channel definition table with IBM WebSphere MQ classes for JMS”](#) on page 875.

Using a client channel definition table with IBM WebSphere MQ classes for JMS

An IBM WebSphere MQ classes for JMS application can use client connection channel definitions that are stored in a client channel definition table (CCDT). You configure a `ConnectionFactory` object to use the CCDT. There are some restrictions on its use.

As an alternative to creating a client connection channel definition by setting certain properties of a `ConnectionFactory` object, a IBM WebSphere MQ classes for JMS application can use client connection channel definitions that are stored in a client channel definition table. These definitions are created by IBM WebSphere MQ Script commands (MQSC) or IBM WebSphere MQ Programmable Command Format (PCF) commands. When the application creates a `Connection` object, IBM WebSphere MQ classes for JMS searches the client channel definition table for a suitable client connection channel definition, and uses the channel definition to start an MQI channel. For more information about client channel definition tables and how to construct one, see [Client channel definition table](#).

To use a client channel definition table, the `CCDTURL` property of a `ConnectionFactory` object must be set to a `URL` object. The `URL` object encapsulates a uniform resource locator (URL) that identifies the name and location of the file containing the client channel definition table and specifies how the file can be accessed. You can set the `CCDTURL` property by using the IBM WebSphere MQ JMS administration tool, or an application can set the property by creating a `URL` object and calling the `setCCDTURL()` method of the `ConnectionFactory` object.

For example, if the file `ccdt1.tab` contains a client channel definition table and is stored on the same system on which the application is running, the application can set the `CCDTURL` property in the following way:

```
java.net.URL chanTab1 = new URL("file:///home/admdata/ccdt1.tab");
factory.setCCDTURL(chanTab1);
```

As another example, suppose the file `ccdt2.tab` contains a client channel definition table and is stored on a system that is different from the one on which the application is running. If the file can be accessed using the FTP protocol, the application can set the `CCDTURL` property in the following way:

```
java.net.URL chanTab2 = new URL("ftp://ftp.server/admdata/ccdt2.tab");
factory.setCCDTURL(chanTab2);
```

In addition to setting the `CCDTURL` property of the `ConnectionFactory` object, the `QMANAGER` property of the same object must be set to one of the following values:

- The name of a queue manager
- An asterisk (*) followed by the name of a queue manager group
- An asterisk (*)
- An empty string, or a string containing all blank characters

These are the same values that can be used for the *QMgrName* parameter on an `MQCONN` call issued by a client application that is using Message Queue Interface (MQI). For more information about the meaning of these values therefore, see `MQCONN`. You can set the `QMANAGER` property by using the WebSphere MQ JMS administration tool or IBM WebSphere MQ Explorer. Alternatively, an application can set the property by calling the `setQueueManager()` method of the `ConnectionFactory` object.

If an application then creates a `Connection` object from the `ConnectionFactory` object, IBM WebSphere MQ classes for JMS accesses the client channel definition table identified by the `CCDTURL` property, uses the `QMANAGER` property to search the table for a suitable client connection channel definition, and then uses the channel definition to start an MQI channel to a queue manager.

Note that the `CCDTURL` and `CHANNEL` properties of a `ConnectionFactory` object cannot both be set when the application calls the `createConnection()` method. If both properties are set, the method throws an exception. The `CCDTURL` or `CHANNEL` property is considered to be set if its value is anything other than null, an empty string, or a string containing all blank characters.

When IBM WebSphere MQ classes for JMS finds a suitable client connection channel definition in the client channel definition table, it uses only the information extracted from the table to start an MQI channel. Any channel related properties of the `ConnectionFactory` object are ignored.

In particular, note the following points if you are using Secure Sockets Layer (SSL):

- An MQI channel uses SSL only if the channel definition extracted from the client channel definition table specifies the name of a CipherSpec supported by IBM WebSphere MQ classes for JMS.
- A client channel definition table also contains information about the location of Lightweight Directory Access Protocol (LDAP) servers that hold certificate revocation lists (CRLs). IBM WebSphere MQ classes for JMS uses only this information to access LDAP servers that hold CRLs.
- A client channel definition table can also contain the location of an Online Certificate Status Protocol (OCSP) responder. IBM WebSphere MQ classes for JMS cannot use the OCSP information in a client channel definition table file. However, you can configure OCSP as described in the section [Using Online Certificate Protocol](#).

For more information about using SSL with a client channel definition table, see [Using the extended transactional client with SSL channels](#).

Note also the following points if you are using channel exits:

- An MQI channel uses only the channel exits and associated user data specified by the channel definition extracted from the client channel definition table.

- A channel definition extracted from a client channel definition table can specify channel exits that are written in Java. This means, for example, that the SCYEXIT parameter on the DEFINE CHANNEL command to create a client connection channel definition can specify the name of a class that implements the WMQSecurityExit interface. Similarly, the SENDEXIT parameter can specify the name of a class that implements the WMQSendExit interface, and the RCVEXIT parameter can specify the name of a class that implements the WMQReceiveExit interface. For more information about how to write a channel exit in Java, see [“Writing channel exits in Java for WebSphere MQ classes for JMS” on page 872.](#)

The use of channel exits written in a language other than Java is also supported. For information about how to specify the SCYEXIT, SENDEXIT, and RCVEXIT parameters on the DEFINE CHANNEL command for channel exits written in another language, see [DEFINE CHANNEL](#).

Automatic JMS client reconnection

Configure your JMS client to reconnect automatically following a network, queue manager, or server failure.

Use the CONNECTIONNAMELIST and CLIENTRECONNECTOPTIONS properties of the MQConnectionFactory class to configure a client connection to automatically reconnect following a connection failure or an administrative request to reconnect client applications after stopping a queue manager.

The full list of connection names in a connectionNameList is only accessible to the set/getconnectionNameList methods that can handle a list of connection names. Methods such as get/setHostname that do not handle lists of names, access the first name in the list.

Automatically reconnectable client connections only become reconnectable once the connection has been established.

Whether an application continues to work correctly after being reconnected automatically depends on its design. Read the related topics to understand how to design reconnectable clients. Some existing clients might work correctly without modification following automatic reconnection.

Automatic client reconnect is not supported by WebSphere MQ classes for Java.

In order to prevent all the clients attached to a failed queue manager from reconnecting simultaneously, the reconnection attempts are delayed by intervals that are partly fixed and partly random.

By default, the reconnection attempts happen at the following intervals:

1. The first attempt is made after an initial delay of one second, plus a random element up to 250 milliseconds.
2. The second attempt is made two seconds, plus a random interval of up to 500 milliseconds, after the first attempt fails.
3. The third attempt is made four seconds, plus a random interval of up to one second, after the second attempt fails.
4. The fourth attempt is made eight seconds, plus a random interval of up to two seconds, after the third attempt fails.
5. The fifth attempt is made 16 seconds, plus a random interval of up to four seconds, after the fourth attempt fails.
6. The sixth attempt, and all subsequent attempts are made 25 seconds, plus a random interval of up to six seconds and 250 milliseconds, after the previous attempt fails.

This process of reconnection continues, until the client is successfully reconnected to the queue manager, or until the maximum reconnection interval elapses.

If you need to increase the default values, in order to more accurately reflect the amount of time required for the queue manager to recover, or the standby queue manager to activate, change the delay values in the MQCLIENT.INI file by using the **ReconDe1ay** attribute.

Related concepts

[Automated client reconnection](#)

Related tasks

[Configuring a client using a configuration file](#)

Sharing a TCP/IP connection in IBM WebSphere MQ classes for JMS

Multiple instances of an MQI channel can be made to share a single TCP/IP connection.

Applications that are running inside the same Java runtime environment, and that use the IBM WebSphere MQ classes for JMS or the IBM WebSphere MQ resource adapter to connect to a queue manager by using the CLIENT transport, can be made to share the same channel instance.

There is a one-to-one relationship between channel instances and TCP/IP connections. One TCP/IP connection is created for every channel instance.

If a channel is defined with the **SHARECNV** parameter set to a value greater than 1, then that number of conversations can share a channel instance. To enable a connection factory or an activation specification to use this function, set the **SHARECONVALLOWED** property to YES.

Every JMS connection and JMS session that is created by a JMS application creates its own conversation with the queue manager.

When an activation specification starts up, the IBM WebSphere MQ classes for JMS resource adapter starts a conversation with the queue manager for the activation specification to use. Every server session in the server session pool that is associated with the activation specification also starts a conversation with the queue manager.

The SHARECNV attribute is a best effort approach to connection sharing. Therefore, when a SHARECNV value greater than 0 is used with the IBM WebSphere MQ classes for JMS, it is not guaranteed that a new connection request will always share an already established connection.

Calculating the number of channel instances

Use the following formulae to determine the maximum number of channel instances that are created by an application:

Activation specifications

Number of channel instances = (**<maxPoolDepth>** + 1) / **<SHARECNV>**

Where **<maxPoolDepth>** is the value of the **maxPoolDepth** property and **<SHARECNV>** is the value of **SHARECNV** property on the channel that is used by the activation specification.

Other JMS applications

Number of channel instances = (**<JMS connections>** + **<JMS sessions>**) / **<SHARECNV>**

Where **<JMS connections>** is the number of connections that are created by the application, where **<JMS sessions>** is the number of JMS sessions that are created by the application, and **<SHARECNV>** is the value of **SHARECNV** property on the channel that is used by the activation specification.

Examples

The following examples show how to use the formulae to calculate the number of channel instances that are created on a queue manager by applications by using either the IBM WebSphere MQ classes for JMS or the IBM WebSphere MQ classes for JMS resource adapter.

JMS application example

A JMS application connection connects to a queue manager by using the CLIENT transport, and creates a JMS connection and three JMS sessions. The channel that the application is using to connect to the queue manager has the **SHARECNV** property set to the value of 10. When the application is running, there are four conversations between the application and the queue manager and one channel instance. The four conversations all share the channel instance.

Activation specification example

An activation specification connects to a queue manager by using the CLIENT transport. The activation specification is configured with the **maxPoolDepth** property set to 10. The channel that the activation specification is configured to use has the **SHARECNV** property set to 10. When the activation specification is running, and processing 10 messages concurrently, the number of conversations between the activation specification and the queue manager is 11 (10 conversations for the server sessions, and one for the activation specification). The number of channel instances that are used by the activation specification is 2.

Activation specification example

An activation specification connects to a queue manager by using the CLIENT transport. The activation specification is configured with the **maxPoolDepth** property set to 5. The channel that the activation specification is configured to use has the **SHARECNV** property set to 0. When the activation specification is running, and processing 5 messages concurrently, the number of conversations between the activation specification and the queue manager is 6 (five conversations for the server sessions, and one for the activation specification). The number of channel instances that are used by the activation specification is 6, because the **SHARECNV** property on the channel is set to 0, every conversation uses its own channel instance.

Specifying a range of ports for client connections in WebSphere MQ classes for JMS

Use the LOCALADDRESS property to specify a range of ports that your application can bind to.

When a WebSphere MQ classes for JMS application attempts to connect to a WebSphere MQ queue manager in client mode, a firewall might allow only those connections that originate from specified ports or a range of ports. In this situation, you can use the LOCALADDRESS property of a ConnectionFactory, QueueConnectionFactory, or TopicConnectionFactory object to specify a port, or a range of ports, that the application can bind to.

You can set the LOCALADDRESS property by using the WebSphere MQ JMS administration tool, or by calling the setLocalAddress() method in a JMS application. Here is an example of setting the property from within an application:

```
mqConnectionFactory.setLocalAddress("192.0.2.0(2000,3000)");
```

When the application connects to a queue manager subsequently, the application binds to a local IP address and port number in the range 192.0.2.0(2000) to 192.0.2.0(3000).

In a system with more than one network interface, you can also use the LOCALADDRESS property to specify which network interface must be used for a connection.

For a real-time connection to a broker, the LOCALADDRESS property is relevant only when multicast is used. In this case, you can use the property to specify which local network interface must be used for a connection, but the value of the property must not contain a port number, or a range of port numbers.

Connection errors might occur if you restrict the range of ports. If an error occurs, a JMSException is thrown with an embedded MQException that contains the WebSphere MQ reason code MQRC_Q_MGR_NOT_AVAILABLE and the following message:

```
Socket connection attempt refused due to LOCAL_ADDRESS_PROPERTY restrictions
```

An error might occur if all the ports in the specified range are in use, or if the specified IP address, host name, or port number is not valid (a negative port number, for example).

Because WebSphere MQ classes for JMS might create connections other than those required by an application, always consider specifying a range of ports. In general, every session created by an application requires one port and WebSphere MQ classes for JMS might require three or four additional ports. If a connection error does occur, increase the range of ports.

Connection pooling, which is used by default in WebSphere MQ classes for JMS, might have an effect on the speed at which ports can be reused. As a result, a connection error might occur while ports are being freed.

Channel compression in WebSphere MQ classes for JMS

A WebSphere MQ classes for JMS application can use WebSphere MQ facilities to compress a message header or data.

Compressing the data that flows on a WebSphere MQ channel can improve the performance of the channel and reduce network traffic. Using function supplied with WebSphere MQ, you can compress the data that flows on message channels and MQI channels. On either type of channel, you can compress header data and message data independently of each other. By default, no data is compressed on a channel.

A WebSphere MQ classes for JMS application specifies the techniques that can be used for compressing header or message data on a connection by creating a `java.util.Collection` object. Each compression technique is an `Integer` object in the collection, and the order in which the application adds the compression techniques to the collection is the order in which the compression techniques are negotiated with the queue manager when the application creates the connection. The application can then pass the collection to a `ConnectionFactory` object by calling the `setHdrCompList()` method, for header data, or the `setMsgCompList()` method, for message data. When the application is ready, it can create the connection.

The following code fragments illustrate the approach described. The first code fragment shows you how to implement header data compression:

```
Collection headerComp = new Vector();
headerComp.add(new Integer(WMQConstants.WMQ_COMPHDR_SYSTEM));
.
.
.
((MQConnectionFactory) cf).setHdrCompList(headerComp);
.
.
.
connection = cf.createConnection();
```

The second code fragment shows you how to implement message data compression:

```
Collection msgComp = new Vector();
msgComp.add(new Integer(WMQConstants.WMQ_COMPMSG_RLE));
msgComp.add(new Integer(WMQConstants.WMQ_COMPMSG_ZLIBHIGH));
.
.
.
((MQConnectionFactory) cf).setMsgCompList(msgComp);
.
.
.
connection = cf.createConnection();
```

In the second example, the compression techniques are negotiated in the order RLE, then ZLIBHIGH, when the connection is created. The compression technique that is selected cannot be changed during the lifetime of the `Connection` object. To use compression on a connection, the `setHdrCompList()` and the `setMsgCompList()` methods must be called before creating the `Connection` object.

Putting messages asynchronously in IBM WebSphere MQ classes for JMS

Normally, when an application sends messages to a destination, the application has to wait for the queue manager to confirm that it has processed the request. You can improve messaging performance in some circumstances by choosing instead to put messages asynchronously. When an application puts a message asynchronously, the queue manager does not return the success or failure of each call, but you can instead check for errors periodically.

Whether a destination returns control to the application, without determining whether the queue manager has received the message safely, depends upon the following properties:

- The JMS Destination Property `PUTSYNCAALLOWED` (short name - `PAALD`).

`PUTSYNCAALLOWED` controls whether JMS applications can put messages asynchronously, if the underlying queue or topic that the JMS Destination represents, allows this option.

- The IBM WebSphere MQ queue or topic property `DEFPRESP` (Default put response type).

`DEFPRESP` specifies whether applications that put messages to the queue, or publish messages to the topic, can make use of the asynchronous put functionality.

The following table shows the possible values for the `PUTASYNCALLOWED` and `DEFPRESP` properties, and what values are required for the asynchronous put functionality to be enabled:

Table 131. PUTASYNCALLOWED and DEFPRESP properties determining if messages are put asynchronously.

WebSphere MQ queue property	PUTASYNCALLOWED = NO	PUTASYNCALLOWED = YES	PUTASYNCALLOWED = AS_DEST or AS_Q_DEF or AS_T_DEF
DEFPRESP=SYNC	Asynchronous put functionality not enabled	Asynchronous put functionality enabled	Asynchronous put functionality not enabled
DEFPRESP=ASYN	Asynchronous put functionality not enabled	Asynchronous put functionality enabled	Asynchronous put functionality enabled

For messages sent in a transacted session, the application ultimately determines whether the queue manager has received the messages safely when it calls `commit()`.

If an application sends persistent messages within a transacted session, and one or more of the messages are not received safely, the transaction fails to commit and produces an exception. However, if an application sends nonpersistent messages within a transacted session, and one or more of the messages are not received safely, the transaction commits successfully. The application does not receive any feedback that the nonpersistent messages did not arrive safely.

For nonpersistent messages sent in a session that is not transacted, the `SENDCHECKCOUNT` property of the *ConnectionFactory* object specifies how many messages are to be sent, before IBM WebSphere MQ classes for JMS checks that the queue manager has received the messages safely.

If a check discovers that one or more messages were not received safely, and the application has registered an exception listener with the connection, IBM WebSphere MQ classes for JMS calls the `onException()` method of the exception listener to pass a JMS exception to the application.

The JMS exception has an error code of `JMSWMQ0028` and this code displays the following message:

At least one asynchronous put message failed or gave a warning.

The JMS exception also has a linked exception that provides more details. The default value of the `SENDCHECKCOUNT` property is zero, which means that no such checks are made.

This optimization is of most benefit to an application that connects to a queue manager in client mode, and needs to send a sequence of messages in rapid succession, but does not require immediate feedback from the queue manager for each message sent. However, an application can still use this optimization even if it connects to a queue manager in bindings mode, but the expected performance benefit is not as great.

Using read ahead with WebSphere MQ classes for JMS

The read ahead functionality that is provided by WebSphere MQ allows non-persistent messages that are received outside of a transaction to be sent to the IBM WebSphere MQ classes for JMS before an application requests them. The IBM WebSphere MQ classes for JMS store the messages in an internal buffer, and pass the messages to the application when the application asks for them.

IBM WebSphere MQ classes for JMS applications that use `MessageConsumers` or `MessageListeners` to receive messages from a destination outside of a transaction can use the read ahead functionality. Using read ahead allows applications that use these objects to benefit from improved performance when they receive messages.

Whether an application that uses `MessageConsumers` or `MessageListeners` can use read ahead depends upon the following properties:

- The JMS Destination Property READAHEADALLOWED (short name - RAALD). READAHEADALLOWED controls whether JMS applications can use read ahead when getting or browsing non-persistent messages outside of a transaction, if the underlying queue or topic that the JMS Destination represents, allows this option.
- The IBM WebSphere MQ queue or topic property DEFREADA (Default read ahead). DEFREADA specifies whether applications that are receiving or browsing non-persistent messages outside of a transaction can use read ahead.

The following table shows the possible values for the READAHEADALLOWED and DEFREADA properties, and what values are required for the read ahead functionality to be enabled:

<i>Table 132. READAHEADALLOWED and DEFREADA properties determining if read ahead is used when receiving or browsing non-persistent messages outside of a transaction.</i>			
WebSphere MQ destination property	READAHEADALLOWED = YES	READAHEADALLOWED = NO	AS_DEST or AS_Q_DEF or AS_T_DEF
WebSphere MQ queue property			
DEFREADA = NO	Read ahead functionality enabled	Read ahead functionality not enabled	Read ahead functionality not enabled
DEFREADA = YES	Read ahead functionality enabled	Read ahead functionality not enabled	Read ahead functionality enabled
DEFREADA = DISABLED	Read ahead functionality not enabled	Read ahead functionality not enabled	Read ahead functionality not enabled

If the read ahead functionality is enabled, when a `MessageConsumer` or `MessageListener` is created by an application, the IBM WebSphere MQ classes for JMS create an internal buffer for the destination that the `MessageConsumer` or `MessageListener` is monitoring. There is one internal buffer for each `MessageConsumer` or `MessageListener`. The queue manager starts sending non-persistent messages to the IBM WebSphere MQ classes for JMS when the application calls one of the following methods:

- `MessageConsumer.receive()`
- `MessageConsumer.receive(long timeout)`
- `MessageConsumer.receiveNoWait()`
- `Session.setMessageListener(MessageListener listener)`

The IBM WebSphere MQ classes for JMS automatically returns the first message back to the application, by the method call that the application has made. The other non-persistent messages are stored by the IBM WebSphere MQ classes for JMS in the internal buffer that was created for the destination. When the application requests the next message to process, the IBM WebSphere MQ classes for JMS will return the next message in the internal buffer.

The IBM WebSphere MQ classes for JMS requests more non-persistent messages from the queue manager when the internal buffer is empty.

The internal buffer that is used by the IBM WebSphere MQ classes for JMS is deleted when an application closes a `MessageConsumer`, or the JMS Session that a `MessageListener` is associated with.

For `MessageConsumers`, any unprocessed messages in the internal buffer is lost.

When using `MessageListeners`, what happens to the messages in the internal buffer depends upon the JMS destination property READAHEADCLOSEPOLICY (short name - RACP). The default value of the property is DELIVER_ALL, which means that the JMS session that was used to create the `MessageListener` is not closed until all of the messages in the internal buffer are delivered to the application. If the property is set to DELIVER_CURRENT, then the JMS session will be closed after the current message has been processed by the application and all of the remaining messages in the internal buffer are discarded.

Retained publications in WebSphere MQ classes for JMS

A WebSphere MQ classes for JMS client can be configured to use retained publications.

A publisher can specify that a copy of a publication must be retained so that it can be sent to future subscribers who register an interest in the topic. This is done in WebSphere MQ classes for JMS by setting the integer property `JMS_IBM_RETAIN` to the value 1. Constants have been defined for these values in the `com.ibm.msg.client.jms.JmsConstants` interface. For example, if you have created a message `msg`, to set it as a retained publication use the following code:

```
// set as a retained publication
msg.setIntProperty(JmsConstants.JMS_IBM_RETAIN, JmsConstants.RETAIN_PUBLICATION);
```

You can now send the message as normal. `JMS_IBM_RETAIN` can also be queried in a received message. It is therefore possible to query whether a received message is a retained publication.

XA support in WebSphere MQ classes for JMS

JMS supports XA-compliant transactions in bindings and client modes with a supported transaction manager.

If you require XA functionality in an application server environment, you must configure your application appropriately. Refer to your application server's own documentation for information about how to configure applications to use distributed transactions.

Using a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker

A WebSphere MQ classes for JMS application can use a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker for publish/subscribe messaging. Both the broker and WebSphere MQ classes for JMS must be configured to enable a real-time connection.

When an application uses a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker, the application and the broker exchange messages using WebSphere MQ Real-Time Transport. Depending on the configuration, messages can also be delivered to the application using WebSphere MQ Multicast Transport.

For information about how an application can connect to a WebSphere MQ queue manager and use WebSphere MQ Enterprise Transport to exchange messages with a broker of WebSphere Event Broker or WebSphere Message Broker, see the documentation for previous releases of WebSphere MQ classes for JMS. Note that, in order to use WebSphere MQ Enterprise Transport, an application must connect to a queue manager using a connection factory running in WebSphere MQ messaging provider migration mode.

Configuring a broker of WebSphere Event Broker or WebSphere Message Broker for a real-time connection

For a WebSphere MQ classes for JMS application to use a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker, you must configure the broker by creating and deploying a message flow to read messages from the TCP/IP port on which the broker is listening and publish the messages. Depending on your requirements, you might need to configure the broker in additional ways.

To configure the broker, you must create and deploy one of the following message flows:

- A message flow that contains a `Real-timeOptimizedFlow` message processing node
- A message flow that contains a `Real-timeInput` message processing node and a `Publication` message processing node

You must configure the `Real-timeOptimizedFlow` or `Real-timeInput` node to listen on the TCP/IP port used for real-time connections. By default, the port number for real-time connections is 1506.

You must also configure the broker if you have any of the following requirements:

- If you want the application to connect to the broker using Secure Sockets Layer (SSL) authentication

- If you want the application to connect to the broker using HTTP tunnelling
- If you want messages to be delivered to a message consumer using multicast

For information about how to configure a broker, see the *WebSphere Event Broker product documentation* or *WebSphere Message Broker product documentation*.

Configuring WebSphere MQ classes for JMS for a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker

For a WebSphere MQ classes for JMS application to use a real-time connection to a broker of WebSphere Event Broker or WebSphere Message Broker, WebSphere MQ classes for JMS must be configured by setting certain properties of the connection factory. Depending on your requirements, WebSphere MQ classes for JMS might need to be configured in additional ways.

To configure WebSphere MQ classes for JMS, the following properties of the connection factory must be set:

- The **TRANSPORT** property must be set to **DIRECT**.

However, for an application to connect using HTTP tunnelling, the **TRANSPORT** property must be set to **DIRECTHTTP** instead. See [“Using HTTP tunnelling”](#) on page 885.

- The **HOSTNAME** property must be set to the host name or IP address of the system on which the broker is running.
- The **PORT** property must be set to the number of the port on which the broker is listening for real-time connections.

An application can set these properties dynamically at run time by using the IBM JMS extensions or the WebSphere MQ JMS extensions. Alternatively, if the connection factory is an administered object, an administrator can set these properties by using the WebSphere MQ JMS administration tool or WebSphere MQ Explorer.

For information about properties, and the methods used by applications to set their values, see [Properties of IBM WebSphere MQ classes for JMS objects](#). For information about how to use the WebSphere MQ JMS administration tool, see [“Using the WebSphere MQ JMS administration tool”](#) on page 893. For information about how to use WebSphere MQ Explorer, see the help provided with WebSphere MQ Explorer.

If you have any of the following requirements, WebSphere MQ classes for JMS requires additional configuration:

- If you want an application to connect to the broker using Secure Sockets Layer (SSL) authentication
- If you want an application to connect to the broker using HTTP tunnelling
- If you want an application to connect to the broker through a proxy server
- If you want messages to be delivered to a message consumer using multicast

The following sections describe how to configure WebSphere MQ classes for JMS for each of these requirements.

Using Secure Sockets Layer (SSL) authentication

SSL authentication can be used on a real-time connection to a broker. Only authentication is supported for this type of connection. You cannot use SSL to encrypt and decrypt the message data that flows between the application and the broker or to detect tampering of the data.

Note the difference between this situation and that when an application connects to a queue manager in client mode. In the latter case, you can use the WebSphere MQ SSL support to encrypt and decrypt the message data that flows between the application and the queue manager and to detect tampering of the data, as well as to provide authentication.

If you want to protect message data on a real-time connection to a broker, you can use the function provided by the broker instead. You can assign a quality of protection (QoP) value to each topic with messages you want to protect. You can therefore select a different level of message protection for each

topic. For more information about the message protection provided by a broker, see the *WebSphere Event Broker product documentation* or *WebSphere Message Broker product documentation*.

To use SSL authentication on a real-time connection to a broker, the `DIRECTAUTH` property of the connection factory must be set to `CERTIFICATE`.

If you want to use SSL for mutual authentication, the Authentication Protocol Type property of the broker must specify the option `R` for symmetric SSL. If you want to use SSL only for authenticating the broker, the Authentication Protocol Type property of the broker must specify the option `S` for asymmetric SSL. But, in this case, the application must connect to the broker by calling `createConnection()` with a user ID and password as parameters, as in the following example:

```
factory.createConnection("user1", "user1pw");
```

The broker then uses the user ID and password, instead of SSL, to authenticate the application. For more information about how to configure the broker for SSL authentication, see the *WebSphere Event Broker product documentation* or *WebSphere Message Broker product documentation*.

Notes:

1. The value of the `DIRECTAUTH` property determines whether SSL authentication is used on a real-time connection to a broker, not the value of the `SSLCIPHERSUITE` property.
2. When SSL authentication is used on a real-time connection to a broker, the `SSLPEERNAME` and `SSLCRL` properties are used to perform the same checks as those performed when an application connects to a queue manager in client mode.
3. WebSphere MQ classes for JMS can use the same Java Secure Socket Extension (JSSE) keystore and truststore configuration to provide the SSL support in either of the following situations:
 - When an application uses a real-time connection to a broker
 - When an application connects to a queue manager in client mode

Using HTTP tunnelling

A WebSphere MQ classes for JMS application can connect to a broker using HTTP tunnelling, which means that the application connects to the broker using the HTTP protocol as though connecting to a website.

To use HTTP tunnelling on a real-time connection to a broker, the `TRANSPORT` property of the connection factory must be set to `DIRECTHTTP`.

HTTP tunnelling cannot be used in conjunction with SSL authentication, connecting through a proxy server, or delivering messages using multicast. The supported version of the HTTP protocol is 1.0. HTTP version 1.1 is not supported.

Connecting through a proxy server

A WebSphere MQ classes for JMS application can use a real-time connection to a broker by connecting through a proxy server. WebSphere MQ classes for JMS connects directly to the proxy server and uses the Internet protocol defined in RFC 2817 to ask the proxy server to forward the connection request to the broker.

To connect to a broker through a proxy server, the following properties of the connection factory must be set:

- The `PROXYHOSTNAME` property must be set to the host name or IP address of the system on which the proxy server is running.
- The `PROXYPORT` property must be set to the number of the port on which the proxy server is listening.

If the `PROXYHOSTNAME` property is not set, or is set to the empty string, WebSphere MQ classes for JMS attempts to connect directly to the broker using only the `HOSTNAME` and `PORT` properties, and does not attempt to connect through a proxy server.

Delivering messages using multicast

Using a real-time connection to a broker, messages can be delivered to a message consumer using multicast.

To enable multicast, the MULTICAST property of the Topic object must be set to the required multicast option. Alternatively, if the MULTICAST property of the Topic object is set to ASCF, the MULTICAST property of the connection factory must be set to the required multicast option.

WebSphere MQ classes for JMS supports both the Packet Transfer Layer (PTL) and the Pragmatic General Multicast (PGM) multicast protocols, and includes support for both implementations of the PGM protocol, PGM/IP and PGM UDP encapsulated. However, PGM/IP support is available only on the following platforms:

- AIX (32-bit only)
- Linux (x86 platform)
- Linux (zSeries platform, 32-bit only)
- Solaris SPARC (32-bit only)
- Windows (32-bit only)
- z/OS

WebSphere MQ classes for JMS Application Server Facilities

This topic describes how WebSphere MQ classes for JMS implements the ConnectionConsumer class and advanced functionality in the Session class. It also summarizes the function of a server session pool.

WebSphere MQ classes for JMS supports the Application Server Facilities (ASF) that are specified in the *Java Message Service Specification, Version 1.1* (see Sun's Java website at <https://java.sun.com>). This specification identifies three roles within this programming model:

- **The JMS provider** supplies ConnectionConsumer and advanced Session functionality.
- **The application server** supplies ServerSessionPool and ServerSession functionality.
- **The client application** uses the functionality that the JMS provider and application server supply.

The information in this topic does not apply if an application uses a real-time connection to a broker.

The JMS ConnectionConsumer

The ConnectionConsumer interface provides a high-performance method to deliver messages concurrently to a pool of threads.

The JMS specification enables an application server to integrate closely with a JMS implementation by using the ConnectionConsumer interface. This feature provides concurrent processing of messages. Typically, an application server creates a pool of threads, and the JMS implementation makes messages available to these threads. A JMS-aware application server (such as WebSphere Application Server) can use this feature to provide high-level messaging functionality, such as message driven beans.

Normal applications do not use the ConnectionConsumer, but expert JMS clients might use it. For such clients, the ConnectionConsumer provides a high-performance method to deliver messages concurrently to a pool of threads. When a message arrives on a queue or a topic, JMS selects a thread from the pool and delivers a batch of messages to it. To do this, JMS runs an associated MessageListener's onMessage() method.

You can achieve the same effect by constructing multiple Session and MessageConsumer objects, each with a registered MessageListener. However, the ConnectionConsumer provides better performance, less use of resources, and greater flexibility. In particular, fewer Session objects are required.

Planning an application with ASF

This section tells you how to plan an application including:

- [“General principles for point-to-point messaging using ASF” on page 887](#)
- [“General principles for publish/subscribe messaging using ASF” on page 888](#)
- [“Removing messages from the queue in ASF” on page 888](#)
- Handling poison messages in ASF. See [“Handling poison messages in IBM WebSphere MQ classes for JMS” on page 850](#).

General principles for point-to-point messaging using ASF

Use this topic for general information about point-to-point messaging using ASF.

When an application creates a `ConnectionConsumer` from a `QueueConnection` object, it specifies a JMS queue object and a selector string. The `ConnectionConsumer` then begins to provide messages to sessions in the associated `ServerSessionPool`. Messages arrive on the queue, and if they match the selector, they are delivered to sessions in the associated `ServerSessionPool`.

In WebSphere MQ terms, the queue object refers to either a QLOCAL or a QALIAS on the local queue manager. If it is a QALIAS, that QALIAS must refer to a QLOCAL. The fully resolved WebSphere MQ QLOCAL is known as the *underlying QLOCAL*. A `ConnectionConsumer` is said to be *active* if it is not closed and its parent `QueueConnection` is started.

It is possible for multiple `ConnectionConsumers`, each with different selectors, to run against the same underlying QLOCAL. To maintain performance, unwanted messages must not accumulate on the queue. Unwanted messages are those for which no active `ConnectionConsumer` has a matching selector. You can set the `QueueConnectionFactory` so that these unwanted messages are removed from the queue (for details, see [“Removing messages from the queue in ASF” on page 888](#)). You can set this behavior in one of two ways:

- Use the JMS administration tool to set the `QueueConnectionFactory` to `MRET(NO)`.
- In your program, use:

```
MQQueueConnectionFactory.setMessageRetention(WMQConstants.WMQ_MRET_NO)
```

If you do not change this setting, the default is to retain such unwanted messages on the queue.

When you set up the WebSphere MQ queue manager, consider the following points:

- The underlying QLOCAL must be enabled for shared input. To do this, use the following MQSC command:

```
ALTER QLOCAL(your.qlocal.name) SHARE GET(ENABLED)
```

- Your queue manager must have an enabled dead-letter queue. If a `ConnectionConsumer` experiences a problem when it puts a message on the dead-letter queue, message delivery from the underlying QLOCAL stops. To define a dead-letter queue, use:

```
ALTER QMGR DEADQ(your.dead.letter.queue.name)
```

- The user that runs the `ConnectionConsumer` must have authority to perform `MQOPEN` with `MQOO_SAVE_ALL_CONTEXT` and `MQOO_PASS_ALL_CONTEXT`. For details, see the WebSphere MQ documentation for your specific platform.
- If unwanted messages are left on the queue, they degrade the system performance. Therefore, plan your message selectors so that between them, the `ConnectionConsumers` will remove all messages from the queue.

For details about MQSC commands, see the [MQSC reference](#).

General principles for publish/subscribe messaging using ASF

ConnectionConsumers receive messages for a specified Topic. A ConnectionConsumer can be durable or non-durable. You must specify which queue or queues the ConnectionConsumer uses.

When an application creates a ConnectionConsumer from a TopicConnection object, it specifies a Topic object and a selector string. The ConnectionConsumer then begins to receive messages that match the selector on that Topic, including any retained publications for the topic subscribed to.

Alternatively, an application can create a durable ConnectionConsumer that is associated with a specific name. This ConnectionConsumer receives messages that have been published on the Topic since the durable ConnectionConsumer was last active. It receives all such messages that match the selector on the Topic. However, if the ConnectionConsumer is using read-ahead, it can lose nonpersistent messages that are in the client buffer when it closes.

If WebSphere MQ classes for JMS is in WebSphere MQ messaging provider migration mode, a separate queue is used for non-durable ConnectionConsumer subscriptions. The CCSUB configurable option on the TopicConnectionFactory specifies the queue to use. Normally, the CCSUB specifies a single queue for use by all ConnectionConsumers that use the same TopicConnectionFactory. However, it is possible to make each ConnectionConsumer generate a temporary queue by specifying a queue name prefix followed by an asterisk (*).

If WebSphere MQ classes for JMS is in WebSphere MQ messaging provider migration mode, the CCDSUB property of the Topic specifies the queue to use for durable subscriptions. Again, this can be a queue that already exists or a queue name prefix followed by an asterisk (*). If you specify a queue that already exists, all durable ConnectionConsumers that subscribe to the Topic use this queue. If you specify a queue name prefix followed by an asterisk (*), a queue is generated the first time that a durable ConnectionConsumer is created with a particular name. This queue is reused later when a durable ConnectionConsumer is created with the same name.

When you set up the WebSphere MQ queue manager, consider the following points:

- Your queue manager must have an enabled dead-letter queue. If a ConnectionConsumer experiences a problem when it puts a message on the dead-letter queue, message delivery from the underlying QLOCAL stops. To define a dead-letter queue, use:

```
ALTER QMGR DEADQ(your.dead.letter.queue.name)
```

- The user that runs the ConnectionConsumer must have authority to perform MQOPEN with MQOO_SAVE_ALL_CONTEXT and MQOO_PASS_ALL_CONTEXT. For details, see the WebSphere MQ documentation for your platform.
- You can optimize performance for an individual ConnectionConsumer by creating a separate, dedicated, queue for it. This is at the cost of extra resource usage.

Removing messages from the queue in ASF

When an application uses ConnectionConsumers, JMS might need to remove messages from the queue in a number of situations.

These situations are as follows:

Badly formatted message

A message might arrive that JMS cannot parse.

Poison message

A message might reach the backout threshold, but the ConnectionConsumer fails to requeue it on the backout queue.

No interested ConnectionConsumer

For point-to-point messaging, when the QueueConnectionFactory is set so that it does not retain unwanted messages, a message arrives that is unwanted by any of the ConnectionConsumers.

In these situations, the ConnectionConsumer attempts to remove the message from the queue. The disposition options in the report field of the message's MQMD set the exact behavior. These options are:

MQRO_DEAD_LETTER_Q

The message is requeued to the queue manager's dead-letter queue. This is the default.

MQRO_DISCARD_MSG

The message is discarded.

The ConnectionConsumer also generates a report message, and this also depends on the report field of the message's MQMD. This message is sent to the message's ReplyToQ on the ReplyToQmgr. If there is an error while the report message is being sent, the message is sent to the dead-letter queue instead. The exception report options in the report field of the message's MQMD set details of the report message. These options are:

MQRO_EXCEPTION

A report message is generated that contains the MQMD of the original message. It does not contain any message body data.

MQRO_EXCEPTION_WITH_DATA

A report message is generated that contains the MQMD, any MQ headers, and 100 bytes of body data.

MQRO_EXCEPTION_WITH_FULL_DATA

A report message is generated that contains all data from the original message.

default

No report message is generated.

When report messages are generated, the following options are honored:

- MQRO_NEW_MSG_ID
- MQRO_PASS_MSG_ID
- MQRO_COPY_MSG_ID_TO_CORREL_ID
- MQRO_PASS_CORREL_ID

If a poison message cannot be requeued, perhaps because the dead-letter queue is full or authorization is wrongly specified, what happens depends on the persistence of the message. If the message is nonpersistent, the message is discarded and no report message is generated. If the message is persistent, delivery of messages to all connection consumers listening on that destination stops. Such connection consumers must be closed and the problem resolved before they can be re-created and message delivery restarted.

It is important to define a dead-letter queue, and to check it regularly to ensure that no problems occur. Particularly, ensure that the dead-letter queue does not reach its maximum depth, and that its maximum message size is large enough for all messages.

When a message is requeued to the dead-letter queue, it is preceded by a WebSphere MQ dead-letter header (MQDLH). See [MQDLH - Dead-letter header](#) for details about the format of the MQDLH. You can identify messages that a ConnectionConsumer has placed on the dead-letter queue, or report messages that a ConnectionConsumer has generated, by the following fields:

- PutApplType is MQAT_JAVA (0x1C)
- PutApplName is "MQ JMS ConnectionConsumer"

These fields are in the MQDLH of messages on the dead-letter queue, and the MQMD of report messages. The feedback field of the MQMD, and the Reason field of the MQDLH, contain a code describing the error. For details about these codes, see ["Reason and feedback codes in ASF" on page 890](#). Other fields are as described in [MQDLH - Dead-letter header](#).

Handling poison messages in ASF

Within the Application Server Facilities, poison message handling is handled slightly differently to elsewhere in WebSphere MQ classes for JMS.

For information about poison message handling in WebSphere MQ classes for JMS, see ["Handling poison messages in IBM WebSphere MQ classes for JMS" on page 850](#).

When you use Application Server Facilities (ASF), the `ConnectionConsumer`, rather than the `MessageConsumer`, processes poison messages. The `ConnectionConsumer` requeues messages according to the `BackoutThreshold` and `BackoutRequeueQName` properties of the queue.

When an application uses `ConnectionConsumers`, the circumstances in which a message is backed out depend on the session that the application server provides:

- When the session is non-transacted, with `AUTO_ACKNOWLEDGE` or `DUPS_OK_ACKNOWLEDGE`, a message is backed out only after a system error, or if the application terminates unexpectedly.
- When the session is non-transacted with `CLIENT_ACKNOWLEDGE`, unacknowledged messages can be backed out by the application server calling `Session.recover()`.

Typically, the client implementation of `MessageListener` or the application server calls `Message.acknowledge()`. `Message.acknowledge()` acknowledges all messages delivered on the session so far.

- When the session is transacted, unacknowledged messages can be backed out by the application server calling `Session.rollback()`.
- If the application server supplies an `XASession`, messages are committed or backed out depending on a distributed transaction. The application server takes responsibility for completing the transaction.

The embedded JMS provider in WebSphere Application Server, Version 5.0 and Version 5.1 handles poison messages in a different way to that just described for WebSphere MQ classes for JMS. For information about how the embedded JMS provider handles poison messages, see the relevant WebSphere Application Server product documentation.

Error handling

This section covers various aspects of error handling, including [“Recovering from error conditions in the ASF” on page 890](#) and [“Reason and feedback codes in ASF” on page 890](#).

Recovering from error conditions in the ASF

If a `ConnectionConsumer` experiences a serious error, message delivery to all `ConnectionConsumers` with an interest in the same `QLOCAL` stops. When this occurs, any `ExceptionListener` that is registered with the affected `Connection` is notified. There are two ways in which an application can recover from these error conditions.

Typically, a serious error of this nature occurs if the `ConnectionConsumer` cannot requeue a message to the dead-letter queue, or it experiences an error when reading messages from the `QLOCAL`.

Because any `ExceptionListener` that is registered with the affected `Connection` is notified, you can use them to identify the cause of the problem. In some cases, the system administrator must intervene to resolve the problem.

Use one of the following techniques to recover from these error conditions:

- Call `close()` on all affected `ConnectionConsumers`. The application can create new `ConnectionConsumers` only after all affected `ConnectionConsumers` are closed and any system problems are resolved.
- Call `stop()` on all affected `Connections`. After all `Connections` are stopped and any system problems are resolved, the application can `start()` its `Connections` successfully.

Reason and feedback codes in ASF

Use reason and feedback codes to determine the cause of an error. Common reason codes generated by the `ConnectionConsumer` are given here.

To determine the cause of an error, use the following information:

- The feedback code in any report messages
- The reason code in the `MQDLH` of any messages in the dead-letter queue

`ConnectionConsumers` generate the following reason codes.

MQRC_BACKOUT_THRESHOLD_REACHED (0x93A; 2362)

Cause

The message has reached the Backout Threshold defined on the QLOCAL, but no Backout Queue is defined.

On platforms where you cannot define the Backout Queue, the message has reached the JMS-defined backout threshold of 20.

Action

If this is not wanted, define the Backout Queue for the relevant QLOCAL. Also look for the cause of the multiple backouts.

MQRC_MSG_NOT_MATCHED (0x93B; 2363)

Cause

In point-to-point messaging, there is a message that does not match any of the selectors for the ConnectionConsumers monitoring the queue. To maintain performance, the message is requeued to the dead-letter queue.

Action

To avoid this situation, ensure that ConnectionConsumers using the queue provide a set of selectors that deal with all messages, or set the QueueConnectionFactory to retain messages.

Alternatively, investigate the source of the message.

MQRC_JMS_FORMAT_ERROR (0x93C; 2364)

Cause

JMS cannot interpret the message on the queue.

Action

Investigate the origin of the message. JMS normally delivers messages of an unexpected format as a BytesMessage or TextMessage. Occasionally, this fails if the message is very badly formatted.

Other codes that appear in these fields are caused by a failed attempt to requeue the message to a Backout Queue. In this situation, the code describes the reason that the requeue failed. To diagnose the cause of these errors, refer to [API reason codes](#).

If the report message cannot be put on the ReplyToQ, it is put on the dead-letter queue. In this situation, the feedback field of the MQMD is completed as described in this topic. The reason field in the MQDLH explains why the report message could not be placed on the ReplyToQ.

The function of a server session pool in AFS

This topic summarizes the function of a server session pool.

[Figure 165 on page 892](#) summarizes the principles of ServerSessionPool and ServerSession functionality.

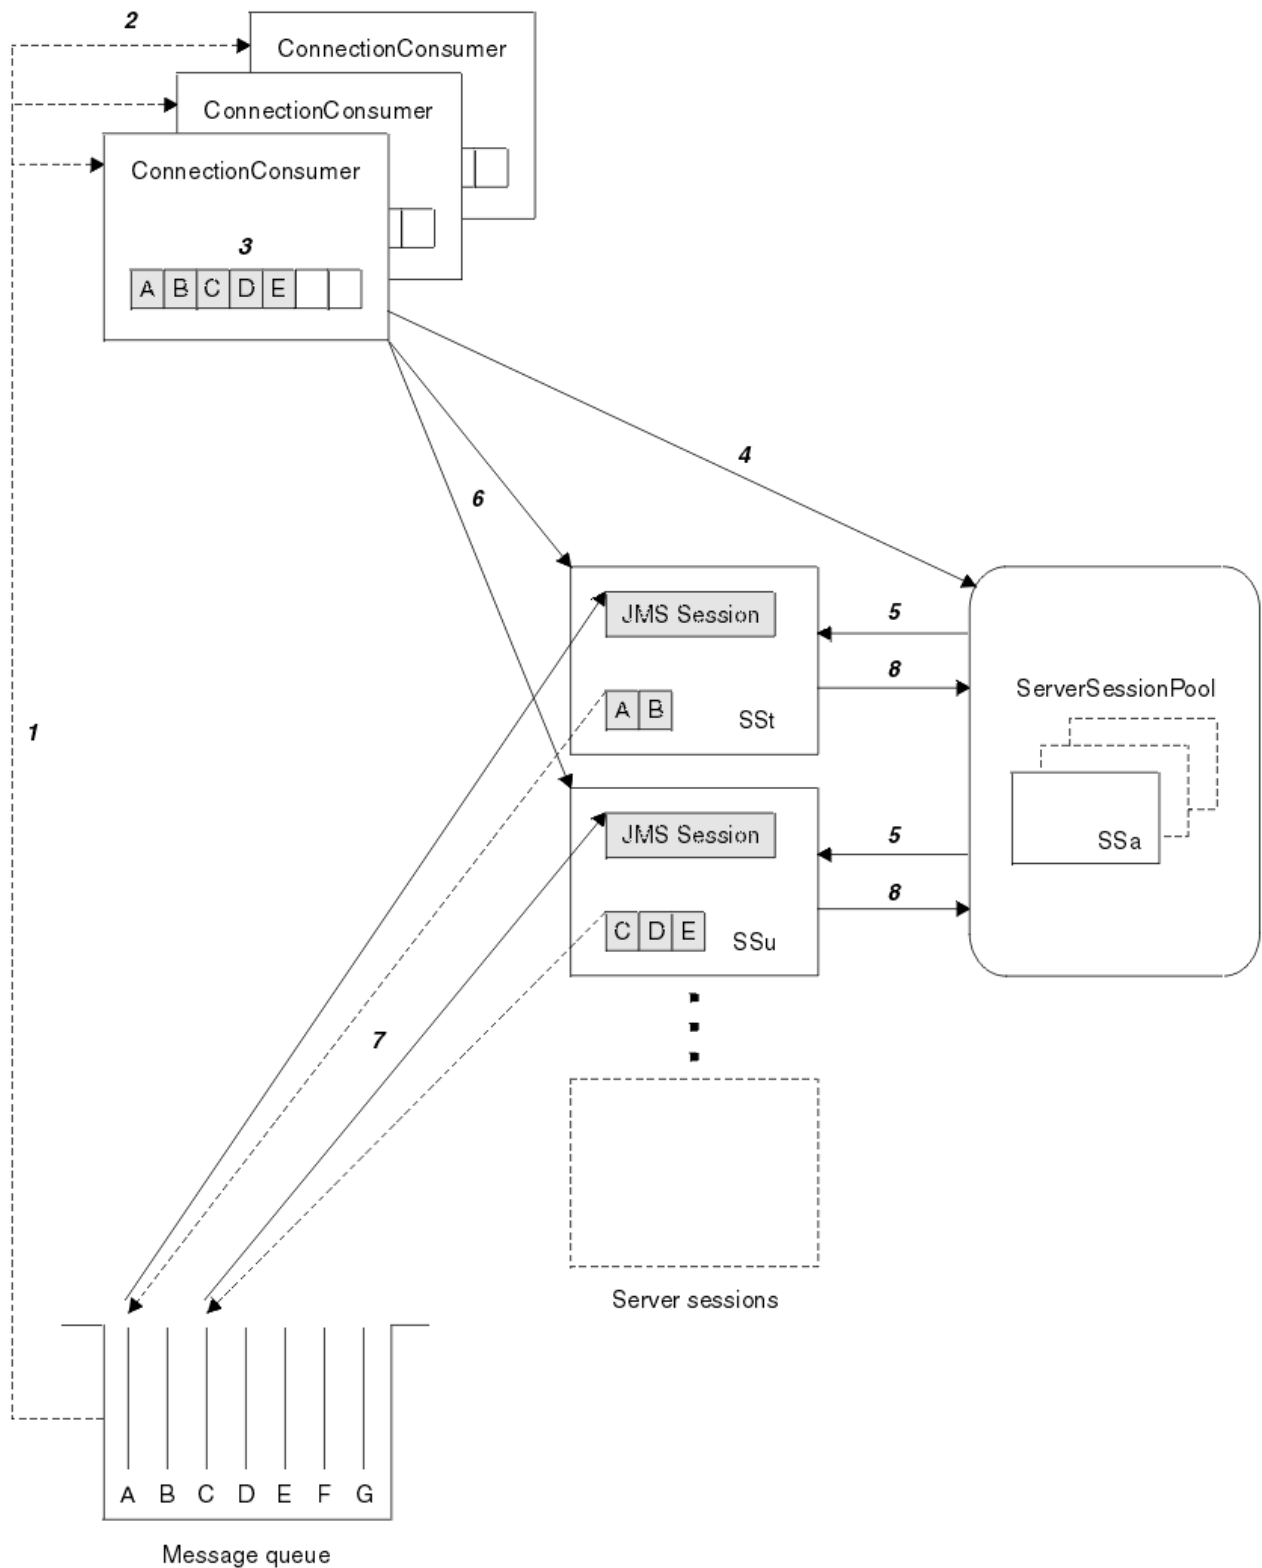


Figure 165. ServerSessionPool and ServerSession functionality

1. The ConnectionConsumers get message references from the queue.
2. Each ConnectionConsumer selects specific message references.
3. The ConnectionConsumer buffer holds the selected message references.
4. The ConnectionConsumer requests one or more ServerSessions from the ServerSessionPool.

5. ServerSessions are allocated from the ServerSessionPool.
6. The ConnectionConsumer assigns message references to the ServerSessions and starts the ServerSession threads running.
7. Each ServerSession retrieves its referenced messages from the queue. It passes them to the onMessage method from the MessageListener that is associated with the JMS Session.
8. After it completes its processing, the ServerSession is returned to the pool.

An application server normally supplies ServerSessionPool and ServerSession functionality.

Using the WebSphere MQ JMS administration tool

Use the administration tool to define the properties of eight types of WebSphere MQ classes for JMS object and to store them within a JNDI namespace. Applications can then use JNDI to retrieve these administered objects from the namespace.

The WebSphere MQ classes JMS objects that you can administer by using the tool are:

- MQConnectionFactory
- MQQueueConnectionFactory
- MQTopicConnectionFactory
- MQQueue
- MQTopic
- MQXAConnectionFactory
- MQXAQueueConnectionFactory
- MQXATopicConnectionFactory

For details of these objects, see [“Administering JMS objects” on page 897](#).

The property types and values you require to use this tool are listed in [Properties of IBM WebSphere MQ classes for JMS objects](#).

The tool also allows administrators to manipulate directory namespace subcontexts within the JNDI. See [“Manipulating subcontexts with the WebSphere MQ JMS administration tool” on page 897](#).

You can also create and configure JMS administered objects with the WebSphere MQ Explorer.

Invoking the IBM WebSphere MQ classes for JMS administration tool

The administration tool has a command line interface. You can use this interactively, or use it to start a batch process.

The interactive mode provides a command prompt where you can enter administration commands. In the batch mode, the command to start the tool includes the name of a file that contains an administration command script.

Interactive mode

To start the tool in interactive mode, enter the command:

```
JMSAdmin [-t] [-v] [-cfg config_filename]
```

where:

-t

Enables trace (default is trace off)

The trace file is generated in "%MQ_JAVA_DATA_PATH%\errors (Windows) or /var/mqm/trace (UNIX). The name of the trace file is of the form:

mqjms_PID.trc

where *PID* is the process ID of the JVM.

-v

Produces verbose output (default is terse output)

-cfg config_filename

Names an alternative configuration file. If this parameter is omitted, the default configuration file, `JMSAdmin.config`, is used. (See [“Configuring the JMS Administration tool”](#) on page 894)

A command prompt is displayed, which indicates that the tool is ready to accept administration commands. This prompt initially appears as:

```
InitCtx>
```

indicating that the current context (that is, the JNDI context to which all naming and directory operations currently refer) is the initial context defined in the `PROVIDER_URL` configuration parameter (see [“Configuring the JMS Administration tool”](#) on page 894).

As you traverse the directory namespace, the prompt changes to reflect this, so that the prompt always displays the current context.

Batch mode

To start the tool in batch mode, enter the command:

```
JMSAdmin <test.scp
```

where *test.scp* is a script file that contains administration commands (see [“Administration commands in the WebSphere MQ JMS administration tool”](#) on page 896). The last command in the file must be the `END` command.

Configuring the JMS Administration tool

The WebSphere MQ JMS Administration tool uses a configuration file to set the values of certain properties. A sample file is supplied, which you can tailor to your system.

The configuration file is a plain-text file that consists of a set of key-value pairs, separated by the equal sign (=). This is shown in the following example:

```
#Set the service provider
INITIAL_CONTEXT_FACTORY=com.sun.jndi.ldap.LdapCtxFactory
#Set the initial context
PROVIDER_URL=ldap://polaris/o=ibm_us,c=us
#Set the authentication type
SECURITY_AUTHENTICATION=none
```

(A # in the first column of the line indicates a comment, or a line that is not used.)

A sample configuration file is supplied with WebSphere MQ. The file is called `JMSAdmin.config`, and is found in the `<MQ_JAVA_INSTALL_PATH>/bin` directory. Edit this file to suit the setup of your system.

Configure the administration tool with values for the following properties:

INITIAL_CONTEXT_FACTORY

The service provider that the tool uses. The supported values for this property are as follows:

- `com.sun.jndi.ldap.LdapCtxFactory` (for LDAP)
- `com.sun.jndi.fscontext.RefFSContextFactory` (for file system context)

You can also use an `InitialContextFactory` that is not in the preceding list. See [“Using an unlisted InitialContextFactory with the WebSphere MQ JMS administration tool”](#) on page 895 for more details.

PROVIDER_URL

The URL of the initial context of the session; the root of all JNDI operations carried out by the tool. Two forms of this property are supported:

- ldap://hostname/contextname
- file:[drive:]/pathname

The format of the LDAP URL can vary, depending on your LDAP provider. See your LDAP documentation for more information.

SECURITY_AUTHENTICATION

Whether JNDI passes security credentials to your service provider. This property is used only when an LDAP service provider is used. This property can take one of three values:

- none (anonymous authentication)
- simple (simple authentication)
- CRAM-MD5 (CRAM-MD5 authentication mechanism)

If a valid value is not supplied, the property defaults to none. See [“Configuring security for the JMS administration tool” on page 895](#) for more details about security with the administration tool.

These properties are set in a configuration file. When you invoke the tool, you can specify this configuration by using the `-cfg` command-line parameter, as described in [“Invoking the IBM WebSphere MQ classes for JMS administration tool” on page 893](#). If you do not specify a configuration file name, the tool attempts to load the default configuration file (`JMSAdmin.config`). It searches for this file first in the current directory, and then in the `<MQ_JAVA_INSTALL_PATH>/bin` directory, where `<MQ_JAVA_INSTALL_PATH>` is the path to your WebSphere MQ classes for JMS installation.

Using an unlisted InitialContextFactory with the WebSphere MQ JMS administration tool

Two `InitialContextFactory` values are supported. You can use other JNDI contexts by setting parameters in the JMS administration configuration file.

You can use the administration tool to connect to JNDI contexts other than those listed in [“Configuring the JMS Administration tool” on page 894](#) by using three parameters defined in the `JMSAdmin` configuration file.

To use a different `InitialContextFactory`:

1. Set the `INITIAL_CONTEXT_FACTORY` property to the required class name.
2. Define the behavior of the `InitialContextFactory` using the `USE_INITIAL_DIR_CONTEXT`, `NAME_PREFIX` and `NAME_READABILITY_MARKER` properties.

The settings for these properties are described in the sample configuration file comments.

You do not need to define the three properties listed here if you use one of the supported `INITIAL_CONTEXT_FACTORY` values. However, you can give them values to override the system defaults. If you omit one or more of the three `InitialContextFactory` properties, the administration tool provides suitable defaults based on the values of the other properties.

Configuring security for the JMS administration tool

Use the `SECURITY_AUTHENTICATION` property to determine whether security credentials are passed to the service provider.

The `SECURITY_AUTHENTICATION` property is described in [“Configuring the JMS Administration tool” on page 894](#). Its effect is as follows:

- If you set this parameter to none, JNDI does not pass any security credentials to the service provider, and *anonymous authentication* is performed.
- If you set the parameter to either `simple` or `CRAM-MD5`, security credentials are passed through JNDI to the underlying service provider. These security credentials are in the form of a user distinguished name (User DN) and password.

If security credentials are required, you are prompted for them when the tool initializes. Avoid this by setting the PROVIDER_USERDN and PROVIDER_PASSWORD properties in the JMSAdmin configuration file.

Note: If you do not use these properties, the text typed, *including the password*, is echoed to the screen. This might have security implications.

The tool does no authentication itself; the task is delegated to the LDAP server. The LDAP server administrator must set up and maintain access privileges to different parts of the directory. See your LDAP documentation for more information. If authentication fails, the tool displays an appropriate error message and terminates.

More detailed information about security and JNDI is in the documentation at Sun's Java website (<https://java.sun.com>).

Administration commands in the WebSphere MQ JMS administration tool

The administration tool accepts commands consisting of an administration verb and its appropriate parameters.

When the command prompt is displayed, the tool is ready to accept commands. Administration commands are generally of the following form:

```
verb [param]*
```

where **verb** is one of the administration verbs listed in [Table 133 on page 896](#). All valid commands contain one verb, which appears at the beginning of the command in either its standard or short form.

The parameters a verb can take depend on the verb. For example, the END verb cannot take any parameters, but the DEFINE verb can take any number of parameters. Details of the verbs that take at least one parameter are discussed in the related topics.

Table 133. Administration verbs		
Verb	Short form	Description
ALTER	ALT	Change at least one of the properties of an administered object
DEFINE	DEF	Create and store an administered object, or create a subcontext
DISPLAY	DIS	Display the properties of one or more stored administered objects, or the contents of the current context
DELETE	DEL	Remove one or more administered objects from the namespace, or remove an empty subcontext
CHANGE	CHG	Alter the current context, allowing the user to traverse the directory namespace anywhere below the initial context (pending security clearance)
COPY	CP	Make a copy of a stored administered object, storing it under an alternative name
MOVE	MV	Alter the name under which an administered object is stored
END		Close the administration tool

Verb names are not case-sensitive.

Typically, to terminate commands, you press the carriage return key. However, you can override this by typing the plus sign (+) directly before the carriage return. This enables you to enter multiline commands, as shown in the following example:

```
DEFINE Q(BookingsInputQueue) +
```



```
QMGR(QM.POLARIS.TEST) +  
QUEUE(BOOKINGS.INPUT.QUEUE) +  
PORT(1415) +  
CCSID(437)
```

Lines beginning with any of the following characters are treated as comments and are ignored: * # / .

Manipulating subcontexts with the WebSphere MQ JMS administration tool

Use the verbs **CHANGE**, **DEFINE**, **DISPLAY** and **DELETE** to manipulate directory namespace subcontexts.

The use of these verbs is described in [Table 134 on page 897](#).

Table 134. Syntax and description of commands used to manipulate subcontexts	
Command syntax	Description
DEFINE CTX(ctxName)	Attempts to create a child subcontext of the current context, having the name ctxName. Fails if there is a security violation, if the subcontext already exists, or if the name supplied is not valid.
DISPLAY CTX	Displays the contents of the current context. Administered objects are annotated with a, subcontexts with [D]. The Java type of each object is also displayed.
DELETE CTX(ctxName)	Attempts to delete the current context's child context having the name ctxName. Fails if the context is not found, is non-empty, or if there is a security violation.
CHANGE CTX(ctxName)	<p>Alters the current context, so that it now refers to the child context having the name ctxName. One of two special values of ctxName can be supplied:</p> <p>=UP moves to the parent of the current context</p> <p>=INIT moves directly to the initial context</p> <p>Fails if the specified context does not exist, or if there is a security violation.</p>

Administering JMS objects

This section describes the eight types of object that the administration tool can handle. It includes details about each of their configurable properties and the verbs that can manipulate them.

You can also create and configure JMS administered objects with the WebSphere MQ Explorer.

JMS object types

The table shows the eight types of administered objects.

The Keyword column shows the strings that you can substitute for *TYPE* in the commands shown in [Table 136 on page 899](#).

Table 135. The JMS object types that are handled by the administration tool		
Object Type	Keyword	Description
MQConnectionFactory	CF	The WebSphere MQ implementation of the JMS ConnectionFactory interface. This represents a factory object for creating connections in both the point-to-point and publish/subscribe domains.

Table 135. The JMS object types that are handled by the administration tool (continued)

Object Type	Keyword	Description
MQQueueConnectionFactory	QCF	The WebSphere MQ implementation of the JMS QueueConnectionFactory interface. This represents a factory object for creating connections in the point-to-point domain.
MQTopicConnectionFactory	TCF	The WebSphere MQ implementation of the JMS TopicConnectionFactory interface. This represents a factory object for creating connections in the publish/subscribe domain.
MQQueue	Q	The WebSphere MQ implementation of the JMS Queue interface. This represents a destination for messages in the point-to-point domain.
MQTopic	T	The WebSphere MQ implementation of the JMS Topic interface. This represents a destination for messages in the publish/subscribe domain.
MQXAConnectionFactory ^{“1” on page 898}	XACF	The WebSphere MQ implementation of the JMS XAConnectionFactory interface. This represents a factory object for creating connections in both the point-to-point and publish/subscribe domains, and where the connections use the XA versions of JMS classes.
MQXAQueueConnectionFactory ^{“1” on page 898}	XAQCF	The WebSphere MQ implementation of the JMS XAQueueConnectionFactory interface. This represents a factory object for creating connections in the point-to-point domain that use the XA versions of JMS classes.
MQXATopicConnectionFactory ^{“1” on page 898}	XATCF	The WebSphere MQ implementation of the JMS XATopicConnectionFactory interface. This represents a factory object for creating connections in the publish/subscribe domain that use the XA versions of JMS classes.
Note: 1. These classes are provided for use by vendors of application servers. They are unlikely to be directly useful to application programmers.		

Verbs used with JMS objects

You can use the verbs ALTER, DEFINE, DISPLAY, DELETE, COPY, and MOVE to manipulate administered objects in the directory namespace.

Table 136 on page 899 summarizes the use of these verbs. Substitute *TYPE* with the keyword that represents the required administered object, as listed in Table 135 on page 897.

Table 136. Syntax and description of commands used to manipulate administered objects

Command syntax	Description
ALTER <i>TYPE</i> (name) [property]*	Attempts to update the properties of the administered object with the ones supplied. Fails if there is a security violation, if the specified object cannot be found, or if the new properties supplied are not valid.
DEFINE <i>TYPE</i> (name) [property]*	Attempts to create an administered object of type <i>TYPE</i> with the supplied properties, and store it under the name name in the current context. Fails if there is a security violation, if the supplied name is not valid or an object of that name exists, or if the properties supplied are not valid.
DISPLAY <i>TYPE</i> (name)	Displays the properties of the administered object of type <i>TYPE</i> , bound under the name name in the current context. Fails if the object does not exist, or if there is a security violation.
DELETE <i>TYPE</i> (name)	Attempts to remove the administered object of type <i>TYPE</i> , having the name name, from the current context. Fails if the object does not exist, or if there is a security violation.
COPY <i>TYPE</i> (nameA) <i>TYPE</i> (nameB)	Makes a copy of the administered object of type <i>TYPE</i> , having the name nameA, naming the copy nameB. This all occurs within the scope of the current context. Fails if the object to be copied does not exist, if an object of name nameB exists, or if there is a security violation.
MOVE <i>TYPE</i> (nameA) <i>TYPE</i> (nameB)	Moves (renames) the administered object of type <i>TYPE</i> , having the name nameA, to nameB. This all occurs within the scope of the current context. Fails if the object to be moved does not exist, if an object of name nameB exists, or if there is a security violation.

Creating objects with the WebSphere MQ JMS administration tool

Create objects and store them in a JNDI namespace using the DEFINE command,

Use the following command syntax:

```
DEFINE TYPE(name) [property]*
```

That is, the DEFINE verb, followed by a *TYPE*(name) administered object reference, followed by zero or more *properties* (see [Properties of IBM WebSphere MQ classes for JMS objects](#)).

LDAP naming considerations for JMS objects

To store your objects in an LDAP environment, you must give them names that comply with certain conventions. The administration tool can help you obey naming conventions by adding a default prefix.

One naming convention is that object and subcontext names must include a prefix, such as cn= (common name), or ou= (organizational unit).

The administration tool simplifies the use of LDAP service providers by allowing you to refer to object and context names without a prefix. If you do not supply a prefix, the tool automatically adds a default prefix to the name you supply. For LDAP, this is cn=.

You can change the default prefix by setting the NAME_PREFIX property in the JMSAdmin configuration file, as described in [“Using an unlisted InitialContextFactory with the WebSphere MQ JMS administration tool”](#) on page 895.

This is shown in the following example.

```

InitCtx> DEFINE Q(testQueue)

InitCtx> DISPLAY CTX

Contents of InitCtx

a cn=testQueue com.ibm.mq.jms.MQQueue

1 Object(s)
0 Context(s)
1 Binding(s), 1 Administered

```

Although the object name supplied (testQueue) does not have a prefix, the tool automatically adds one to ensure compliance with the LDAP naming convention. Likewise, submitting the command DISPLAY Q(testQueue) also causes this prefix to be added.

You might need to configure your LDAP server to store Java objects. For information to assist with this configuration, see the documentation for your LDAP server.

Sample error conditions creating a JMS object

A number of common error conditions can arise when you create an object.

The following are examples of these error conditions:

CipherSpec mapped to CipherSuite

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) SSLCIPHERSUITE(RC4_MD5_US)
WARNING: Converting CipherSpec RC4_MD5_US to
CipherSuite SSL_RSA_WITH_RC4_128_MD5

```

Invalid property for object

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) PRIORITY(4)
Unable to create a valid object, please check the parameters supplied
Invalid property for a QCF: PRI

```

Invalid type for property value

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) CCSID(english)
Unable to create a valid object, please check the parameters supplied
Invalid value for CCS property: English

```

Property clash - client/bindings

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) HOSTNAME(polaris.hursley.ibm.com)
Unable to create a valid object, please check the parameters supplied
Invalid property in this context: Client-bindings attribute clash

```

Property clash - Exit initialization

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) SECEXITINIT(initStr)
Unable to create a valid object, please check the parameters supplied
Invalid property in this context: ExitInit string supplied
without Exit string

```

Property value outside valid range

```

InitCtx/cn=Trash> DEFINE Q(testQ) PRIORITY(12)
Unable to create a valid object, please check the parameters supplied
Invalid value for PRI property: 12

```

Unknown property

```

InitCtx/cn=Trash> DEFINE QCF(testQCF) PIZZA(ham and mushroom)
Unable to create a valid object, please check the parameters supplied
Unknown property: PIZZA

```

The following are examples of error conditions that might arise on Windows when looking up JNDI administered objects from a JMS application.

1. If you are using the WebSphere JNDI provider, `com.ibm.websphere.naming.WsnInitialContextFactory`, you must use a forward slash (/) to access administered objects defined in subcontexts; for example, `jms/MyQueueName`. If you use a backslash (\), an `InvalidNameException` is thrown.
2. If you are using the Sun JNDI provider, `com.sun.jndi.fscontext.RefFSContextFactory`, you must use a backslash (\) to access administered objects defined in subcontexts; for example, `ctx1\\fred`. If you use a forward slash (/), a `NameNotFoundException` is thrown.

Using WebSphere MQ Explorer for JMS configuration

Use the IBM WebSphere MQ Explorer graphical user interface to create JMS objects from WebSphere MQ objects, and WebSphere MQ objects from JMS objects, as well as for administering and monitoring other WebSphere MQ objects.

Before you begin

Before you create and configure JMS administered objects with the WebSphere MQ Explorer, add an initial context to define the root of the JNDI namespace in which the JMS objects are stored in the naming and directory service. For more information, refer to the IBM WebSphere MQ Explorer user assistance for JMS administered objects.

About this task

You can perform the following tasks with the IBM WebSphere MQ Explorer, either contextually from an existing object in the IBM WebSphere MQ Explorer, or from within a create new object wizard. Refer to the WebSphere MQ Explorer help for examples of the WebSphere MQ Explorer user assistance for some typical tasks.

Procedure

- Create a JMS Connection Factory from any of the following WebSphere MQ objects:
 - a) A WebSphere MQ queue manager, whether on your local computer or on a remote system.
 - b) A WebSphere MQ channel
 - c) A WebSphere MQ listener
- Add a WebSphere MQ queue manager to WebSphere MQ Explorer using a JMS Connection Factory
- Create a JMS queue from a WebSphere MQ queue
- Create a WebSphere MQ queue from a JMS queue
- Create a JMS topic from a WebSphere MQ topic, which can be a WebSphere MQ object or a dynamic topic
- Create a WebSphere MQ topic from a JMS topic

Using the WebSphere MQ Headers package

The WebSphere MQ Headers package provides a set of helper interfaces and classes that you can use to manipulate the WebSphere MQ headers of a message. Typically, you use the WebSphere MQ Headers package because you want to perform administrative services by using the command server (by using Programmable Command Format (PCF) messages).

About this task

The WebSphere MQ Headers package is located in the `com.ibm.mq.headers` and `com.ibm.mq.pcf` packages. You can use this facility for both of the two alternative APIs that WebSphere MQ provides for use in Java applications:

- WebSphere MQ classes for Java (also referred to as WebSphere MQ Headers Base Java).
- WebSphere MQ classes for Java Message Service (WebSphere MQ classes for JMS, also referred to as WebSphere MQ JMS).

WebSphere MQ Base Java applications typically manipulate `MQMessage` objects, and the Headers support classes can directly interact with these objects, since they natively understand the WebSphere MQ Base Java interfaces.

In WebSphere MQ JMS, the payload for a message is typically a `String` or a byte array object, which can be manipulated with `DataInput` and `DataOutput` streams. The WebSphere MQ Headers package can be used to interact with these data streams and is suitable for manipulating any MQ messages that are sent and received by WebSphere MQ JMS applications.

Therefore, although the WebSphere MQ Headers package contains references to the WebSphere MQ Base Java package, it is also intended for use within WebSphere MQ JMS applications and is suitable for use within Java Platform, Enterprise Edition (Java EE) environments.

A typical way in which you might use the WebSphere MQ Headers package is to manipulate administration messages in Programmable Command Format (PCF), for example for any of the following reasons:

- To access details about a WebSphere MQ resource.
- To monitor the depth of a queue.
- To inhibit access to a queue.

By using PCF messages with the WebSphere MQ JMS API, this kind of administration of application-centric resources can be performed from within Java EE applications without having to resort to using the WebSphere MQ Base Java API.

Procedure

- To use the WebSphere MQ Headers package to manipulate message headers for WebSphere MQ classes for Java, see [“Using with WebSphere MQ classes for Java” on page 902](#).
- To use the WebSphere MQ Headers package to manipulate message headers for JMS, see [“Using with WebSphere MQ classes for JMS” on page 903](#).

Using with WebSphere MQ classes for Java

WebSphere MQ classes for Java applications typically manipulate `MQMessage` objects, and the Headers support classes can directly interact with these objects, since they natively understand the WebSphere MQ classes for Java interfaces.

About this task

WebSphere MQ provides some sample applications that demonstrate how to use the WebSphere MQ Headers package with the WebSphere MQ Base Java API (WebSphere MQ classes for Java).

The samples show two things:

- How to create a PCF message to perform an administrative action and parse the response message.
- How to send this PCF message using the WebSphere MQ classes for Java.

Depending on which platform you are using, these samples are installed under the `pcf` directory in the `samples` or `tools` directory of your WebSphere MQ installation (see [“Installation directories for WebSphere MQ classes for Java” on page 628](#)).

Procedure

1. Create a PCF message to perform an administrative action and parse the response message.
2. Send this PCF message using the WebSphere MQ classes for Java.

Related concepts

[“Handling WebSphere MQ message headers with WebSphere MQ classes for Java” on page 648](#)

Java classes are provided representing different types of message header. Two helper classes are also provided.

[“Handling PCF messages with WebSphere MQ classes for Java” on page 653](#)

Java classes are provided to create and parse PCF-structured messages, and to facilitate sending PCF requests and collecting PCF responses.

Using with WebSphere MQ classes for JMS

To use the WebSphere MQ Headers with the WebSphere MQ classes for JMS, you carry out the same essential steps as for WebSphere MQ classes for Java . The PCF message can be created and the response parsed in exactly the same way by using the WebSphere MQ Headers package and the same sample code as for WebSphere MQ classes for Java .

About this task

To send a PCF message using the WebSphere MQ API, the message payload must be written into a JMS Bytes Message, and sent using the standard JMS APIs. The only consideration is that the message must not contain a JMS RFH2 or any other headers with specific values in the MQMD.

To send a PCF message, complete the following steps. The way in which the PCF message is created, and information is extracted from the response message is the same as for WebSphere MQ classes for Java (see [“Using with WebSphere MQ classes for Java” on page 902](#)).

Procedure

1. Create a JMS Queue Destination that represents the SYSTEM.ADMIN.COMMAND.QUEUE.

WebSphere MQ JMS applications send the PCF messages to the SYSTEM.ADMIN.COMMAND.QUEUE, and need access to a JMS Destination object that represents this queue. The Destination must have the following properties set:

```
WMQ_MQMD_WRITE_ENABLED = YES
WMQ_MESSAGE_BODY = MQ
```

If you are using WebSphere Application Server, you must define these properties as custom properties on the Destination.

To create the destination programmatically from within an application, use the following code:

```
Queue q1 = session.createQueue("SYSTEM.ADMIN.COMMAND.QUEUE");
((MQQueue) q1).setIntProperty(WMQConstants.WMQ_MESSAGE_BODY,
    WMQConstants.WMQ_MESSAGE_BODY_MQ);
((MQQueue) q1).setMQMDWriteEnabled(true);
```

2. Convert a PCF message into a JMS Bytes message containing the correct MQMD values.

A JMS Bytes message needs to be created, and the PCF Message written to it. A response queue needs to be created, but this needs to have no specific settings.

The following sample code snippet shows how to create a JMS Bytes Message and write a `com.ibm.mq.headers.pcf.PCFMessage` object into it. The `PCFMessage` object (`pcfCmd`) has previously been built using the WebSphere MQ Headers package. (Note the package to load the `PCFMessage` is `com.ibm.mq.headers.pcf.PCFMessage`).

```
// create the JMS Bytes Message
final BytesMessage msg = session.createBytesMessage();

// Create the wrapping streams to put the bytes into the message payload
ByteArrayOutputStream baos = new ByteArrayOutputStream();
DataOutput dataOutput = new DataOutputStream(baos);

// Set the JMSReplyTo so the answer comes back
msg.setJMSReplyTo(new MQQueue("adminResp"));

// write the pcf into the stream
pcfCmd.write(dataOutput);
```

```

baos.flush();
msg.writeBytes(baos.toByteArray());

// we have taken control of the MD, so need to set all
// flags in the MD that we require - main one is the format
msg.setJMSPriority(4);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_PERSISTENCE,
    CMQC.MQPER_NOT_PERSISTENT);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_EXPIRY, 300);
msg.setIntProperty(WMQConstants.JMS_IBM_MQMD_REPORT,
    CMQC.MQRO_PASS_CORREL_ID);
msg.setStringProperty(WMQConstants.JMS_IBM_MQMD_FORMAT, "MQADMIN");

// and send the message
sender.send(msg);

```

3. Send the message, and receive the response using the standard JMS APIs.
4. Convert the response message into a PCF message for processing.

To retrieve the response message and process it as a PCF message, use the following code:

```

// Get the message back
BytesMessage msg = (BytesMessage) consumer.receive();

// get the size of the bytes message & read into an array
int bodySize = (int) msg.getBodyLength();
byte[] data = new byte[bodySize];
msg.readBytes(data);

// Read into Stream and DataInput Stream
ByteArrayInputStream bais = new ByteArrayInputStream(data);
DataInput dataInput = new DataInputStream(bais);

// Pass to PCF Message to process
PCFMessage response = new PCFMessage(dataInput);

```

Related concepts

[“JMS messages” on page 772](#)

JMS messages are composed of a header, properties, and a body. JMS defines five types of message body.

Using Web services in WebSphere MQ

You can develop IBM WebSphere MQ applications for Web services using the IBM WebSphere MQ transport for SOAP or the IBM WebSphere MQ bridge for HTTP.

The IBM WebSphere MQ transport for SOAP provides a JMS transport for SOAP. The IBM WebSphere MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

For more information about the IBM WebSphere MQ transport for SOAP, see [“WebSphere MQ transport for SOAP” on page 905](#).

With IBM WebSphere MQ bridge for HTTP, client applications can exchange messages with IBM WebSphere MQ without the need to install a WebSphere MQ MQI client. You can call WebSphere MQ from any platform or language with HTTP capabilities.

For more information about the IBM WebSphere MQ bridge for HTTP, see [“WebSphere MQ bridge for HTTP” on page 977](#).

Related concepts

[“Application development concepts” on page 8](#)

You can use a choice of procedural or object-oriented languages to write IBM WebSphere MQ applications. Use the links in this topic for information about IBM WebSphere MQ concepts that are useful for application developers.

[“Deciding which programming language to use” on page 75](#)

Use this information to find out about programming languages and frameworks supported by IBM WebSphere MQ, and some considerations for using them.

[“Designing IBM WebSphere MQ applications” on page 85](#)

When you have decided how your applications can take advantage of the platforms and environments available to you, you need to decide how to use the features offered by WebSphere MQ.

[“Sample WebSphere MQ programs” on page 91](#)

Use this collection of topics to learn about sample WebSphere MQ programs on different platforms.

[“Writing a queuing application” on page 186](#)

Use this information to learn about writing queuing applications, connecting to and disconnecting from a queue manager, publish/subscribe, and opening and closing objects.

[“Writing client applications” on page 335](#)

What you need to know to write client applications on WebSphere MQ.

[“Writing publish/subscribe applications” on page 266](#)

Start writing publish/subscribe WebSphere MQ applications.

[“Building an IBM WebSphere MQ application” on page 409](#)

Use this information to learn about building an IBM WebSphere MQ application on different platforms.

[“Handling program errors” on page 523](#)

This information explains errors associated with your applications MQI calls either when it makes a call, or when its message is delivered to its final destination.

WebSphere MQ transport for SOAP

The WebSphere MQ transport for SOAP provides a JMS transport for SOAP. The WebSphere MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

Introduction to IBM WebSphere MQ transport for SOAP

The IBM WebSphere MQ transport for SOAP provides a JMS transport for SOAP. The WebSphere MQ SOAP sender and listener provide a means to call Web services.

The WebSphere MQ SOAP listener supports services hosted by .NET Framework 1, .NET Framework 2, and Axis 1.4. The WebSphere MQ SOAP sender supports web services clients running on .NET Framework 1, .NET Framework 2, Axis 1.4 and Axis2. Clients can be either a WebSphere MQ server or client application. The IBM WebSphere MQ transport for SOAP is also integrated into other environments such as Microsoft Windows Communication Foundation, WebSphere Application Server, and CICS Transaction Server.

The integration into Microsoft Windows Communication Foundation is part of the IBM WebSphere MQ support for .NET Framework 3.

The IBM WebSphere MQ transport for SOAP is a set of protocols and tools to transport SOAP messages by using JMS over IBM WebSphere MQ. It is packaged in different ways for different application environments as shown in [Table 137 on page 905](#).

<i>Table 137. IBM WebSphere MQ transport for SOAP application environments</i>		
	Integrated with additional WebSphere MQ components	Integrated into a framework
Provided as part of WebSphere MQ installation	.NET Framework 1 .NET Framework 2 Axis 1.4	Windows Communication Foundation (.NET Framework 3) Axis2 (Client only)

Table 137. IBM WebSphere MQ transport for SOAP application environments (continued)

	Integrated with additional WebSphere MQ components	Integrated into a framework
Provided in another software package		WebSphere Application Server CICS Transaction Server 4.1 WebSphere ESB WebSphere Process Server for Multiplatforms

Integration of the IBM WebSphere MQ transport for SOAP into an application framework simplifies the development and deployment of web services to IBM WebSphere MQ.

With additional IBM WebSphere MQ SOAP components, you can interact directly with the WebSphere MQ SOAP components to develop and deploy services. Use the IBM WebSphere MQ SOAP tools to configure and deploy the web services and the web service clients to IBM WebSphere MQ.

In the integrated environments, development and deployment is simpler. You use the same tools for development and deployment as you would to develop and deploy a SOAP HTTP web service. You must still configure the IBM WebSphere MQ queues, channels, and queue managers that you require using WebSphere MQ tools.

You can mix and match IBM WebSphere MQ SOAP clients and servers from any of these environments.

Benefits

The WebSphere MQ transport for SOAP offers existing IBM WebSphere MQ users the following principal benefits:

Using your IBM WebSphere MQ network to connect existing web services.

The services might be ones you have written, or services that are provided as interfaces to other packaged software applications you have deployed.

The benefit comes from using your existing WebSphere MQ network to connect web services. The IBM WebSphere MQ transport has the advantage of being a managed and reliable queued messaging service.

Writing new applications, or converting existing applications, to use SOAP rather than IBM WebSphere MQ interfaces.

Typically, applications require a specific WebSphere MQ adapter to be developed to integrate with another application. Adapters have two parts: the connector piece, that puts and gets messages to and from the transport, and the adapter piece that converts data to and from application-specific formats. Integrating each pair of applications is a new challenge.

The benefit of SOAP comes from standardizing on SOAP for defining application interfaces, and then having a choice of transports. You do not need to write application-specific adapters, and you can choose whether to use IBM WebSphere MQ or HTTP as the connector. Which transport you choose depends on what qualities of service and connectivity you require.

For existing SOAP over HTTP users, the benefit of WebSphere MQ transport for SOAP comes from using a managed and reliable asynchronous transport. The benefits are twofold:

A truly asynchronous programming model for availability and performance.

By using an asynchronous client interface, the client and service applications do not have to be available at the same time. Requests sent by the client will be stored until the service is available to process them.

A ready built managed network that is designed to be reliable and available.

By choosing IBM WebSphere MQ as a transport, you are getting the advantage of using a managed network that provides reliable messaging.

In contrast, transports such as HTTP and FTP over TCP/IP are unmanaged. An unmanaged network is ideal for unpredictable connections: there are fewer management tasks.

Summary

IBM WebSphere MQ transport for SOAP provides the following components:

- The SOAP/JMS transport binding is used in WSDL documents to bind a SOAP service to a JMS transport. The WebSphere MQ implementation of the SOAP/JMS binding uses a URI that takes either of two forms:

WebSphere MQ transport for SOAP

```
jms:/queue?&Name=Value&Name=Value...
```

WebSphere MQ wire format for W3C candidate recommendation

```
jms:queue:qName?connectionFactory=connectQueueManager(qMgrName)&Name=Value&Name=Value...
```

- The mapping of a SOAP message onto a WebSphere MQ message.
- Two IBM WebSphere MQ SOAP listeners to receive SOAP requests, one for Java and one for .NET Framework 1 or .NET Framework 2. The listeners use .NET or Axis 1.4 to process the SOAP request.
- Two IBM WebSphere MQ SOAP senders to create IBM WebSphere MQ SOAP requests. web services clients register with a sender to process `jms: SOAP` requests.
- Integration with Windows Communication Foundation (WCF), sometimes known as .NET 3, to send and receive WebSphere MQ Transport for SOAP messages.
- Integration of the client with Axis2, sometimes known as JAX-WS, to send WebSphere MQ Transport for SOAP or W3C SOAP JMS messages.
- The command **amqwdployMQService**, which creates development and run time components and scripts to deploy a web service by using the IBM WebSphere MQ transport for SOAP.
- Sample Java and .NET client and service code.
- A script to set the class path, and other utility scripts.

In the integrated environments the sender and listener are integrated into each environment, as are extensions to the development and deployment tools.

Integration of SOAP and WebSphere MQ

The WebSphere MQ transport for SOAP extends SOAP, and Web services tools and run time, with WebSphere MQ as an alternative transport to HTTP for SOAP. You do not need to modify existing Web services to use WebSphere MQ transport for SOAP as a transport. The transport uses a custom URI format for SOAP/JMS. The W3C URI format for SOAP/JMS is supported in a limited way by Axis2 clients.

An additional line of code has to be added to clients in the .NET Framework 1, .NET Framework 2, and Axis 1.4 environments. No additional code is required in Axis 2 and Windows Communication Foundation (WCF) clients. The WebSphere MQ SOAP listener runs services in the .NET Framework 1, .NET Framework 2, and Axis 1.4 environments. The WebSphere MQ transport for SOAP is integrated into some other application server environments, including WCF, CICS, and WebSphere Application Server.

What is SOAP?

SOAP⁹ describes the standardized format of the messages and interaction protocols that applications use to exchange requests, replies, and datagrams. SOAP is independent of the transport used to transfer the messages, and of the application environment that sends and receives the messages. The W3C defines SOAP Version 1.2 succinctly:

*SOAP Version 1.2 provides the definition of the XML-based information which can be used for exchanging structured and typed information between peers in a decentralized, distributed environment.*¹⁰

⁹ Historically, the acronym stood for Simple Object Access Protocol.

¹⁰ [W3C: SOAP Version 1.2 Part 0](#)

To use SOAP it must be bound to a transport, such as HTTP, e-mail, or WebSphere MQ.

A SOAP protocol binding framework is the set of rules for carrying a SOAP message on top of another protocol, such as HTTP. [SOAP Version 1.2 Part 2: Adjuncts \(Second Edition\)](#) describes the SOAP HTTP binding.

The W3C candidate recommendation, 4 June 2009, [SOAP over Java Message Service 1.0](#), describes the recommendation for the SOAP JMS binding. As JMS is an API specification, and not a transport protocol, the JMS SOAP recommendation does not describe the wire-format of SOAP JMS messages. It describes the SOAP interaction protocols and the JMS API binding. Consequently, when using the JMS SOAP recommendation you must still use the same JMS implementation for the SOAP client and the SOAP server. It does enable a SOAP JMS application to be run on any implementation of JMS. A JMS implementation can be plugged into a J2EE application server, if both the server and the JMS implementation comply with the JCA specification. WebSphere MQ JMS complies with the JCA specification and can be plugged into a compliant application server.

WebSphere MQ transport for SOAP binding is like the proposed W3C standard, but it is not the same. Its usage is described in the topic [MQRFH2 SOAP settings](#). Unlike the W3C candidate recommendation, the SOAP binding is not formally specified. Effectively, it is the HTTP binding, and the service address takes the form, `jms:/queue?name=value&name=value...`, rather than `http://authority/path?query#fragment`. `jms:` is not an officially registered IANA URI scheme.

What is a Web service?

SOAP enables programs written in different languages, running on different platforms, to communicate using various transport protocols. SOAP is the protocol specification. A Web service is an application that provides a service through a SOAP interface that can be accessed using internet protocols.

An important goal of SOAP is to provide services that clients can use easily. Once you have designed a client to use a service, you can program the call to invoke the service without reference to external documentation. Service interfaces are described in XML, in a WSDL document. The query, `http://authority/path?wsdl`, returns the WSDL description of a SOAP service.

Tip: When you deploy a Web service to use WebSphere MQ, also deploy the service to HTTP so that the standard WSDL query works.

Developing Web services

Web services have a client and a service part. The service is written first, either starting from the interface description in WSDL, or by following the rules for writing the service class. Web service toolkits have utilities to generate WSDL from the interface definition of a class; for example **java2wsdl** or **disco**. They also have tools to generate or class skeletons from WSDL interface descriptions; for example **wsdl2java**, **wsimport**, or **wsdl**. The former is known as bottom up development, and the latter top down.

The **amqwdployWMQService** command in WebSphere MQ transport for SOAP uses these tools to generate WSDL, client stubs, and client proxies.

Web services are typically written using an integrated development environment targeted at a particular application server environment:

Eclipse IDE for Java EE Developers

Creates Web services for Axis 2. Supports JAX-RPC and JAX-WS

Rational® Application Developer V7.5

Creates Web services for WebSphere Application Server V7 and previous versions, and also for Axis. Supports JAX-RPC and JAX-WS.

WebSphere Integration Developer V6.2

Creates Web services for WebSphere Process Server and WebSphere ESB. Supports JAX-RPC and JAX-WS.

Visual Studio 2008 (Version 9)

Creates Web services for .NET Framework 3.5 and earlier (Windows Communication Foundation)

Visual Studio 2005 (Version 8)

Creates Web services for .NET Framework 2 and earlier

You can use any of these tools in combination with WebSphere MQ transport for SOAP. Once you have developed a service to use with HTTP, use the **amqwdployMQService** tool to deploy the services to use WebSphere MQ as a transport. You can write a new client using the output from the tool, or modify your existing clients to use the WebSphere MQ transport for SOAP.

If WebSphere MQ transport for SOAP is integrated into the application environment, then you do not need use the **amqwdployMQService** tool or modify the client code. The client SOAP layer directs client requests that have a URI with the prefix `jms:` to the WebSphere MQ transport for SOAP. The server SOAP layer calls WebSphere MQ transport for SOAP to wait for `jms:` SOAP requests, and returns responses to WebSphere MQ transport for SOAP.

Typically, .NET services have been developed bottom up using Web service annotations in code, and Java services top down, using WSDL interface definitions. The difference in approaches is narrowing, as Java Standard Edition Version 6 supports JAX-WS 2.0, and uses annotations to qualify the definition of service interfaces. It is now as easy to develop Java services bottom up as top down. Which approach you choose is a matter of development method.

The Web services client is written after the service, using the WSDL service definition and generated client stubs and proxies. In some applications, the service definition is not known when the client is written. The client retrieves the service WSDL and creates service requests dynamically. More commonly the service definition is known, but the address to which the service is deployed, is not. The Web service toolkit generates interfaces for the client to use to make service requests. The client provides the service address when it is required. In the third case, the WSDL contains all the information a client needs. The WSDL contains both the interface and address of the service. The code generated by the Web service toolkit has all the information needed by the client to make requests of a service.

You can use any of these three styles with WebSphere MQ transport for SOAP.

Web service application environments

Web service toolkits require a mapping from the WSDL definition of a service to the streams of bytes that are transferred in SOAP requests and responses. The byte stream is defined by the SOAP specification, and is contained in the SOAP envelope. The SOAP envelope is shown in [Figure 166 on page 909](#).

```
<?xml version='1.0'?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
<soap:Header> <!-- optional -->
<!-- headers... -->
</soap:Header>
<soap:Body>
<!-- payload or fault message -->
</soap:Body>
</soap:Envelope>
```

Figure 166. SOAP envelope

Mapping from the SOAP envelope to the language binding and back is part standardized, and part proprietary. The mapping is fundamental to the .NET architecture, and is provided as part of the Common Language Runtime (CLR). The mapping is standardized in Java by JAX specifications. Because the Java mappings are standardized, Java Web service clients and services are portable between different Java-based application environments. JAX-RPC (sometimes called JAX-WS 1.0) is the mapping most in use today. It is supported by Axis 1.4. JAX-WS (sometimes called JAX-WS 2.0) is a greatly improved standard and is likely to replace JAX-RPC rapidly. JAX-WS is supported by Axis 2.0. WebSphere MQ 7.0.1 does not support JAX-WS and Axis 2.

WebSphere MQ transport for SOAP does not alter the contents of the SOAP envelope, and the contents do not affect the transport. The language bindings do affect the WebSphere MQ transport for SOAP. WebSphere MQ 7.0.1 supports .NET Framework 1, .NET Framework 2, and Axis 1.4 using the code and utilities shipped with WebSphere MQ transport for SOAP. Support for the WebSphere transport for SOAP

in .NET Framework 3 and 3.5 is implemented using the WebSphere MQ custom channel for Windows Communication Foundation.

Other SOAP development and run time environments might ship support for WebSphere MQ transport for SOAP, and support different languages. For example, Web services running on CICS supports languages such as COBOL and PL/1.

Note: The mapping used makes no difference to the interoperability of Web services. You can mix and match clients and services written using .NET, JAX-RPC, and JAX-WS mappings.

What is WebSphere MQ transport for SOAP?

WebSphere MQ transport for SOAP is a SOAP binding and a Web services toolkit. Together, they enable applications to exchange SOAP messages using WebSphere MQ rather than HTTP. [Figure 167 on page 910](#) shows WebSphere MQ as an alternative to HTTP as a SOAP transport.

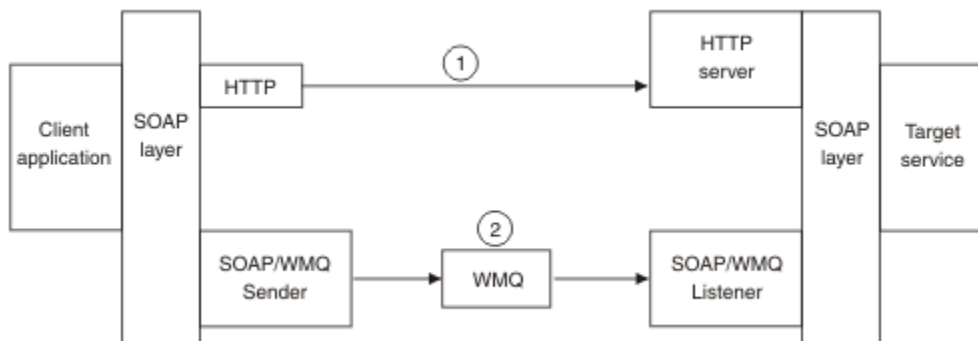


Figure 167. Overview of WebSphere MQ transport for SOAP

SOAP over HTTP is shown as (1) in the diagram. The client SOAP layer converts a request into a SOAP message and the HTTP component sends over TCP/IP. The HTTP server component listens for HTTP requests, typically on the TCP/IP port 80. If the request is for a SOAP service, the HTTP server component calls the SOAP layer to convert the SOAP request into method call. It then returns the response.

SOAP over WebSphere MQ is shown as (2). The client application registers the WebSphere MQ SOAP sender component as a handler for the `jms:` protocol with the SOAP layer. The SOAP layer passes SOAP messages addressed to `jms:` to the WebSphere MQ SOAP sender. The sender uses the URI in the message to place the message on the request queue with the required qualities of service. The corresponding WebSphere MQ SOAP listener waits for messages on its request queue and calls the SOAP layer to process requests and return responses.

The SOAP sender and listener are normal WebSphere MQ programs. They can be connected to the same queue manager, as in [Figure 168 on page 911](#), or connected to different queue managers; see [Figure 169 on page 912](#). The client can be connected by a client connection.




Figure 168. Queues used by SOAP/WebSphere MQ (single queue manager)




Figure 169. Queues used by SOAP/WebSphere MQ (separate queue managers)

W3C candidate recommendation for binding SOAP to JMS.

The W3C candidate recommendation defines the SOAP over JMS binding; SOAP over Java Message Service 1.0. Also useful for its examples is [URI Scheme for Java\(tm\) Message Service 1.0](#)¹¹.

Some application frameworks, such as WebSphere Application Server v7, have support for the W3C candidate recommendation. Send SOAP requests formatted with a URI compatible with the W3C candidate recommendation using the Axis2 client; see [W3C SOAP over JMS URI for the WebSphere MQ Axis 2 client](#). The Axis2 client sends a SOAP request formatted with either a W3C or a WebSphere MQ transport for SOAP based on URI in the SOAP request.

Axis2 client support for the W3C recommendation is introduced in the 7.0.1.3 fixpack. Support for other clients, and for the SOAP listeners provided by WebSphere MQ is not provided.

Related concepts

[Implementation of WebSphere transport for SOAP on .NET Framework 1, .NET 2 and Axis 1.4](#)

You might want to write your own WebSphere MQ SOAP sender and listener. Use the implementation of WebSphere MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4 as a guide.

[WebSphere MQ transport for SOAP and Web services reliable messaging](#)

¹¹ Look for *URI Scheme for JMS*, in the W3C specification references, for the latest draft.

Web services reliable messaging is a protocol for reliably exchanging Web service requests and responses over an unreliable connection. It is best-suited to solve problems of short-lived connection disruption.

Implementation of WebSphere transport for SOAP on .NET Framework 1, .NET 2 and Axis 1.4

You might want to write your own WebSphere MQ SOAP sender and listener. Use the implementation of WebSphere MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4 as a guide.

1. A client program uses the appropriate Web services framework in the same way as it would for the HTTP transport. It must also register the `jms:` prefix. The prefix is registered using either the `com.ibm.mq.soap.Register.extension()` Java method or the `IBM.WMQSOAP.Register.Extension()` CLR method.
2. The Axis 1.4 or .NET Framework 1 or 2 framework marshals the call into a SOAP request message exactly as for SOAP/HTTP.
3. A WebSphere MQ service is identified by a URI prefixed with `jms:`. When the framework identifies the `jms:` URI, it calls the WebSphere MQ transport sender code; `com.ibm.mq.soap.transport.jms.WMQSender` (for Axis 1.4) or `IBM.WMQSOAP.MQWebRequest` (for .NET1 and 2). If the framework encounters a URI with an `http:` prefix, it calls the standard SOAP over HTTP sender.
4. The SOAP message is transported by the WebSphere MQ SOAP sender using the request queue. The **SimpleJavaListener** (for Java) or **amqwSOAPNETListener** (for .NET) receives the request message.

The WebSphere MQ SOAP listeners are stand-alone processes and are multithreaded with a customizable number of threads.

5. The WebSphere MQ SOAP listener reads the incoming SOAP request, and passes it to the appropriate Web service infrastructure.
6. The Web service infrastructure parses the SOAP request message and invokes the service. The procedure is the same as for a message that arrived on an HTTP transport.
7. The infrastructure formats the response into a SOAP response message and returns it to the WebSphere MQ SOAP listener.
8. The listener places the message on the response queue, and the message is transferred to the WebSphere MQ SOAP sender. The sender passes it to the client Web service infrastructure.
9. The client infrastructure parses the response SOAP message and hands the result back to the client application.

Each application context is served by a separate WebSphere MQ request queue.

The application context is controlled in Axis 1.4 by ensuring that the WebSphere MQ SOAP listener and service execute in the appropriate directory. Axis 1.4 sets the correct `CLASSPATH` for the directory.

Application context is controlled in .NET by the WebSphere MQ SOAP listener executing the service in a context created by a call to `ApplicationHost.CreateApplicationHost`. The call specifies the target execution directory. Each service then operates in the directory in which it was deployed.

amqwdeployWMQService generates the request and response queues. It also generates the infrastructure necessary for handling the queues and deploying services to Axis 1.4.

Related concepts

[Integration of SOAP and WebSphere MQ](#)

[WebSphere MQ transport for SOAP and Web services reliable messaging](#)

Web services reliable messaging is a protocol for reliably exchanging Web service requests and responses over an unreliable connection. It is best-suited to solve problems of short-lived connection disruption.

WebSphere MQ transport for SOAP and Web services reliable messaging

Web services reliable messaging is a protocol for reliably exchanging Web service requests and responses over an unreliable connection. It is best-suited to solve problems of short-lived connection disruption.

WebSphere MQ for SOAP takes advantage of using a WebSphere MQ managed and reliable network for passing SOAP messages. Transports such as HTTP and FTP are unmanaged. Unmanaged networks are ideal for unpredictable connections, where the difficulties and costs of managing connections outweigh the benefits of not losing requests and responses.

To overcome the problem of losing files when connections break in unmanaged networks, services like managed FTP build a management layer on top of FTP. The management layer takes over the burden of checking that files have transferred successfully from users, and retransmits missing files if necessary. To use managed FTP, you must have the management software installed at both ends of the connection.

Web services reliable messaging (WSRM) takes a different approach to solving the problem of unreliable connections. Its goal is to transfer Web service requests and responses reliably, without both ends of the connection having to use the same software. Any software, by implementing the Web services reliable messaging protocol, can exchange messages reliably with other software.

When a connection fails, a sender and receiver must preserve the context of the WSRM message transfer, using a generated URI as a key. The sender and receiver keep attempting to establish a new connection. If a new connection is successfully established, the transfer completes. The WSRM specification does not specify how context is preserved, or when a new connection is attempted.

You might decide that only short-lived outages are of interest. For longer outages, you might be prepared to discard transfers that could not be restarted after a time. Similarly, you might be prepared to discard transfers if either the client or service fails. Leaving the user responsible for assuring transfers, places fewer demands on managing the coordination of client and service.

If network outages are long-lived, over 30 minutes or so, or if the client or server fails, there is an increased likelihood that some connections are never re-established. You can no longer rely on WSRM restoring message transfer automatically in an unmanaged way. You have to consider managing the failed WSRM connections, which implies developing software to manage the network of clients and services.

Using WSRM to overcome short outages might significantly reduce dealing with lost messages on a mobile network. If you do not have to assure message delivery, the benefits of reducing message loss might justify the additional cost of developing a WSRM implementation.

SOAP over JMS provides assured message delivery and deals with longer duration outages of the client, the server, and the network. If you are seeking a more reliable quality of service for SOAP than HTTP, which solution do you choose: WebSphere MQ transport for SOAP or WSRM? The answer depends on many factors. Some of the factors to consider are listed:

1. Whether the unreliability is due to connection failure.
2. How long-lived connection failures are.
3. If you can manage both the client and server side of the connection.
4. If the user or an administrator is ultimately responsible for the delivery of messages.

Related concepts

[Integration of SOAP and WebSphere MQ](#)

[Implementation of WebSphere transport for SOAP on .NET Framework 1, .NET 2 and Axis 1.4](#)

You might want to write your own WebSphere MQ SOAP sender and listener. Use the implementation of WebSphere MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4 as a guide.

Installing and verifying WebSphere MQ Web services

Use the instructions in these topics to install and verify the WebSphere MQ transport for SOAP.

Installing WebSphere MQ Web transport for SOAP

Use these instructions to install the WebSphere MQ Web transport for SOAP. The installation creates tools to run Web service clients or services using WebSphere MQ as the SOAP transport. The tools are used in the .NET Framework 1, .NET 2, Axis 1.4, or Axis2 SOAP environments.

Before you begin

Check the prerequisite products at [System Requirements for IBM WebSphere MQ](#). The installation process does not check for the presence and availability of the prerequisite software. You must verify that the prerequisites are installed.

WebSphere MQ provides a copy of the Axis 1.4 run time. Use this version with WebSphere MQ rather than any other you might have installed. IBM does not provide technical support for Apache Axis. Contact the Apache Software Foundation if you have technical problems with it.

To run Web services in the .NET Framework 3 SOAP environment, WebSphere MQ uses Windows Communication Foundation. The WebSphere MQ custom channel for Windows Communication Foundation runs Web service clients and services using WebSphere MQ as a transport for SOAP messages.

About this task

You can install the WebSphere MQ Web transport for SOAP as either a WebSphere MQ MQI client or Server application. If you have already installed WebSphere MQ as a client or a server on your computer, check that you have installed the components listed.

`MQ_INSTALLATION_PATH` represents the high-level directory in which WebSphere MQ is installed.

Carry out the following installation steps.

Procedure

1. Select the "Java and .Net Messaging and Web services" component for installation.
2. On Solaris and HP-UX, select the "Java Runtime environment" component for installation.
3. Select the Development toolkit for installation.
4. Install and verify WebSphere MQ as described in the Quick Beginning for your platform.
5. Copy the Apache Axis 1.4 run time, `axis.jar` from the `prereqs/axis` directory on the WebSphere MQ installation media. Copy it to the installation directory described in [Table 138 on page 916](#), [Table 139 on page 916](#), or [Table 140 on page 916](#).

Windows

```
Copy D:\PreReqs\axis\axis.jar MQ_INSTALLATION_PATH\java\lib\soap
```

AIX

```
cp -f PreReqs/axis/axis.jar MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chown mqm:mqm MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chmod 444 MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
```

HP-UX, Solaris, and Linux (all platforms) installation directories

```
cp -f PreReqs/axis/axis.jar MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chown mqm:mqm MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
chmod 444 MQ_INSTALLATION_PATH/java/lib/soap/axis.jar
```

- On Windows 2003, run **Aspnet_regiis.exe** to update the script maps to point to the version of the Common Language Run time you are using.
Look for the **Aspnet_regiis.exe** utility in %SystemRoot%\Microsoft.NET\Framework*version-number* .
- Set the environment variable, WMQSOAP_HOME , to point to the WebSphere MQ installation directory.

Results

Table 138. Windows installation directories	
Location	Contents
MQ_INSTALLATION_PATH\programs\bin	Binary, command, DLL, and executable files
MQ_INSTALLATION_PATH\programs\java\lib	.jar files
MQ_INSTALLATION_PATH\programs\java\lib\soap	SOAP .jar files
MQ_INSTALLATION_PATH\programs\soap\samples	Samples and IVT

Table 139. AIX installation directories	
Location	Contents
MQ_INSTALLATION_PATH/bin	Shell scripts
MQ_INSTALLATION_PATH/java/lib	.jar files
MQ_INSTALLATION_PATH/java/lib/soap	axis.jar and other JAX-RPC .jar files
MQ_INSTALLATION_PATH/samp/soap	Samples and IVT

Table 140. HP-UX, Solaris, and Linux (all platforms) installation directories	
Location	Contents
MQ_INSTALLATION_PATH/bin	Shell scripts
MQ_INSTALLATION_PATH/java/lib	.jar files
MQ_INSTALLATION_PATH/java/lib/soap	axis.jar and other JAX-RPC .jar files
MQ_INSTALLATION_PATH/samp/soap	Samples and IVT

What to do next

- For .NET only, you must register the WebSphere MQ transport for SOAP files with the Global Assembly Cache. If .NET is already installed when you install WebSphere MQ, registration is performed automatically at installation. If you install .NET after WebSphere MQ, the registration is performed automatically when the IVT is first run.

You can run **amqiregisterdotnet.cmd** to perform registration of the .NET assemblies. You can also run **amqiregisterdotnet.cmd** to force reregistration at any stage. Once made, this registration survives system restarts and subsequent reregistration is not normally necessary.
- Run the Installation Verification Test, as described in [“Verifying the IBM WebSphere MQ transport for SOAP”](#) on page 917.
- If you intend to develop Axis2 client, you must download Axis2 1.4.1 from Apache; see [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse”](#) on page 938.

Verifying the IBM WebSphere MQ transport for SOAP

Verify the IBM WebSphere MQ transport for SOAP using the **runivt** command. The command runs a number of demonstration applications and ensures that the environment is correctly set up after installation.

Before you begin

Before you run the **runivt** command, ensure that you have the following runtime environments:

- To run on Axis only: you must have a Java SDK (within SOE) available on your system. You must also include the location of the `java.exe` and `javac.exe` commands in the systems **PATH** environment variable.
- To run a test on .NET only (supported only on Windows): you must have both a Java SDK and the .NET compilers and tools on your system. To do so, access either a Visual Studio command prompt or the Microsoft Windows SDK command prompt, then add the location of the `java.exe` and `javac.exe` files to the **PATH** environment variable.
- To run all available the tests: for Windows platforms, the environment must be configured as described in the .NET test run. On UNIX and Linux platforms, the environment must be configured as described in the Axis only test run.

About this task

Instead of running the verification test on both .NET and Axis, you might want to run the test only on Axis, or only on .NET.

If you experience problems with the tests and want to start again:

1. Stop the queue manager `WMQSOAP.DEMO.QM` using the `immediate` option.
2. Stop the listener that has been started in a different window.
3. Delete the queue manager.
4. Delete the temporary samples directory you created and start again.

On UNIX and Linux platforms, you must run the command using an X Windows system session.

The **runivt** command changes the contents of the `soap/samples` directory. To keep the installation image unchanged, copy the samples directory to a temporary location, and run the verification test from the temporary location.

You can run the installation verification as many times as you like.

Carry out the following steps to verify the installation of IBM WebSphere MQ transport for SOAP on .NET Framework 1, .NET Framework 2, and Axis 1.4:

Procedure

1. Copy the `./tools/soap/samples` directory tree to a temporary location.
2. Start a command window with the temporary directory as the current directory.
3. Use the **runivt** command to start the installation test. The `runivt` script compiles a test class, the sample client, and services before deploying and running them. For the test class, the sample client, and services to run, complete the installation steps outlined in [Installing WebSphere\(r\) MQ Web transport for SOAP](#) and ensure that the command prompt used to run the `runivt` command has the required runtime environment set. Use any of the following methods to run the **runivt** command:
 - Run a test on Axis only: `runivt Axis`.
 - Run a test on .NET only (supported only on Windows): `runivt DotNet`.
 - Run all available the tests: `runivt`.

For more information about the `runivt` command syntax and parameters, see [runivt: IBM WebSphere MQ transport for SOAP installation verification test](#). The tests that you can run are listed in the file `ivttests.txt` on Windows and `ivttests_unix.txt` on UNIX and Linux platforms.

Related reference

[runivt: WebSphere MQ transport for SOAP installation verification test](#)

Developing Web services for WebSphere MQ transport for SOAP

Use your normal Web service development environment to develop services for use with the WebSphere MQ transport for SOAP.

Before you begin

1. If you are planning to use the command-line tools supplied with WebSphere MQ transport for SOAP:
 - a. Create a deployment directory for the service.
 - b. Start a command window in the directory.
 - c. For .NET, csc.exe and wsdl.exe must be in the path, and be from the same version of .NET Framework.
 - d. For Java,
 - i) Run the **amqwsctcp** command to set up the classpath.
 - ii) An IBM JRE and a JDK at the same version level must be in the current path. The version level must be at least 5.0.
 - iii) Customize the classpath to include the locations of any additional .jar libraries, and directories containing .java packages, including for the service you are developing. Put the current directory "." in the classpath.
 - iv) Create a directory, relative to the current directory of the command window, corresponding to the package name of the service you are developing.
2. Alternatively, use workbench tools that support Web services development. The example development tasks use Microsoft Visual Studio 2008, Eclipse IDE for Java EE Developers and WebSphere Application Server Community Edition.

About this task

Existing Web services need no modification to work with WebSphere transport for SOAP. The tools supplied with WebSphere MQ transport for SOAP deploy a Web service, and run it using a WebSphere MQ SOAP listener. The tools also generate WSDL, .NET client stubs, and .java proxy classes to develop WebSphere MQ transport for SOAP clients.

Follow these steps to create a service, and prepare it for deployment and the generation of clients. Follow the steps in the related tasks to create a service using Eclipse or Microsoft Visual Studio 2008.

Procedure

1. Develop the service using your normal development environment.
2. Test the service using an HTTP Web services client
3. Follow these steps to prepare the deployment directory:
 - For Java
 - a. Copy the .java file defining the service interface into the deployment directory.
 - b. Copy any .class files for the service into the directory corresponding to the package name.
 - c. Check that the classpath can locate all the classes that are required: compile the service .java file using **javac**.
 - For .NET
 - a. Copy the .asmx file defining the service into the deployment directory, and
 - b. If you are using the code-behind model, copy any .dll files into a *deployment directory\bin* directory.

Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse

Develop an Axis 1.4 Web service to run using WebSphere MQ as your service provider. Use your normal Web service development environment to create a service for deployment to Axis 1.4.

Before you begin

Take into account the requirements for deploying a Web server to the WebSphere MQ SOAP listener for Axis 1.4.

- The WebSphere MQ SOAP listener for Axis 1.4 requires an IBM JRE at version 5.0 or greater. The JRE, and JDK used for development, must be at the same version level.
- The WebSphere MQ SOAP listener for Axis 1.4 requires the `axis.jar` installed with WebSphere MQ. Change the build path in your development environment to reference the `axis.jar` file installed with WebSphere MQ, rather than the `axis.jar` files installed with the development environment.
- Up to, and including, WebSphere MQ V7.0.1, the WSDL generated for the deployed service is RPC/encoded. From V7.1, you can also request RPC/literal style WSDL. The generated WSDL is used only for deployment. You can define the service using WS-I compliant WSDL.

About this task

Create the service using your normal Web services development environment.

In this task, we use the freely available open source Eclipse Java EE IDE for Web Developers, known as Galileo. For the application server, we use WebSphere Application Server Community Edition v2.1 (Community Edition), based on Geronimo. See the related tasks for information about how to obtain, install, and configure the IDE and server.

Test the service using HTTP as a transport using the Web Services Explorer provided in the IDE before. Alternatively, generate an HTTP client proxy and test the service using your own client code.

You might follow these steps to develop a bottom-up Web service. Use the sample program `StockQuoteAxis.java` as an example.

Procedure

1. Start Eclipse IDE for Java EE Developers with a new workspace.
2. Configure the workspace to use Java50,
WebSphere Application Server Community Edition 2.1.4 does not work with Java60.
 - a) **Window > Preferences > Java > Installed JREs > Add... > Standard VM > Next > Directory ...**
 - b) Browse to the install directory of **Java50 > OK > Finish**
 - c) Check the **Java50** JRE > OK
3. Add the Community Edition runtime environment and start Community Edition.
 - a) **Window > Preferences > Server > Runtime Environments > Add...**
 - b) Select **IBM WASCE v2.1** from the **New Server Runtime Environment** list > Check **Create a new local server > Next**
If **IBM WASCE 2.1** is not in the list you need to complete two other tasks:
 - i) Install WebSphere Application Server Community Edition.
 - ii) Install the Eclipse update for Community Edition.Find the details at [WebSphere Application Server Community Edition](#)
 - c) Browse to the Application Server Installation Directory > **OK > Finish > OK**.
 - d) Right-click **IBM WASCE v2.1 server** in the servers view > **Start**

Tip: You can manage WASCE in Eclipse: Right-click **IBM WASCE v2.1 server > Launch WASCE Console** . the default **Username** and **Password** are system and manager .

4. Set up the server and runtime for Web services.
 - a) **Window > Preferences > Web services > Server and Runtime**
 - b) Select **IBM WASCE v2.1 Server** as the server.
 - c) Leave **Apache Axis** as the Web service runtime.
5. Create a Dynamic Web Project.
 - a) **File > New > Dynamic Web Project.**
Name the project StockQuoteAxis.
 - b) Check **Add project to an EAR > New...**
 - c) In the **EAR Application Project** page type the **Project name** StockQuoteAxisEAR > **Finish**
Reply **OK** in response to the dialog box suggesting you switch to the Java EE perspective, or stay Java perspective to follow these instructions exactly.
 - d) **IBM WASCE 2.1 server** is selected as the target runtime. Accept it, and the other defaults > **Finish.**
Reply **OK** in response to the dialog box suggesting you switch to the Java EE perspective, or stay Java perspective to follow these instructions exactly.
6. Import the StockQuoteAxis.java sample program
 - a) Open the **StockQuoteAxis** Web project > Right click the **src** folder > **Import...**
 - b) Select **General > File System > Next**
 - c) Browse to `MQ_INSTALLATION_PATH\tools\soap\samples\java\server` > check **StockQuoteAxis.java** > **Finish**
MQ_INSTALLATION_PATH represents the high level directory where WebSphere MQ is installed. You must highlight the server directory to see the files it contains.
7. Correct the compile error by moving StockQuoteAxis.java into its correct package.
 - a) Open StockQuoteAxis.java and right click the problem > **Quick Fix**
 - b) Double click **Move 'StockQuoteAxis.java' to package 'soap.server'** > **Save.**
8. Create a Web service from StockQuoteAxis.java.
 - a) Right-click StockQuoteAxis.java > **Web Services > Create Web service > Next.**
Accept the default configuration for the service:

Web service type
Bottom up Java beanWeb Service

Service implementation
soap.server.StockQuoteAxis

Server
IBM WASCE v2.1 server

Web service runtime
Apache Axis

Service project
StockQuoteAxis

Service EAR project
StockQuoteAxisEAR

Configuration
No client generation
9. Select the methods to access and the style of Web service > **Next.**
Start the server if prompted.
 - a) Leave all the methods selected.
 - b) Select document/literal (wrapped) style.
10. **Finish**

When the service has deployed, look in the WebContent\wsdl folder in the StockQuoteAxis Web project and find the generated StockQuoteAxis.wsdl file.

11. Test the service using HTTP with the Web Services Explorer.

- a) Right-click StockQuoteAxis.wsdl > **Test with Web Services Explorer**.
- b) Click the operation **getQuote** in the **StockQuoteAxisSoapBinding** actions in the **Web Services Explorer** window.
- c) Type **ibm** in the **symbol** input field > **Go**

12. Test the service using WebSphere MQ transport for SOAP.

To deploy the service, use **SimpleJavaListener**, which is in com.ibm.mq.soap.jar. You have to add the WebSphere MQ Java and SOAP libraries to the build path.

- a) Right click the **StockQuoteAxis** Web project > **Build Path** > **Configure Build Path...**
- b) Click the **Libraries** tab > **Add External Jars....** Browse to *MQ_INSTALLATION_PATH*\java\lib. and select all the .jar files > **Open** > **Add External Jars....** Browse to *WMQ Install directory*\java\lib\soap and select all the .jar files > **Open** > **OK**.
MQ_INSTALLATION_PATH represents the high level directory where WebSphere MQ is installed.
- c) In the Project Explorer, Right click StockQuoteAxis\Java Resources\Libraries\com.ibm.mq.soap.jar\com.ibm.mq.soap.transport.jms\SimpleJavaListener.class\ SimpleJavaListener > **Run as...> Run Configurations...**

Tip:

If there is no configuration for SimpleJavaListener, click the **New configuration** icon on the **Create, manage, and run configurations** page of the **Run configurations** wizard, The **SimpleJavaListener** has no command to stop it. To monitor or stop **SimpleJavaListener**, open the **Debug perspective** in Eclipse.

- d) Open the **(x)= Arguments** tab. In the **Program arguments** input area type the parameters to **SimpleJavaListener**.

For this example, type

```
-u "jms:/queue?destination=REQUESTAXIS@QM1&connectionFactory=()
&initialContextFactory=com.ibm.mq.jms.NoJndi
&targetService=StockQuoteAxis
&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE" -n 10
```

Note: The target service is StockQuoteAxis to match the target service name created in the service deployment descriptor, StockQuoteAxis\WebContent\WEB-INF\server-config.wsdd. amqwdployWMQService creates a target service called soap.server.StockQuoteAxis. In this example, you are using the same StockQuoteAxis.class and service-config.wsdd as the HTTP server.

- e) On the same tab, configure **Working directory** to reference the server-config.wsdd file: `${workspace_loc:StockQuoteAxis/WebContent/WEB-INF}`

a) **Run**

Errors are written to the console. If the console remains blank, **SimpleJavaListener** has started ok.

- b) To test the deployment, run a StockQuoteAxis client, developed in the task, [“Developing a JAX-RPC client for WebSphere transport for SOAP using Eclipse” on page 931](#).

Example: StockQuoteAxis sample program

The sample Java Web service, StockQuoteAxis.java, is installed in *WMQ install directory*\tools\soap\samples\java\server.StockQuoteAxis.java, [Figure 170 on page 922](#), has four methods:

1. float getQuote(String symbol)

2. void getQuoteOneWay(String symbol).
 3. int asyncQuote(int delay)
 4. float getQuoteTran(String symbol)
-

```
package soap.server;
import java.lang.Thread;
import java.io.FileWriter;
public class StockQuoteAxis {
    public float getQuote(String symbol) throws Exception {
        return ((float) 55.25);
    }
    public void getQuoteOneWay(String symbol) throws Exception {
        try {
            // Write the results for this service to a file
            FileWriter f = new FileWriter("getQuoteOneWay.txt", true);
            f.write("One way service result via proxy is: 44.44\n");
            f.close();
        } catch (Exception ee) {
            System.out.println("Error writing result file in getQuoteOneWay");
            ee.printStackTrace();
        }
    }
    public int asyncQuote(int delay) {
        try {
            Thread.sleep(delay);
        } catch (Exception e) {
            System.out.println("Exception in asyncQuote during sleep");
        }
        return delay;
    }
    public float getQuoteTran(String symbol) throws Exception {
        if (symbol.equalsIgnoreCase("ROLLBACK")) {
            System.out.println("Rollback was requested,
                exiting from service by calling System.exit().");
            System.exit(0);
        }
        return ((float) 55.25);
    }
}
```

Figure 170. StockQuoteAxis

What to do next

Deploy the service using WebSphere MQ Transport for SOAP, instead of HTTP using the command **amqwdployWMQService**.

The command has an option, `axisDeploy`, which deploys the service by creating an Apache Axis 1.4 deployment descriptor. The WebSphere MQ SOAP listener runs the service. The SOAP listener is called `SimpleJavaListener` and is supplied with the WebSphere MQ transport for SOAP.

Related tasks

[Developing a .NET 1 or 2 service for WebSphere MQ transport for SOAP using Microsoft Visual Studio 2008](#)

[Develop the SampleStockQuote Web service for .NET 1 or .NET 2 using Microsoft Visual Studio 2008](#)

[Developing a JAX-WS EJB Web service for W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a JEE application server. This task is step 2 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol.

Developing a .NET 1 or 2 service for WebSphere MQ transport for SOAP using Microsoft Visual Studio 2008

Develop the SampleStockQuote Web service for .NET 1 or .NET 2 using Microsoft Visual Studio 2008

About this task

Create the StockQuote service with a code-behind implementation using Visual Studio 2008.

Procedure

1. Create a template for the service, and check that it runs on HTTP.
 - a) Start Visual Studio 2008 > **File** > **New** > **Project....** Select **C#** Project Type, **.NET Framework 2**, and **ASP.NET Web Service Application**. Type the **Name:** and **Solution Name:** StockQuoteDotNet > **OK**
 - b) Right click **Service1.asmx** in the **Solution Explorer** > **Rename** > StockQuote.asmx.
 - c) Change the code fragment `public class Service1` to `public class StockQuote`.
 - d) Right click **StockQuote.asmx** in the **Solution Explorer** > **Open with...** > **XML Editor**. Change `Class="StockQuoteDotNet.Service1"` to `Class="StockQuoteDotNet.StockQuote"`
 - e) Change the code fragment `[WebService(Namespace = "http://tempuri.org/")]` to `[WebService(Namespace = "http://stock.samples/")]`.
 - f) Remove the line of code `[ToolboxItem(false)]`.
 - g) Check everything is correct so far: **Debug** > **Start Debugging (F5)**. Verify the output in Explorer.
2. Add the methods from the sample `SQDNNonInline.asmx.cs`, and test the service on HTTP.
 - a) Open `MQ_INSTALLATION_PATH\tools\soap\samples\dotnet\SQDNNonInline.asmx.cs` and replace the `HelloWorld` method with the four `Quote` methods; see Figure 171 on page 924. `MQ_INSTALLATION_PATH` represents the directory where WebSphere MQ is installed.
 - b) **Build** > **Rebuild** the solution > Right click one of the **Thread** in error > **Resolve** > Using **System.Threading**.
 - c) Press F5 to start debugging.

The service is not conformant to WS-I Basic Profile v1.1. You have the choice of either changing the `WebMethod` annotation from `[SoapRpcMethod]` to `[SoapDocumentMethod]`, or removing the annotation `[WebServiceBinding(ConformsTo = WsiProfiles.BasicProfile1_1)]`.
 - d) Press F5 to verify the implementation using HTTP.
3. Generate WSDL, clients, and run the service using WebSphere MQ transport for SOAP.
 - a) Open a command window in the project directory tree, where the `StockQuote.asmx` is stored.
 - b) (Optional) Use `amqswdeployWMQService` to generate artifacts. The queue manager must be started:

```
amqswdeployWMQService -f StockQuote.asmx
-u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
&targetService=StockQuote.asmx"
StockQuote.asmx StockQuote.wsd1
```

All artifacts are created in the `./generated` directory tree.

- c) (Optional) Generate just the WSDL for calling the service using WebSphere MQ transport for SOAP.

```
amqswsdl -u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
```

```
&targetService=StockQuote.asmx"
StockQuote.asmx StockQuote.wsdl
```

- a) Run the .NET listener. Either use `.\generated\server\startWMQNListener.cmd` or type the command:

```
amqSOAPNETListener -u "jms:/queue?initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=()
&destination=REQUESTDOTNET@QM1
&targetService=StockQuote.asmx"
```

4. Test the service using a client generated from the WSDL, or using the clients generated by **amqwdployWMQService**.

Sample code

The sample .NET Web service, `StockQuoteDotNet`, is installed in `MQ_INSTALLATION_PATH\tools\soap\samples\dotnet`. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed. The Web service binding of the published samples differ slightly from the binding used in the task. The task uses the defaults used in Visual Studio 2008.

There are two examples of .NET Framework 1 and .NET Framework 2 Web services. `StockQuoteDotNet.asmx`, is an inline service. `SQDNNNoninline.asmx`, is a code-behind Web service implemented by `SQDNNNoninline.asmx.cs`.

`StockQuoteDotNet` has four methods:

1. `float getQuote(String symbol)`
2. `void getQuoteOneWay(String symbol).`
3. `int asyncQuote(int delay)`
4. `float getQuoteDOC(String symbol)`

```
<%@ WebService Language="C#" Class="StockQuoteDotNet" %>
using System;
using System.Web.Services;
using System.Web.Services.Protocols;
using System.Web.Services.Description;
using System.Threading;
[WebService (Namespace="http://stock.samples")]
public class StockQuoteDotNet {
    [WebMethod] [SoapRpcMethod(OneWay=true)]
    public void getQuoteOneWay(String symbol) {
        if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(5000);
        System.Console.WriteLine("getQuoteOneWay was invoked.");
    }
    [WebMethod] [SoapRpcMethod]
    public float getQuote(String symbol) {
        if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(10000);
        return 88.88F;
    }
    [WebMethod] [SoapRpcMethod]
    public int asyncQuote(int delay) {
        Thread.Sleep(delay);
        return delay;
    }
    [WebMethod]
    public float getQuoteDOC(String symbol) {
        return 77.77F;
    }
}
```

Figure 171. Inline service: `StockQuoteDotNet.asmx`

```
<%@ WebService Language="C#" Codebehind="SQDNNonInline.asmx.cs" Class="SQDNNonInline" %>
```

Figure 172. Code-behind: Design *SQDNNonInline.asmx*

```
using System;
using System.Web.Services;
using System.Web.Services.Protocols;
using System.Web.Services.Description;
using System.Threading;

[WebService(Namespace = "http://stock.samples")]
public class SQDNNonInline : System.Web.Services.Protocols.SoapHttpClientProtocol
{
    [WebMethod]
    [SoapRpcMethod(OneWay = true)]
    public void getNonInlineQuoteOneWay(String symbol)
    {
        if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(5000);
        System.Console.WriteLine("getNonInlineQuoteOneWay was invoked.");
    }

    [WebMethod]
    [SoapRpcMethod]
    public float getNonInlineQuote(String symbol)
    {
        if (symbol.ToUpper().Equals("DELAY")) Thread.Sleep(10000);
        return 88.88F;
    }

    [WebMethod]
    [SoapRpcMethod]
    public int asyncNonInlineQuote(int delay)
    {
        Thread.Sleep(delay);
        return delay;
    }

    [WebMethod]
    public float getNonInlineQuoteDOC(String symbol)
    {
        return 77.77F;
    }
}
```

Figure 173. Code-behind: Implementation: *SQDNNonInline.asmx.cs*

Related tasks

[Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse](#)

Develop an Axis 1.4 Web service to run using WebSphere MQ as your service provider. Use your normal Web service development environment to create a service for deployment to Axis 1.4.

[Developing a JAX-WS EJB Web service for W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a JEE application server. This task is step 2 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol.

Developing a JAX-WS EJB Web service for W3C SOAP over JMS

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a JEE application server. This task is step 2 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol.

Before you begin

Use Rational Application Developer to create the EJB Web service. The Web service wizard in Rational Application Developer has an option to create a Web service using the W3C candidate recommendation

for the SOAP over JMS binding. Rational Application Developer 7.54 is required. The exercise used Rational Application Developer included in Rational Software Architect for WebSphere Software v7.5.5.1. The EJB is deployed to WebSphere Application Server from Rational Application Developer as part of this task. You must complete [“Configuring WebSphere Application Server to use W3C SOAP over JMS” on page 961](#)

To create the WSDL actually used in the task, you must first complete the task, [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#). Then you can either import the WSDL from the Dynamic Web project in the Eclipse Galileo workspace, or from the running HTTP Web service deployed to WASCE.

WebSphere Application Server might still be running as a result of doing [“Configure WebSphere Application Server resources” on page 963](#). If it is not, you can start it from the Servers view in RAD.

About this task

In this task, you redeploy the StockQuoteAxis service from running as a JAX-RPC Axis service run by the **SimpleJavaListener** using WebSphere MQ transport for SOAP, to being a JAX-WS service running in WebSphere Application server using the W3C SOAP over JMS protocol.

There are two parts to migrating the service from the **SimpleJavaListener** to WebSphere Application Server:

1. Generate the Web service interface from the WSDL for the service using the Top-down EJB Web service wizard in Rational Application Developer.
2. Implementing the service by importing the WebSphere MQ SOAP sample StockQuoteAxis.java.

An alternative approach would have been to generate the service bottom up, from the StockQuoteAxis.java. However, to be sure that the interface to the migrated service is identical, the top-down approach is better, as it uses the same WSDL.

The Web service is developed for the EJB container and not the Web container because the JMS support is part of the EJB container.

Procedure

1. Start Rational Application Developer, and verify WebSphere Application Server is running.
 - a) Start Rational Application Developer in a new workspace.
 - b) Open the Java EE perspective.
 - c) Open the **Servers** tab, and check WebSphere Application Server is running.
 - If there is no WebSphere Application Server v7.0 in the view, right-click in the view > **New** > **Server**. Follow the choices in the wizard to create a WebSphere Application Server v7.0 instance.
 - If the server is present, but not started, click the arrow head to start it.
 - To verify the properties and get quick access to the server logs, right-click **WebSphere Application Server v7.0 at localhost** > **Properties** > **WebSphere Application Server**.
 - To administer the server, either use an external browser and open the URL, <http://localhost:9061/ibm/console/unsecureLogon.jsp>, or right-click **WebSphere Application Server v7.0 at localhost** > **Run administrative console**.
 - The default setting is to publish automatically. Many people prefer deploy updates to the server manually. Double-click **WebSphere Application Server v7.0 at localhost**, and expand the **Publishing** twisty in the **Overview** window. Click **Never publish automatically**.
 - Another default that you might want to alter is to clear the **Terminate server on workbench shutdown** checkbox in the **Overview** window.
2. Create the JEE projects

You must create an Enterprise Application Project (EAR) and an Enterprise Java Bean (EJB) Project.

 - a) **File** > **New** > **Enterprise Application Project**. Name the project W3CJMSEAR > **Finish**.

The defaults must identify WebSphere Application Server v7.0 as the target runtime, and EAR version 5.0. The default configuration must be selected.

- b) **File > New > EJB Project**. Name the project W3CJMSEJB. Select W3CEARJMS as the **EAR Project Name > Next**.

The default EJB module version is 3.0 and the default configuration is used again.

- c) Clear the **Create an EJB Client JAR module** checkbox > **Finish**.

3. Generate and deploy the EJB Web service from the StockQuoteAxis WSDL.

- a) **Run > Launch the Web Services Explorer**.

- b) Select the WSDL page using the icons in the **Web Services Explorer** window > click **WSDL main** in the Navigator.

- c) In the **Actions** window, type in or browse to the WSDL URL for StockQuoteAxis.wsd1.

If you have WASCE running with StockQuoteAxis deployed as an HTTP service, the URL is:

```
http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl
```

If you have the WSDL in the file system, the URL might be:

```
File:\Dirpath\StockQuoteAxis\WebContent\wsdl\StockQuoteAxis.wsd1
```

- d) Click the line containing the imported URL in the Navigator tree.

It is the line immediately following **WSDL Main**, if it is the first WSDL you have imported into Web Services Explorer.

- e) In the **Actions** window, click **Launch Web Service Wizard > Web Service Skeleton > Go**

- f) In the Web Service wizard, select **Top down EJB Web Service**

Select or verify the Configuration using information from [Table 141 on page 927](#) Check **Overwrite files without warning > Next**.

Table 141. Top down EJB Web service configuration

Field	Value
Server	WebSphere Application Server v7.0
Web service run time	IBM WebSphere JAX-WS
Service project	W3CJMSEJB
Service EAR project	W3CJMSEAR
Configuration:	No client generation

- g) On the page subtitled, **Specify options for creating a WebSphere JAX-WS EJB Top Down Web Service**, check the **Switch to JMS binding** box. Also check **Enable Wrapper Style**, **Copy WSDL to project**, and **Generate Web Service Deployment Descriptor > Next**.

- h) On the page titled, **WebSphere JAX-WS JMS Binding Configuration**, check **Use SOAP/JMS interoperability protocol** and provide values from [Table 142 on page 927](#), leaving other fields blank > **Next**.

Table 142. WebSphere JAX-WS JMS Binding Configuration

Field	Value
JMS destination	queue
Destination JNDI name:	requestaxis
JMS connection factory	qm1
Reply to Name	W3CJMSEAR

Table 142. WebSphere JAX-WS JMS Binding Configuration (continued)	
Field	Value
Configuration:	replyaxis

- a) On the page titled, **WebSphere JAX-WS Router Project Configuration**, type `qm1as` in the **ActivationSpec JNDI name** field > **Next**.
RAD takes about 30 seconds to a minute to generate and deploy the project.
 - b) Ignore the options in the **Web Service Publication** page > **Finish**.
4. Check the generated WSDL.

You asked for the service-specific WSDL to be generated and saved in the project.

- a) In the Enterprise Explorer navigator, open the folder **W3CJMSEJB** > **ejbmodule** > **META-INF** > **wsdl**. Double click `StockQuoteAxis.wsdl` to open it in the WSDL editor.

Inspect the binding; you see the JMS url:

```
jms:jndi:requestaxis?jndiConnectionFactoryName=qm1&targetService=StockQuoteAxis
```

5. Optional step: Bind the EJB to SOAP over HTTP using JAX-WS.

Providing two bindings to the EJB gives clients the choice of SOAP bindings to call the Web service. It also provides clients with the means to query the Web server to obtain its WSDL, using HTTP.

The steps to bind an EJB to SOAP over HTTP are not included as part of the task.

6. Implement and redeploy `StockQuoteAxis` using the sample `StockQuoteAxis.java`
 - a) In the Enterprise Explorer navigator, open the folder **W3CJMSEJB** > **Services**. Double click `StockQuoteAxisService` to open the implementation class in a Java editor.
 - b) Open the `StockQuoteAxis.java` sample program in the *WebSphere MQ Installation directory*\tools\soap\samples\java\server folder > Select all the methods, but not the class name > **Copy**.
 - c) In `StockQuoteAxisSoapBindingImpl.java` select all the methods, but not the class name, and paste in the methods from `StockQuoteAxis.java`.
 - d) Add a print statement to output to the WebSphere Application Server console when the service is called.
Change the `getQuote(String symbol)` method:

```
public float getQuote(String symbol) {
    System.out.println("StockQuoteAxisSoapBindingImpl called with symbol: "
        + symbol);
    return ((float) 55.25);
}
```

- e) Correct the imports: **Source** > **Organize imports** > **Save**.
- f) Correct the three errors due to the implementation not matching the interface.

The errors are due to three of the methods in `StockQuoteAxis.java` throwing exceptions, and the WSDL for the service not containing any fault messages. The problem is diagnosed as being a mismatch between the method signatures and the method Web service annotations.

Either annotate the methods with `@WebFault` and regenerate the WSDL, or keep the interface unchanged, and remove the exceptions.

To keep the interface the same, remove the three `throws` exception from the method signatures > **Save**.

What to do next

[“Deploying to an Axis2 client using W3C SOAP over JMS” on page 972](#)

Related tasks

[Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse](#)

Develop an Axis 1.4 Web service to run using WebSphere MQ as your service provider. Use your normal Web service development environment to create a service for deployment to Axis 1.4.

Developing a .NET 1 or 2 service for WebSphere MQ transport for SOAP using Microsoft Visual Studio 2008

Develop the SampleStockQuote Web service for .NET 1 or .NET 2 using Microsoft Visual Studio 2008

Developing WebSphere MQ Web service clients for WebSphere MQ transport for SOAP

Use your normal development environment to develop Web service clients for use with the WebSphere MQ transport for SOAP.

Before you begin

Create the service. You can use one of the examples in [“Developing Web services for WebSphere MQ transport for SOAP”](#) on page 918.

Make choices about how you are going to develop, deploy and use the clients, and where to go to get the WSDL for client generation.

Decide on your approach to developing clients and services for WebSphere MQ transport for SOAP.

There are two approaches.

1. Use standard development tools, develop an HTTP service and client, and then use the URL for WebSphere MQ transport for SOAP.
2. Use the tools and samples supplied with WebSphere MQ transport for SOAP.

If you take the HTTP route, you can run the service on an HTTP server and also run it using WebSphere MQ transport for SOAP. To run it using WebSphere MQ transport for SOAP, configure the appropriate WebSphere MQ listener for SOAP and set up the paths and deployment descriptors to run the service. The tools provided by WebSphere MQ transport for SOAP do the configuration for you. Alternatively, you can configure the environment to run the listeners.

The tools supplied with WebSphere MQ transport for SOAP are useful in getting started and learning how to deploy the transport. For production work, there are benefits in using standard tools, and deploying the same service accessible to different SOAP transports.

Decide on the type of client to develop

You must decide what type of Web service client to develop. The choice depends on whether you know the service interface and the address of the service.

If the interface is known, use Axis or .NET tools to generate proxy client classes from the service interface. The proxy client classes make it easier to write a client to call the service. If the location of the service is known when you develop the client, then use the static proxy interface. If the location of the service changes, for example if the service is redeployed onto a production server, then use the dynamic proxy interface.

If the service interface is not known at the time you develop a client, on Axis, you can create a Dynamic Invocation Interface (DII) client for Axis 1.4. A DII client uses a generic interface to call any service. To pass parameters to a particular service correctly, you need to build the specific service interface programmatically. Build the interface programmatically in the client, or by loading the WSDL for the service into the client. On Axis2, you can create a Dispatch client. The Dispatch client uses a document model to describe the client request, whereas a DII client uses a call model. Both work on building the request dynamically.

Obtain the WSDL for the service

Except for the case of the service interface being built programmatically, you must first obtain the service WSDL in order to create a Web service client. The service WSDL is obtainable from three different sources:

1. Directly from the Web service implementation using a tool such as **java2wsdl** (Axis) or **disco** (.NET).
2. By querying the Web service using the URL: *Web service http url?wsdl*.
3. From a file, either on a file system, or from a registry such as UDDI or WebSphere Service Registry and Repository.

Note: If the service is not accessible using HTTP, then the WSDL query does not work. The service itself might only be available using the WebSphere MQ transport for SOAP.

The WSDL generated by **amqwdeployMQService** is not the same as WSDL generated using **java2wsdl** or **disco**. The generated WSDL is also different to any WSDL you might have started with to create the service "Top Down". On Axis, the `server-config.wsdd` deployment descriptor maps the SOAP message produced by a client to an operation and service. **amqwdeployMQService** generates a different deployment descriptor from Eclipse.

The WSDL you use to build clients depends on how the service is deployed:

Deployed using amqwdeployMQService

Use the WSDL generated by **amqwdeployMQService**. Specify the `-w` flag and select `rpcLiteral` WSDL. For compatibility, you can select `rpcEncoded` WSDL. `rpcEncoded` WSDL works only with .NET and Axis 1.4 clients.

Manual deployment using SimpleJavaListener

Use one of the following WSDL files:

1. WSDL used to define the service, or stored in a repository.
2. WSDL generated from the service by **java2wsdl**.
3. WSDL queried using the URL *Web service http url ?wsdl*, if available from an HTTP server. You might run a tool such as Web Services Explorer to import the service definition directly into Eclipse.

You might need to change the URI for the service. Change it from the address of the HTTP service, to the URI for the WebSphere MQ transport for SOAP.

Manual deployment using amqSOAPNETListener.

Use one of the following WSDL files:

1. WSDL used to define the service, or stored in a repository.
2. WSDL obtained from the .NET service class (.asmx). using **disco**.
3. WSDL queried using the URL *Web service http url ?wsdl*, if available. You might run a tool such as Web Services Explorer to import the service definition directly into Eclipse.
4. WSDL obtained by running **amqswsdl** against the .NET service class (.asmx).

You might need to change the URI for the service. Change it from the address of the HTTP service, to the URI for the WebSphere MQ transport for SOAP.

Deployed to Windows Communication Foundation

Obtain the service WSDL by using the URL *Web service http url?wsdl*. The service must be defined with the `serviceMetaData` behavior configuration as part of the service definition.

Deployment to a different server platform.

Follow the guidance provided with the platform about how to obtain the correct service WSDL.

About this task

Develop clients using standard development tools. The following tasks illustrate how to build clients for .NET 1 and 2, Axis 1.4 (JAX-RPC) and Axis2 (JAX-WS). For Windows Communication Foundation, see the related tasks links.

Developing a JAX-RPC client for WebSphere transport for SOAP using Eclipse

Develop an Axis 1.4 Web service client to run using the WebSphere MQ transport for SOAP.

Before you begin

You must have the service available. If you are following the task as a practical exercise, use the workspace and service you created in the task, “Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919. You need to have an application server running in eclipse that supports Axis 1.4 Web services. In this task we use the freely available WebSphere Application Server Community Edition Version 2.1.4. It is configured as part of the task “Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919. You might also use Tomcat 6, which is a smaller open source application server.

About this task

The task shows the development of three types of client for the sample StockQuoteAxis service using Eclipse running on Windows. The clients are a static and dynamic client developed using the client proxy, and a DII client.

Two alternative approaches to generating the client proxies from WSDL are illustrated:

1. Generating client proxies using **amqwdeployWMQService**.
2. Importing WSDL into Eclipse, and using the Web service wizard to generate the client proxies.

Procedure

1. Start the Eclipse IDE for Java EE developers.
2. Create a Java project called StockQuoteAxisClient:
 - a) Switch to the Java perspective > **File > New > Java Project**. In the **Project name** field of the **Create a Java Project page** type, StockQuoteAxisEclipseClient. Make sure that the execution environment is either **J2SE1-1.4** or **J2SE-1.5** > **Next**.
 - b) On the **Java Settings** page, select the **Libraries** tab > **Add External JARs...**
 - c) Browse to *MQ_INSTALLATION_PATH*/java/lib and select all the .jar files > **Open**.
MQ_INSTALLATION_PATH is the directory where WebSphere MQ is installed.
 - d) Browse to *MQ_INSTALLATION_PATH*/java/lib/soap and select all the .jar files > **Open**. You must have installed axis.jar from the WebSphere MQ installation media into this directory.
MQ_INSTALLATION_PATH is the directory where WebSphere MQ is installed.
 - e) The **Library** tab now references all the .jar files needed to build the client > **Finish**.
3. Follow one of these two approaches to create proxies in Eclipse for the sample StockQuoteAxis Web service:
 - Generate the client proxies using **amqwdeployWMQService**.
 - a. Create a queue manager. For the task create QM1 as the default queue manager.
 - b. Create a working directory, samples. Copy the StockQuoteAxis.java sample program into samples/soap/server.
 - c. Modify amqwsetcp.cmd in *MQ_INSTALLATION_PATH*/bin to include the current directory in the classpath. *MQ_INSTALLATION_PATH* is the directory where WebSphere MQ is installed.
 - d. Open a command window in samples and run the modified **amqwsetcp** command
 - e. Create WSDL for the StockQuoteAxis service by running the command,

```
amqwdeployWMQService -f soap/server/StockQuoteAxis.java -c genAxisWsd1
-u "jms:/queue?destination=REQUESTAXIS
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```

Remember: Use "/", rather than "." or "\" when using Java commands.

Tip: Rather than import the generated proxies into Eclipse, you can import the generated WSDL from .samples/generated. The resulting proxies differ in two ways:

- i) The package names are different - which you can refactor.
 - ii) The Eclipse generated proxies include an additional helper class, StockQuoteAxisProxy.java
- f. Create the client proxies for the StockQuoteAxis service by running the command:

```
amqwdployWMQService -f soap/server/StockQuoteAxis.java -c genProxiestoAxis  
-u "jms:/queue?destination=REQUESTAXIS  
&initialContextFactory=com.ibm.mq.jms.Nojndi  
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```

g. Import the client proxies into StockQuoteAxisClient:

- i) Right click **StockQuoteAxisClient\src** > Select **File System** > **Next** > **Browse...** > find the folder .\samples\generated\client\remote\soap\server > **OK**.
 - ii) Check **server** in the **Import** page > **Finish**.
- h. Refactor the package name to soap.server.
- i) Right click the package containing the client proxies > **Refactor** > **Rename**. Type the **New name:** soap.server > leave the selected defaults for the other choices > **OK**. All the errors are fixed.

- Generate the client proxies using Eclipse.

You have a choice of ways of obtaining the WSDL for the service. In this example, the service has been deployed to WebSphere Application Server Community Edition and you obtain the WSDL from the Web server. The deployment is described in the task [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#),

- a. In Eclipse, switch to the Web perspective, and check that the WebSphere Application Server Community Edition v2.1 Server is running and StockQuoteAxis is deployed and synchronized.
- b. Import the WSDL into Web Services Explorer:

- i) Click the **Web Services Explorer** icon in the action-bar, or click **Run** > **Launch the Web Services Explorer**.
- ii) Click the WSDL page icon in the Web Services Explorer to switch to the WSDL page.
- iii) Click **WSDL Main** in the Navigator window of the Web Services Explorer.
- iv) Type in the URL of the Web service, followed by ?WSDL. The URL for StockQuoteAxis, deployed in the task [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#), is:

```
http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl
```

c. Generate the client proxies:

- i) In the Web Services Explorer navigator, click **http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl**.
- ii) In the **Actions** window, click **Launch Web Service Wizard** > leave **Web Service Client** selected > **Go**.
- iii) On the first page of the wizard, click the **Client** project link in the configuration > Select the **StockQuoteAxisClient** client project > **OK**.

Tip: The wizard window might lose focus. You need to bring it back into focus manually.

- iv) The Web service runtime must be Apache Axis to generate a JAX-RPC client.
- v) Click **Finish**.

vi) Change the static URL of the service to point to the WebSphere MQ transport for SOAP address for the StockQuoteAxis service. You might choose to skip this step, until you have tested the client with an HTTP server.

a) Open `StockQuoteAxisServiceLocator.java` and find the declaration for `StockQuoteAxis_address`.

b) Change the URL to

```
"jms:/queue?destination=REQUESTAXIS
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QM1)binding(auto))"
```

Tip: Eclipse automatically transforms `&` to `&#amp;`, and the reverse, when you copy and paste strings into `.java` code.

d. Create three Java client classes, each with a main method:

i) Create package. Right click **StockQuoteAxisClient/src** > **New Package**. Name it `soap.client` > **Finish**.

ii) Select **soap.client** > **New** > **Class**. Name the class `SQASStaticClient` > Check **public static void main(string [] args)** > **Finish**

iii) Repeat the procedure to create `SQADynamicClient.java` and `SQADIIClient.java`

e. Write the client code.

Figure 177 on page 936 through Figure 181 on page 938 provide examples of the three styles of client code. The examples use an HTTP URL to test the client using the StockQuoteAxis service deployed to an HTTP server. To run the clients against the StockQuoteAxis service deployed using WebSphere MQ transport for SOAP, change the URL to:

```
"jms:/queue?destination=REQUESTAXIS
connectionFactory=(connectQueueManager(QM1)binding(auto))
initialContextFactory=com.ibm.mq.jms.Nojndi
targetService=soap.server.StockQuoteAxis.java
replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE"
```

– Figure 177 on page 936 and Figure 179 on page 937 use the proxy generated by Eclipse, which has the extra `StockQuoteAxisproxy` helper class that makes coding a little easier.

– Figure 178 on page 937 and Figure 180 on page 937 use the proxy generated by **amqwdeployWMQService**.

– Figure 181 on page 938 uses no proxy classes.

Each of the clients call `com.ibm.mq.soap.Register.extension()` to link to the WebSphere MQ transport for SOAP. The extension is registered in the client deployment descriptor. Client deployment to Axis 1.4 is described in [“Deploying a Web service client to Axis 1.4 to use IBM WebSphere MQ transport for SOAP”](#) on page 967.

f. Run the clients by sending the SOAP request to StockQuoteAxis hosted by the WebSphere Application Server Community Edition server configured in the workspace.

i) Check that the server is running, StockQuoteAxis is deployed and synchronized.

ii) Select or open the client you want to test > Click **Run** in the action bar. Alternatively, click the green Run icon or right click the client in the navigator > **Run As** > **Run Configurations** Configure the parameters you require to run the client.

g. Run the client using the WebSphere MQ transport for SOAP.

The procedure uses **amqwdeployWMQService** to deploy the service, and only works with the client that uses the WSDL or proxies built by **amqwdeployWMQService**. To run the client using the original WSDL, or proxies built by eclipse, deploy the service with its deployment descriptor built by Eclipse. Manually start **SimpleJavaListener** using the service port binding name as the `targetServiceName`.

- i) Follow the instructions in [“Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService”](#) on page 956 to deploy the service to the WebSphere MQ Simple Java SOAP listener. The service deployment works only for the client using the WSDL or client proxies built by **amqwdeployWMQService**.
- ii) In a command window, run **amqwclientconfig** to create the client deployment descriptor file, `client-deploy.wsdd`.
- iii) Import `client-deploy.wsdd` into the root of the Java project you want to test using WebSphere MQ transport for SOAP.
 - a) Right click the Java project **StockQuoteAxisEclipseClient** > **Import** > **File system** > **Next** > **Browse...**
 - b) Browse to the directory containing `client-deploy.wsdd` > **Open** > Select the directory in the **Import** wizard page > check `client-deploy.wsdd` in the right pane.
 - c) Verify **Into folder:** has `StockQuoteAxisEclipseClient` entered > **Finish**.
- iv) Confirm that the working directory for running a Java application in this project is the `StockQuoteAxisEclipseClient` directory:

Right click the Java project **StockQuoteAxisEclipseClient** > **Run as....** > **Run Configurations...** > Select the (x)= **Arguments** tab > Verify that in Working Directory the **Default** radio button is checked, and the path is `StockQuoteAxisEclipseClient`. Alternatively make one of the following choices to select a different location or file containing the client configuration:

 - Check **Other:** > type a directory path of your choice.
 - In the **VM arguments** window, type `-Daxis.ClientConfigFile=full path to client deployment descriptor file`
- v) Make sure the URL is configured to point to the service deployed using WebSphere MQ transport for SOAP. Run the client as described in step ii.

Tip: Typically, you might encounter one of these errors:

- i) Exception: No client transport named 'jms' found!.
- ii) A JMS connection error.
- iii) Exception: The AXIS engine could not find a target service to invoke! `targetService is soap.server.StockQuoteAxis.java`
- iv) Exception: `java.lang.InstantiationException: soap.server.StockQuoteAxis`

Explanations:

- i) `client-config.wsdd` is not found, or does include the line `<transport name="jms" pivot="java:com.ibm.mq.soap.transport.jms.WMQSender"/>` in `client-config.wsdd`.
- ii) Possibly a build path problem - not including the .jar files in `MQ_INSTALLATION_PATH/java/lib`. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.
- iii) Service deployment problem, either with `server-config.wsdd`, or with parameters passed to **SimpleSoapListener**.
- iv) Mismatch between deployment descriptor and the implementation of the service.

If you are having difficulty running the client in Eclipse, try using a command window:

- i) Switch to the `StockQuoteAxisEclipseClient\bin` directory in the workspace directory tree.
- ii) Run **amqwsetcp** and **amqwclientconfig**
- iii) Run `java soap/client/SQASStaticClient`.

Sample JAX-RPC Web service clients

The sample Java Web service clients shipped with WebSphere MQ are installed in `MQ_INSTALLATION_PATH\tools\soap\samples\java\clients`. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.

SQAxis2Axis.java

`SQAxis2Axis.java`, [Figure 174 on page 935](#), is a dynamic proxy client to invoke the `StockQuoteAxis` service. You can override the URL of the service, which is compiled into the dynamic proxy, by providing a URL on the command line.

SQAxis2DotNet.java

`SQAxis2DotNet.java`, [Figure 175 on page 935](#), is a dynamic proxy client to invoke the `StockQuoteDotNet` service. You can override the URL of the service, which is compiled into the dynamic proxy, by providing a URL on the command line.

Wsd1Client.java

`Wsd1Client.java`, [Figure 176 on page 936](#), is a dynamic invocation client to invoke either the `StockQuoteDotNet` or `StockQuoteAxis` service. The client invokes the `StockQuoteAxis` service by default. Add the command-line option `-D` to invoke the `StockQuoteDotNet` service and `-w` to provide a different port to the one in `.\generated\StockQuoteDotNet_Wmq.wsdl`

```
package soap.clients;
import java.net.URL;
import soap.server.*;
public class SQAxis2Axis {
    public static void main(String[] args) {
        com.ibm.mq.soap.Register.extension();
        try {
            StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
            StockQuoteAxis service = null;
            if (args.length == 0)
                service = locator.getSoapServerStockQuoteAxis_Wmq();
            else
                service = locator.getSoapServerStockQuoteAxis_Wmq(
                    new java.net.URL(args[0]));
            System.out.println("Response: " + service.getQuote("XXX"));
        } catch (Exception e) {
            System.out.println("\n>>> EXCEPTION WHILE RUNNING ProxyClient DEMO <<<\n");
            e.printStackTrace();
            System.exit(2);
        }
    }
}
```

Figure 174. `SQAxis2Axis.java`

```
public class SQAxis2DotNet {
    public static void main(String[] args) {
        com.ibm.mq.soap.Register.extension();
        try {
            StockQuoteDotNet locator = new StockQuoteDotNetLocator();
            StockQuoteDotNetSoap_PortType service = null;
            if (args.length == 0)
                service = locator.getStockQuoteDotNetSoap();
            else
                service = locator.getStockQuoteDotNetSoap(new java.net.URL(
                    args[0]));
            System.out.println("Response: " + service.getQuoteDOC("XXX"));
        } catch (Exception e) {
            System.out.println("\n>>> EXCEPTION WHILE RUNNING ProxyClient DEMO <<<\n");
            e.printStackTrace();
            System.exit(2);
        }
    }
}
```

Figure 175. `SQAxis2DotNet.java`

```

package soap.clients;
import com.ibm.mq.soap.*;
import org.apache.axis.utils.Options;
import java.net.URL;
import javax.xml.rpc.Call;
import javax.xml.rpc.Service;
import javax.xml.rpc.ServiceFactory;
import javax.xml.namespace.QName;
public class WsdIClient {
public static void main(String[] args) {
    String wsdlService, wsdlPort, namespace, wsdlSource, wsdlTargetURI, s;
    try {
        Register.extension();
        Options opts = new Options(args);
        if (opts.isFlagSet('D') != 0) {
            wsdlService = "StockQuoteDotNet";
            wsdlPort = "StockQuoteDotNetSoap";
            namespace = "http://stock.samples";
            wsdlSource = "file:generated/StockQuoteDotNet_Wmq.wsdl";
        } else {
            wsdlService = "StockQuoteAxisService";
            wsdlPort = "soap.server.StockQuoteAxis_Wmq";
            namespace = "soap.server.StockQuoteAxis_Wmq";
            wsdlSource = "file:generated/soap.server.StockQuoteAxis_Wmq.wsdl";
        }
        if (null != (s = (opts.isValueSet('w'))))
            wsdlPort = s;
        System.out.println("start WsdIClient demo, wsdl port " + wsdlPort
            + " resolving uri to ...");
        QName servQN = new QName(namespace, wsdlService);
        QName portQN = new QName(namespace, wsdlPort);
        Service service = ServiceFactory.newInstance().createService(
            new URL(wsdlSource), servQN);
        Call call = (Call) service.createCall(portQN, "getQuote");
        wsdlTargetURI = call.getTargetEndpointAddress().toString();
        System.out.println("'" + wsdlTargetURI + "'");
        Object ret = call.invoke(new Object[] { "XXX" });
        System.out.println("Response: " + ret);
    } catch (Exception e) {
        System.out.println("\n>>> EXCEPTION WHILE RUNNING WsdIClient DEMO <<<\n");
        e.printStackTrace();
        System.exit(2);
    }
}
}
}

```

Figure 176. WsdIClient.java

The example clients used in this task:

```

package soap.client;
import soap.server.StockQuoteAxisProxy;
public class SQAStaticClient {
    public static void main(String[] args) {
        try {
            com.ibm.mq.soap.Register.extension();
            StockQuoteAxisProxy sqa = new StockQuoteAxisProxy();
            System.out.println("Static client synchronous result is:"
                + sqa.getQuote("ibm"));
        } catch (Exception e) {
            System.out.println("Exception: " + e);
        }
    }
}
}

```

Figure 177. Static client using Eclipse generated proxy


```

package soap.client;
import soap.server.StockQuoteAxis;
import soap.server.StockQuoteAxisService;
import soap.server.StockQuoteAxisServiceLocator;
public class SQASStaticClient {
    public static void main(String[] args) {
        try {
            com.ibm.mq.soap.Register.extension();
            StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
            StockQuoteAxis sqa = locator.getSoapServerStockQuoteAxis_Wmq();
            System.out.println("Static client synchronous result is: "
                + sqa.getQuote("ibm"));
        } catch (Exception e) {
            System.out.println("Exception: " + e);
        }
    }
}

```

Figure 178. Static client using amqwdployWMQService generated proxy

```

package soap.client;
import soap.server.StockQuoteAxisProxy;
public class SQADynamicClient {
    public static void main(String[] args) {
        try {
            com.ibm.mq.soap.Register.extension();
            StockQuoteAxisProxy sqa = new StockQuoteAxisProxy(
                "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
            System.out.println("Dynamic client synchronous result is: "
                + sqa.getQuote("ibm"));
        } catch (Exception e) {
            System.out.println("Exception: " + e);
        }
    }
}

```

Figure 179. Dynamic client using Eclipse generated proxy

```

package soap.client;

import java.net.URL;
import soap.server.StockQuoteAxis;
import soap.server.StockQuoteAxisService;
import soap.server.StockQuoteAxisServiceLocator;
public class SQADynamicClient {
    public static void main(String[] args) {
        try {
            com.ibm.mq.soap.Register.extension();
            URL sqaURL = new URL(
                "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
            StockQuoteAxisService locator = new StockQuoteAxisServiceLocator();
            StockQuoteAxis sqa = locator.getSoapServerStockQuoteAxis_Wmq(sqaURL);
            System.out.println("Dynamic client synchronous result is: "
                + sqa.getQuote("ibm"));
        } catch (Exception e) {
            System.out.println("Exception: " + e);
        }
    }
}

```

Figure 180. Dynamic client using amqwdployWMQService generated proxy

```

package soap.client;
import java.net.URL;
import javax.xml.namespace.QName;
import javax.xml.rpc.Call;
import javax.xml.rpc.Service;
import javax.xml.rpc.ServiceFactory;
public class SQADIIClient {
    public static void main(String[] args) {
        try {
            com.ibm.mq.soap.Register.extension();
            URL wsdl = new URL(
                "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl");
            Service SQAService = (ServiceFactory.newInstance()).createService(wsdl,
                new QName("http://server.soap", "StockQuoteAxisService"));
            Call SQACall = SQAService.createCall(new QName("http://server.soap",
                "StockQuoteAxis"), "getQuote");
            System.out.println("DII client synchronous result is "
                + SQACall.invoke(new Object[] { "ibm" }));
        } catch (Exception e) {
            System.out.println("Exception: " + e);
        }
    }
}

```

Figure 181. DII client (No proxy)

Related tasks

[Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse](#)

Develop an Axis2 Web service client to run using the WebSphere MQ transport for SOAP. The sample Axis2 clients provided with WebSphere MQ transport for SOAP are listed, and the **wsimport** command used to generate proxies.

[Developing a .NET 1 or 2 client for WebSphere transport for SOAP using Microsoft Visual Studio 2008](#)

Develop an .NET 1 or 2 Web service client to run using the WebSphere MQ transport for SOAP.

Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse

Develop an Axis2 Web service client to run using the WebSphere MQ transport for SOAP. The sample Axis2 clients provided with WebSphere MQ transport for SOAP are listed, and the **wsimport** command used to generate proxies.

Before you begin

Obtain the Axis2 libraries, and configure a development and test environment to run the client.

Note: The naming of versions and releases used by Axis causes confusion. Typically, Axis 1.4 refers to the JAX-RPC implementation, and Axis2 to the JAX-WS implementation.

Axis 1.4 is a version level. If you search for Axis 1.4 on the internet, you are taken to <http://ws.apache.org/axis/>. The page contains a list of preceding versions of Axis (1.2, 1.3) and the April 22, 2006, final release of Axis 1.4. There are later releases of Axis 1.4, that fix bugs, but they are all known as Axis 1.4. It is one of these bug fix releases that is shipped with WebSphere MQ. For Axis 1.4, use the version of `axis.jar` that is shipped with WebSphere MQ rather than the one obtainable from <http://ws.apache.org/axis/>.

The Axis website also refers to Axis 1.1 to refer to all the versions of what is more typically called Axis 1.4. Axis 1.2 is used to refer to what is typically called Axis2.

Axis 1.5 is not a later release of Axis 1.4, it is an Axis2 release. If you search for Axis 1.5 you are directed to <http://ws.apache.org/axis2/>. <https://ws.apache.org/axis2/download.cgi> contains a list of release versions of Axis2, labeled 0.9 to 1.5.1 (and including, confusingly version 1.4). The release version of Axis2 to use with WebSphere MQ transport for SOAP is 1.4.1. Download Axis2 1.4.1 from http://ws.apache.org/axis2/download/1_4_1/download.cgi.

You can choose to generate proxies for the Web service clients for WebSphere MQ transport for SOAP using either **wsimport** or the tooling provided with an IDE. Eclipse IDE for Java EE Developer 3.5 SR1 uses **wsdl2java.wsimport** is supplied with Java 6. You can use Java 5 to run client proxies generated either with **wsimport** or **wsdl2java**.

The sample Web service Axis2 clients provided with WebSphere MQ transport for SOAP were developed using **wsimport**; see [“Sample Axis2 clients” on page 943](#).

The task that follows demonstrates how to generate and use the proxies produced by the Web services wizard that is packaged with Eclipse IDE for Java EE Developers. The sample clients show how to use the proxies produced by **wsimport**.

To use the Web services wizard, you must add an application server that supports Axis2 to the workbench. The steps show how to configure WASCE to support Axis2 using the workbench.

1. Configure the application server used in Eclipse IDE for Java EE Developers to support Axis2. In this example, configure the WASCE 2.1.4, application server, which is part of the workspace created in [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#).
 - a. Open the workspace preferences to configure the server: Open **Window > Preferences**.
 - b. Check the installed JRE is Java50: Click **Installed JREs**.
 - c. Add WASCE as the server: Click **Server > Runtime environments > Add... > IBM > WASCE v2.1 > Next**. The JRE must be Java50 > Browse to the WASCE installation directory > **OK > Finish**. You must have installed the WASCE plug-in for Eclipse Java EE IDE for Web Developers.
 - d. Add Axis2: Click **Web Services > Axis2 Preferences**. On the **Axis2 Runtime** tab > **Browse...** Open the directory containing the many Axis2 jar files > **Apply**.
 - e. Associate WASCE with Axis2: Click **Web Services > Server and Runtime**. Under **Server** select **IBM WASCE v2.1 Server**, and under **Web service runtime**, select **Apache Axis2 > Apply > OK**
 - f. Start the server: Open the Web perspective and open the Servers view. Right click in the Servers view > **New > Server**. **IBM WASCE v2.1 Server** is selected and configured > **Finish**. Start the server.
2. Check that you have deployed the StockQuoteAxis service to WASCE to run the Web service wizard.
3. To test the service with the WebSphere MQ transport for SOAP service, deploy the service to a WebSphere MQ transport for SOAP listener for Axis 1.4; see [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#).

About this task

The Eclipse IDE for Java EE Developers uses Java50 and the Web services wizard to generate the proxy classes for the service. The proxy classes are different to the classes created by the **wsimport** tool provided with Java 6. An alternative approach is to generate the proxy classes using **wsimport** and import the packages it creates into your Eclipse Java EE IDE for Web Developers.

The Web services wizard in the Eclipse IDE for Java EE Developers builds a Web service client in a Web project. You can run the client as a simple Java application; it does not require an application server. You can also transfer the code to a Java project, and configure the build path to include the Axis2 JAR files.

Procedure

1. Create a Web project in a new Enterprise project:
 - a) With nothing selected in the Project Explorer > Right-click the white space > **New > Enterprise Application Project** > Name it StockQuoteAxis2EAR > **Finish**. Reply No to the window giving you the option of opening the Java EE perspective.
The defaults are set to use WASCE.
 - b) Right-click StockQuoteAxis2EAR > **New > Dynamic Web Project**. Name the project StockQuoteAxis2WebClient > Check the EAR membership box to add the project to **StockQuoteAxis2EAR**. WASCE 2.1 is selected as the Target runtime.
 - c) In the Configuration section of the **New Dynamic Web Project** page > **Modify...** > Check the Axis2 Web services project facet. **Dynamic Web Module 2.5, Java 6.0**, and **WASCE deployment 1.2** are already checked. > **OK > Finish**. Reply No to the window giving you the option of opening the Java EE perspective.

2. Import WSDL for the service into the workspace and generate the client proxy:

In this example, the WSDL document contains the HTTP service binding and becomes the target for the static Web client proxy. You can modify the URL in the Web service binding to point to the WebSphere MQ transport for SOAP URL before generating the client proxy. The static Web client proxy is then the service that is deployed to WebSphere MQ transport for SOAP.

- a) Launch the Web Services Explorer: either use the icon in the action bar, or **Run > Launch the Web Services Explorer**.
 - b) Select the WSDL explorer by clicking the WSDL icon in the **Web Services Explorer** window > Click **WSDL Main** in the Navigator window > Type the URL of the StockQuoteAxis WSDL file > **Go**.
In this example, obtain the WSDL directly from the HTTP service: `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis?wsdl`
 - c) In the Navigator, click the line with the URL of the Web service. In the **Actions** window, click **Import WSDL to Workbench** > Select a **StockQuoteAxis2WebClient** as the **Workbench project** > Type the **WSDL file name**, `StockQuoteAxisHTTP.wsdl` > **Go**.
 - d) Right-click **StockQuoteAxisHTTP.wsdl** > **Web Services** > **Generate Client**. Check the configuration information about the Web services page of the wizard is as follows: Server: IBM WASCE v2.1
Server, Web service runtime: Apache Axis2, Client project: StockQuoteAxis2WebClient, Client EAR project: StockQuoteAxisEAR. To correct the configuration, click the lines that are wrong.
 - e) Click **Next** > verify the code generation settings > **Finish**.
Notice that a new package, `soap.server`, is created and it contains the proxies you require.
3. Configure the project to run WebSphere MQ transport for SOAP as the JMS transport.
- WebSphere MQ transport for SOAP provides a `transportSender`, but no `transportReceiver`. In other words, WebSphere MQ transport for SOAP supports Axis2 clients. Currently it does not support Axis2 services.
- a) In the **StockQuoteAxis2WebClient** project, right-click `WebContent\WEB-INF\conf\axis2.xml` > **Open with...** > **XML editor**.
 - b) Search for the last `transportSender` (towards the end of the file) and find the commented out JMS `transportSender` > Right-click the line > **Add before...** > **transportSender**.
 - c) Right-click **transportSender** > **Add Attribute** > **Name** > Right-click **transportSender** > **Add Attribute** > **Class**.
 - d) Right-click **Name** > **Edit Attribute** > Type the **Value**: `jms`
 - e) Right-click **Class** > **Edit Attribute** > Type the **Value**:
`com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender`. > **Save**.
 - f) Add `com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender` to the build path: Right-click **StockQuoteAxis2WebClient** > **Build Path** > **Configure Build Path...** > Click the **Libraries** tab > **Add External JARs...** Select all the JARs in `MQ_INSTALLATION_PATH\java\lib` > **OK**.
`MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.
4. Create a synchronous static client, test it using HTTP, then convert the proxy to run the static client using WebSphere MQ transport for SOAP.
- a) Right-click **Java Resources: src** > **New** > **Package** > Name the package `soap.client` > **Finish**
 - b) Right-click **soap.client** > **New** > **Class** > Name the class `SQA2StaticClient` > **Finish**.
 - c) Replace the class with the following code, then click **Save**.

Figure 182. *SQA2DynamicClient.java*

```
package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
public class SQA2StaticClient {
    public static void main(String[] args) {
        try {
            StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub();
```

```

        GetQuote request = new GetQuote();
        request.setSymbol("ibm");
        System.out.println("Response is: "
            + (stub.getQuote(request)).getGetQuoteReturn());
    } catch (Exception e) {
        System.out.println("Exception: " + e.getMessage());
        e.printStackTrace();
    }
}
}

```

5. Test the client with the StockQuoteAxis service deployed to WASCE, and with WebSphere MQ transport for SOAP.

- a) In the Project Explorer, right-click **SQA2StaticClient** > **Run as...** > **Java Application**.

The result, Response is 55.25, appears in the Console view. You can also select the WASCE console window in the Console view, and see the output on the WASCE server, StockQuoteAxis called with parameter: ibm.

- b) The proxy was built with the service address, `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis`, and so the static client calls the service running on HTTP. You can change the static client to call the service using WebSphere MQ transport for SOAP. The following instructions change the service address in `StockQuoteAxisServiceStub.java` without rebuilding the proxy, and configure the `SQA2StaticClient` runtime parameters to load `axis2.xml`. You configure `axis2.xml` configures Axis2 to use WebSphere MQ transport for SOAP.

- c) Open `StockQuoteAxisServiceStub.java` >

Replace the two occurrences of `http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis` with,

```

jms:/queue?destination=REQUESTAXIS@QM1
&connectionFactory=()
&initialContextFactory=com.ibm.mq.jms.Nojndi
&targetService=StockQuoteAxis
&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE

```

- d) If you run `SQA2StaticClient` now, it throws an exception because it has not found a `transportSender` configured for JMS
The exception is:

```

Exception: null java.lang.NullPointerException at
soap.server.StockQuoteAxisServiceStub.getQuote(StockQuoteAxisServiceStub.java:547)
at soap.client.SQA2StaticClient.main(SQA2StaticClient.java:11)

```

- e) In the Project Explorer, right-click **SQA2StaticClient** > **Run as...** > **Run Configurations....** Switch to the **(x)= Arguments** tab, and in the **VM arguments** input area, type the path to the `axis2.conf` file > **Apply** > **Run**.

The VM argument is: `-Daxis2.xml=${workspace_loc:StockQuoteAxis2WebClient/WebContent/WEB-INF/conf}/axis2.xml`. Or you can provide a standard path to the Axis2 configuration file.

- f) Run `SQA2StaticClient` again. On this run, you are using the WebSphere MQ transport for SOAP. Confirm it by checking there is no new output in the WASCE console. Open the console or command window that is associated with `SimpleJavaListener`, and the output there is `StockQuoteAxis` called with parameter: ibm.

6. Create a dynamic client for HTTP and WebSphere MQ transport for SOAP, and test it.

- a) Right-click **soap.client** > **New** > **Class** > Name the class `SQA2DynamicClient` > **Finish**.

- b) Replace the class with the following code, then click **Save**.

```

package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
public class SQA2DynamicClient {
    public static void main(String[] args) {
        try {
            StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub(
                "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
            GetQuote request = new GetQuote();

```

```

        request.setSymbol("ibm");
        System.out.println("HTTP Sync: "
            + (stub.getQuote(request)).getGetQuoteReturn());
        stub = new StockQuoteAxisServiceStub(
            "jms:/queue?destination=REQUESTAXIS@QM1"
            + "&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.NoJndi"
            + "&targetService=StockQuoteAxis&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE");
        System.out.println("JMS sync: "
            + (stub.getQuote(request)).getGetQuoteReturn());
    } catch (Exception e) {
        System.out.println("Exception: " + e.getMessage());
        e.printStackTrace();
    }
}
}

```

- c) Create a Run configuration for `SQA2DynamicClient.java`, and add the path to `axis2.xml`:
`-Daxis2.xml=${workspace_loc:StockQuoteAxis2WebClient/WebContent/WEB-INF/conf}/axis2.xml`
 - d) Run `SQA2DynamicClient`. Check the console output for the `SQA2DynamicClient`, `WASCE` and **SimpleJavaListener**.
7. Create an asynchronous client, and access the result in a callback handler, and in the main program thread.

The asynchronous client proxies created by the Web service wizard for Eclipse Java EE IDE for Web Developers differ from the proxies created by **wsimport**. The **wsimport** proxies use `Future`, `Response`, and `AsyncHandler` generic types.

The Web service wizard for Eclipse Java EE IDE for Web Developers creates a `StockQuoteAxisServiceCallbackHandler` abstract class. You must extend `StockQuoteAxisServiceCallbackHandler` and create a callback handler.

- a) Right-click **soap.client** > **New** > **Class** > Name the class `SQA2CallbackHandler` > **Finish**.
- b) Replace the class with the following code.

```

package soap.client;
import soap.server.StockQuoteAxisServiceCallbackHandler;
import soap.server.StockQuoteAxisServiceStub.GetQuoteResponse;
public class SQA2CallbackHandler
    extends StockQuoteAxisServiceCallbackHandler {
    private boolean complete = false;
    SQA2CallbackHandler() {
        super();
        System.out.println("Callback constructor");
    }
    public void receiveResultgetQuote(GetQuoteResponse response) {
        System.out.println("Result in Callback " + response.getGetQuoteReturn());
        super.clientData = response;
        complete = true;
    }
    public boolean isComplete() {
        return complete;
    }
}

```

- c) Right-click **soap.client** > **New** > **Class** > Name the class `SQA2AsyncClient` > **Finish**.
- d) Replace the class with the following code.

Figure 183. SQA2AsyncClient.java

```

package soap.client;
import soap.server.StockQuoteAxisServiceStub;
import soap.server.StockQuoteAxisServiceStub.GetQuote;
import soap.server.StockQuoteAxisServiceStub.GetQuoteResponse;
import soap.server.StockQuoteAxisServiceCallbackHandler;
@SuppressWarnings("unused")
public class SQA2AsyncClient {
    public static void main(String[] args) {
        try {
            StockQuoteAxisServiceStub stub = new StockQuoteAxisServiceStub(

```

```

        "http://localhost:8080/StockQuoteAxis/services/StockQuoteAxis");
        GetQuote request = new GetQuote();
        request.setSymbol("ibm");
        System.out.println("HTTP Sync: "
            + (stub.getQuote(request)).getGetQuoteReturn());
        SQA2CallbackHandler callback = new SQA2CallbackHandler();
        stub.startgetQuote(request, callback);
        do {
            System.out.println("Waiting for HTTP callback");
            Thread.sleep(2000);
        } while (!callback.isComplete());
        System.out.println("HTTP poll: "
            + ((GetQuoteResponse) (callback.getClientData()))
            .getGetQuoteReturn());
        stub = new StockQuoteAxisServiceStub(
            "jms:/queue?destination=REQUESTAXIS@QM1"
            + "&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.NoJndi"
            + "&targetService=StockQuoteAxis&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE");
        System.out.println("JMS Sync: "
            + (stub.getQuote(request)).getGetQuoteReturn());
        callback = new SQA2CallbackHandler();
        stub.startgetQuote(request, callback);
        while (!callback.isComplete()) {
            System.out.println("Waiting for JMS callback");
            Thread.sleep(2000);
        }
        System.out.println("JMS poll: "
            + ((GetQuoteResponse) (callback.getClientData())).getGetQuoteReturn());
    } catch (Exception e) {
        System.out.println("Exception: " + e.getMessage());
        e.printStackTrace();
    }
}
}
}

```

The console output is as follows:

```

HTTP Sync: 55.25
Callback constructor
Waiting for HTTP callback
Result in Callback 55.25
HTTP poll: 55.25
JMS Sync: 55.25
Callback constructor
Waiting for JMS callback
Result in Callback 55.25
JMS poll: 55.25

```

Sample Axis2 clients

The sample proxies are generated using the **wsimport** tool that is packaged with Java 6. Six samples are provided:

1. [DynamicProxyClientSync.java](#)
2. [DynamicProxyClientAsyncPolling.java](#)
3. [DynamicProxyClientAsyncCallback.java](#)
4. [DispatchClientSync.java](#)
5. [DispatchClientAsyncPolling.java](#)
6. [DispatchClientAsyncCallback.java](#)

The client samples are generated for the sample StockQuoteAxis server. Generate the WSDL with the **amqwdpoyWMQServer** command, specifying the -w switch to select rpcLiteral style. Use the following command to generate the proxies for the samples:

```
wsimport soap.server.StockQuoteAxis_Wmq.wsdl -d generated -keep -p com.ibm.mq.axis2.samples
```

Figure 184. *DynamicProxyClientSync.java*

```
package com.ibm.mq.axis2.samples;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientSync {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DynamicProxyClientSync");

            System.out.println("Creating proxy instance for service StockQuoteAxisService");
            StockQuoteAxisService stub = new StockQuoteAxisService();
            StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

            System.out.println("Invoking getQuoteOneWay OneWay operation synchronously...");
            service.getQuoteOneWay("48");
            System.out.println(" > getQuoteOneWay has returned");

            System.out.println("Invoking getQuote Request Reply operation synchronously...");
            float result = service.getQuote("48");
            System.out.println(" > getQuote has returned result of " + result);

            System.out.println("End of sample");
        }
        catch (Exception fault) {
            // Identify the cause of the Axis Fault
            System.err.println(fault.toString());
            Throwable e = fault.getCause();
            for (int i = 1; e != null; i++) {
                // The toString method on an MQAxisException will cause the message, explanation and
                user // action.
                System.err.println("Exception(" + i + "): " + e.toString());

                if (e.getCause() != null) {
                    e = e.getCause();
                }
                else {
                    break;
                }
            } // end of for loop
        } // end of catch block
    }
}
```

Figure 185. *DynamicProxyClientAsyncPolling.java*

```
package com.ibm.mq.axis2.samples;

import java.util.concurrent.CancellationException;
import javax.xml.ws.Response;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientAsyncPolling {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DynamicProxyClientAsyncPolling");

            System.out.println("Creating proxy instance for service StockQuoteAxisService");
            StockQuoteAxisService stub = new StockQuoteAxisService();
            StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

            System.out
                .println("Invoking getQuoteAsync Request Reply operation asynchronously by
polling...");
            Response<Float> response = service.getQuoteAsync("49");
```



```

    /** Sleep main thread until response arrives */
    System.out.println("Waiting for response to arrive...");
    while (!response.isDone()) {
        Thread.sleep(100);
    }
    System.out.println(" > Response received");

    /** Retrieve the result */
    try {
        Float result = response.get();
        System.out.println(" > getQuoteAsync call has returned result of " + result);
    }
    catch (CancellationException ce) {
        // processing was cancelled via response.cancel()
    }

    System.out.println("End of sample");
}
catch (Exception fault) {
    // Identify the cause of the Axis Fault
    System.err.println(fault.toString());
    Throwable e = fault.getCause();
    for (int i = 1; e != null; i++) {
        // The toString method on an MQAxisException will cause the message, explanation and
user      // action.
        System.err.println("Exception(" + i + "): " + e.toString());

        if (e.getCause() != null) {
            e = e.getCause();
        }
        else {
            break;
        }
    } // end of for loop
} // end of catch block
}
}

```

Figure 186. *DynamicProxyClientAsyncCallback.java*

```

package com.ibm.mq.axis2.samples;

import java.util.concurrent.Future;

import javax.xml.ws.AsyncHandler;
import javax.xml.ws.Response;

import com.ibm.mq.axis2.samples.proxy.StockQuoteAxis;
import com.ibm.mq.axis2.samples.proxy.StockQuoteAxisService;

public class DynamicProxyClientAsyncCallback implements AsyncHandler<Float> {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DynamicProxyClientAsyncCallback");

            System.out.println("Creating proxy instance for service StockQuoteAxisService");
            StockQuoteAxisService stub = new StockQuoteAxisService();
            StockQuoteAxis service = stub.getSoapServerStockQuoteAxisWmq();

            DynamicProxyClientAsyncCallback handler = new DynamicProxyClientAsyncCallback();

            System.out
                .println("Invoking getQuoteAsync Request Reply operation asynchronously using a
callback...");
            Future<?> monitor = service.getQuoteAsync("50", handler);
            System.out.println(" > Invoke call has returned");

            /** Sleep main thread until handler has been notified */
            System.out.println("Waiting for handler to be called...");
            while (!monitor.isDone()) {
                Thread.sleep(100);
            }

            System.out.println("End of sample");
        }
    }
}

```

```

    }
    catch (Exception fault) {
        // Identify the cause of the Axis Fault
        System.err.println(fault.toString());
        Throwable e = fault.getCause();
        for (int i = 1; e != null; i++) {
            // The toString method on an MQAxisException will cause the message, explanation and
user
            // action.
            System.err.println("Exception(" + i + "): " + e.toString());

            if (e.getCause() != null) {
                e = e.getCause();
            }
            else {
                break;
            }
        } // end of for loop
    } // end of catch block
}

public void handleResponse(Response<Float> response) {
    try {
        Float result = response.get();
        System.out.println(" > Async Handler has received a result of " + result);
    }
    catch (Exception fault) {
        // Identify the cause of the Axis Fault
        System.err.println("Exception in handleResponse");
        System.err.println(fault.toString());
        Throwable e = fault.getCause();
        for (int i = 1; e != null; i++) {
            // The toString method on an MQAxisException will cause the message, explanation and
user
            // action.
            System.err.println("Exception(" + i + "): " + e.toString());

            if (e.getCause() != null) {
                e = e.getCause();
            }
            else {
                break;
            }
        } // end of for loop
    } // end of catch block
}
}
}

```

Figure 187. *DispatchClientSync.java*

```

package com.ibm.mq.axis2.samples;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;
import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientSync {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DispatchClientSync");

            String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
                + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
                +
            "&initialContextFactory=com.ibm.mq.jms.NoJndi&targetService=soap.server.StockQuoteAxis.java";

```

```

        QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
        QName portName = new QName("soap.server.StockQuoteAxis_Wmq",
"soap.server.StockQuoteAxis_Wmq");

        Service service = Service.create(serviceName);
        service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

        /** Create a Dispatch instance from a service **/
        System.out.println("Creating dispatch instance for service StockQuoteAxisService");
        Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
            Service.Mode.MESSAGE);
        System.out.println(" > Dispatch instance created.");

        /*******
         * Create OneWay SOAPMessage request.
         *****/
        MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

        System.out.println("\nCreating a OneWay SOAP Message");
        SOAPMessage request = mf.createMessage();

        /** Obtain the SOAPEnvelope and header and body elements **/
        SOAPPart part = request.getSOAPPart();
        SOAPEnvelope env = part.getEnvelope();
        SOAPHeader header = env.getHeader();
        SOAPBody body = env.getBody();

        /** Construct the message payload **/
        SOAPElement operation = body.addChildElement("getQuoteOneWay", "ns1",
            "soap.server.StockQuoteAxis_Wmq");
        SOAPElement value = operation.addChildElement("in0");
        value.addAttribute(new QName("https://www.w3.org/2001/XMLSchema-instance", "type"),
"string");
        value.addTextNode("XXX");
        request.saveChanges();
        System.out.println(" > SOAP Message created.");

        /** Invoke the service endpoint **/
        System.out.println("Invoking getQuoteOneWay OneWay operation synchronously...");
        dispatch.invokeOneWay(request);
        System.out.println(" > getQuoteOneWay call has returned");

        /*******
         * Create Request Reply SOAPMessage request.
         *****/
        mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

        System.out.println("\nCreating a Request Reply SOAP Message");
        request = mf.createMessage();

        /** Obtain the SOAPEnvelope and header and body elements **/
        part = request.getSOAPPart();
        env = part.getEnvelope();
        header = env.getHeader();
        body = env.getBody();

        /** Construct the message payload **/
        operation = body.addChildElement("getQuote", "ns1", "soap.server.StockQuoteAxis_Wmq");
        value = operation.addChildElement("in0");
        value.addAttribute(new QName("https://www.w3.org/2001/XMLSchema-instance", "type"),
"string");
        value.addTextNode("XXX");
        request.saveChanges();
        System.out.println(" > SOAP Message created.");

        /** Invoke the service endpoint **/
        System.out.println("Invoking getQuote Request Reply operation synchronously...");
        SOAPMessage ans = dispatch.invoke(request);
        System.out.println(" > getQuote call has returned");

        /** Retrieve the result **/
        part = ans.getSOAPPart();
        env = part.getEnvelope();
        body = env.getBody();

        /** Define name of the SOAP folders we are interested in **/
        QName responseName = new QName("soap.server.StockQuoteAxis_Wmq", "getQuoteResponse");
        QName resultName = new QName("getQuoteReturn");

        /** Retrieve result from SOAP envelope **/
        System.out.println("Parsing SOAP response...");
        SOAPElement bodyElement = (SOAPElement) body.getChildElements(responseName).next();

```

```

        SOAPElement responseElement = (SOAPElement)
bodyElement.getChildElements(resultName).next();
        String message = responseElement.getValue();
        System.out.println(" > Response contains result of " + message);

        System.out.println("End of sample");
    }
    catch (Exception fault) {
        // Identify the cause of the Axis Fault
        System.err.println(fault.toString());
        Throwable e = fault.getCause();
        for (int i = 1; e != null; i++) {
            // The toString method on an MQAxisException will cause the message, explanation and
user
            // action.
            System.err.println("Exception(" + i + "): " + e.toString());

            if (e.getCause() != null) {
                e = e.getCause();
            }
            else {
                break;
            }
        } // end of for loop
    } // end of catch block
}
}
}

```

Figure 188. DispatchClientAsyncPolling.java

```

package com.ibm.mq.axis2.samples;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;
import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Response;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientAsyncPolling {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DispatchClientAsyncPolling");

            String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
                + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
                + "&initialContextFactory=com.ibm.mq.jms.Nojndi&targetService=soap.server.StockQuoteAxis.java";

            QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
            QName portName = new QName("soap.server.StockQuoteAxis_Wmq",
"soap.server.StockQuoteAxis_Wmq");

            Service service = Service.create(serviceName);
            service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

            /** Create a Dispatch instance from a service. */
            System.out.println("Creating dispatch instance for service StockQuoteAxisService");
            Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
                Service.Mode.MESSAGE);
            System.out.println(" > Dispatch instance created.");

            /** Create SOAPMessage request. */
            MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

            System.out.println("Creating a Request Reply SOAP Message");
            SOAPMessage request = mf.createMessage();

            /** Obtain the SOAPEnvelope and header and body elements */

```

```

        SOAPPart part = request.getSOAPPart();
        SOAPEnvelope env = part.getEnvelope();
        SOAPHeader header = env.getHeader();
        SOAPBody body = env.getBody();

        /** Construct the message payload */
        SOAPElement operation = body.addChildElement("getQuote", "ns1",
            "soap.server.StockQuoteAxis_Wmq");
        SOAPElement value = operation.addChildElement("in0");
        value.addAttribute(new QName("https://www.w3.org/2001/XMLSchema-instance", "type"),
"string");
        value.addTextNode("XXX");
        request.saveChanges();
        System.out.println(" > SOAP Message created.");

        /** Invoke the service endpoint */
        System.out.println("Invoking getQuote Request Reply operation asynchronously by
polling...");
        Response<SOAPMessage> response = dispatch.invokeAsync(request);
        System.out.println(" > getQuote call has returned");

        /** Sleep main thread until response arrives */
        System.out.println("Waiting for response to arrive...");
        while (!response.isDone()) {
            Thread.sleep(100);
        }
        System.out.println(" > Response received");

        /** retrieve the result */
        SOAPMessage ans = response.get();
        part = ans.getSOAPPart();
        env = part.getEnvelope();
        body = env.getBody();

        /** Define name of the SOAP folders we are interested in */
        QName responseName = new QName("soap.server.StockQuoteAxis_Wmq", "getQuoteResponse");
        QName resultName = new QName("getQuoteReturn");

        /** Retrieve result from SOAP envelope */
        SOAPElement bodyElement = (SOAPElement) body.getChildElements(responseName).next();
        SOAPElement responseElement = (SOAPElement)
bodyElement.getChildElements(resultName).next();
        String message = responseElement.getValue();
        System.out.println(" > Response contains result of " + message);

        System.out.println("End of sample");
    }
    catch (Exception fault) {
        // Identify the cause of the Axis Fault
        System.err.println(fault.toString());
        Throwable e = fault.getCause();
        for (int i = 1; e != null; i++) {
            // The toString method on an MQAxisException will cause the message, explanation and
user
            // action.
            System.err.println("Exception(" + i + "): " + e.toString());

            if (e.getCause() != null) {
                e = e.getCause();
            }
            else {
                break;
            }
        } // end of for loop
    } // end of catch block
}
}
}

```

Figure 189. *DispatchClientAsyncCallback.java*

```

package com.ibm.mq.axis2.samples;

import java.util.concurrent.Future;

import javax.xml.namespace.QName;
import javax.xml.soap.MessageFactory;

```

```

import javax.xml.soap.SOAPBody;
import javax.xml.soap.SOAPConstants;
import javax.xml.soap.SOAPElement;
import javax.xml.soap.SOAPEnvelope;
import javax.xml.soap.SOAPHeader;
import javax.xml.soap.SOAPMessage;
import javax.xml.soap.SOAPPart;
import javax.xml.ws.AsyncHandler;
import javax.xml.ws.Dispatch;
import javax.xml.ws.Response;
import javax.xml.ws.Service;
import javax.xml.ws.soap.SOAPBinding;

public class DispatchClientAsyncCallback implements AsyncHandler<SOAPMessage> {

    public static void main(String[] args) {
        try {
            System.out.println("Starting sample DispatchClientAsyncCallback");

            String endpointUrl = "jms:/queue?destination=SOAPJ.demos@WMQSOAP.DEMO.QM&"
                + "connectionFactory=connectQueueManager(WMQSOAP.DEMO.QM)"
                +
            "&initialContextFactory=com.ibm.mq.jms.NoJndi&targetService=soap.server.StockQuoteAxis.java";

            QName serviceName = new QName("soap.server.StockQuoteAxis_Wmq", "StockQuoteAxisService");
            QName portName = new QName("soap.server.StockQuoteAxis_Wmq",
            "soap.server.StockQuoteAxis_Wmq");

            Service service = Service.create(serviceName);
            service.addPort(portName, SOAPBinding.SOAP11HTTP_BINDING, endpointUrl);

            /** Create a Dispatch instance from a service. */
            System.out.println("Creating dispatch instance for service StockQuoteAxisService");
            Dispatch<SOAPMessage> dispatch = service.createDispatch(portName, SOAPMessage.class,
            Service.Mode.MESSAGE);
            System.out.println(" > Dispatch instance created.");

            /** Create SOAPMessage request. */
            MessageFactory mf = MessageFactory.newInstance(SOAPConstants.SOAP_1_1_PROTOCOL);

            System.out.println("Creating a Request Reply SOAP Message");
            SOAPMessage request = mf.createMessage();

            /** Obtain the SOAPEnvelope and header and body elements */
            SOAPPart part = request.getSOAPPart();
            SOAPEnvelope env = part.getEnvelope();
            SOAPHeader header = env.getHeader();
            SOAPBody body = env.getBody();

            /** Construct the message payload. */
            SOAPElement operation = body.addChildElement("getQuote", "ns1",
            "soap.server.StockQuoteAxis_Wmq");
            SOAPElement value = operation.addChildElement("in0");
            value.addAttribute(new QName("https://www.w3.org/2001/XMLSchema-instance", "type"),
            "string");
            value.addTextNode("XXX");
            request.saveChanges();
            System.out.println(" > SOAP Message created.");

            /** Invoke the service endpoint. */
            DispatchClientAsyncCallback handler = new DispatchClientAsyncCallback();

            System.out
            .println("Invoking getQuote Request Reply operation asynchronously using a
            callback...");
            Future<?> monitor = dispatch.invokeAsync(request, handler);
            System.out.println(" > getQuote call has returned");

            /** Sleep main thread until handler has been notified */
            System.out.println("Waiting for handler to be called...");
            while (!monitor.isDone()) {
                Thread.sleep(100);
            }

            System.out.println("End of sample");
        }
        catch (Exception fault) {
            // Identify the cause of the Axis Fault
            System.err.println(fault.toString());
            Throwable e = fault.getCause();
            for (int i = 1; e != null; i++) {
                // The toString method on an MQAxisException will cause the message, explanation and

```

```

user
    // action.
    System.err.println("Exception(" + i + "): " + e.toString());

    if (e.getCause() != null) {
        e = e.getCause();
    }
    else {
        break;
    }
} // end of for loop
} // end of catch block
}

public void handleResponse(Response<SOAPMessage> response) {
    try {
        // retrieve the result
        SOAPMessage ans = response.get();
        SOAPPart part = ans.getSOAPPart();
        SOAPEnvelope env = part.getEnvelope();
        SOAPBody body = env.getBody();

        /** Define name of the SOAP folders we are interested in */
        QName responseName = new QName("soap.server.StockQuoteAxis_Wmq", "getQuoteResponse");
        QName resultName = new QName("getQuoteReturn");

        /** Retrieve result from SOAP envelope */
        SOAPElement bodyElement = (SOAPElement) body.getChildElements(responseName).next();
        SOAPElement responseElement = (SOAPElement)
bodyElement.getChildElements(resultName).next();
        String result = responseElement.getValue();

        System.out.println(" > Async Handler has received a result of " + result);
    }
    catch (Exception fault) {
        // Identify the cause of the Axis Fault
        System.err.println("Exception in handleResponse");
        System.err.println(fault.toString());
        Throwable e = fault.getCause();
        for (int i = 1; e != null; i++) {
            // The toString method on an MQAxisException will cause the message, explanation and
user
            // action.
            System.err.println("Exception(" + i + "): " + e.toString());

            if (e.getCause() != null) {
                e = e.getCause();
            }
            else {
                break;
            }
        } // end of for loop
    } // end of catch block
}
}

```

Related tasks

[Developing a JAX-RPC client for WebSphere transport for SOAP using Eclipse](#)

[Develop an Axis 1.4 Web service client to run using the WebSphere MQ transport for SOAP.](#)

[Developing a .NET 1 or 2 client for WebSphere transport for SOAP using Microsoft Visual Studio 2008](#)

[Develop an .NET 1 or 2 Web service client to run using the WebSphere MQ transport for SOAP.](#)

Developing a .NET 1 or 2 client for WebSphere transport for SOAP using Microsoft Visual Studio 2008

Develop an .NET 1 or 2 Web service client to run using the WebSphere MQ transport for SOAP.

Before you begin

You can start the development of a .NET 1 or 2 client in a number of different ways:

1. Use **amqwdeployMQService** to generate client stubs from a Web service and import them into Visual Studio.

2. Use **java2wsdl** to generate WSDL from a Java implementation of a Web service, and then use `wsdl.exe`, which is shipped with .NET, to generate client stubs.
3. Generate WSDL from a .NET .asmx implementation of the service using **amqswsdl**, and then use `wsdl.exe`.
4. If you have developed and deployed the service for HTTP, use the **Add Web reference...** wizard in Visual studio to configure the client to access the HTTP service. Alter the URL refer to the service deployed to WebSphere MQ transport for SOAP.

The task uses the service developed in “Developing a .NET 1 or 2 service for WebSphere MQ transport for SOAP using Microsoft Visual Studio 2008” on page 923.

About this task

Follow these steps to create a .NET 1 or 2 Client for HTTP and WebSphere MQ transport for SOAP.

Procedure

1. Create the client console application and modify it to invoke the StockQuote HTTP Web service.
 - a) Right click **Solution 'StockQuoteDotNet'** in the **Solution Explorer** > Add... > New Project. Select the **C#** Project Type, **.NET Framework 2.0**, and **Console Application**. Name the project **StockQuoteClientDotNet** > **OK**
 - b) Right click **Solution 'StockQuoteDotNet'** in the **Solution Explorer** > Add... > New Project. Select the **C#** Project Type, **.NET Framework 2.0**, and **Console Application**. Name the project **StockQuoteClientDotNet** > **OK**
 - c) Right-click **StockQuoteClientDotNet** > **Set as Startup project**.
 - d) Right-click **StockQuoteClientDotNet** > **Add Web Reference...** > Browse to Web services in this solution > Select **StockQuote** > **Add Reference**. Notice you have added a Web reference to local host and a new configuration file `app.config`.
 - e) In the Solution Explorer, change the name of the console application from `Program.cs` to `StockQuoteClientDotNet.cs` > Click **OK** to changing all the usages of `Program.cs` to `StockQuoteClientDotNet.cs`.
 - f) Replace the contents of `StockQuoteClientDotNet.cs` with the code in [Figure 190 on page 952](#).

```
using System;
using StockQuoteClientDotNet.localhost;
namespace StockQuoteClientDotNet {
    class StockQuoteClientDotNet {
        static void Main(string[] args) {
            try {
                StockQuote stockobj = new StockQuote();
                Console.WriteLine("http reply is: "
                    + stockobj.getNonInlineQuote("http request"));
            }
            catch (System.Exception e) {
                Console.WriteLine("Exception thrown: " + e.ToString());
            }
            Console.ReadLine();
        }
    }
}
```

Figure 190. HTTP StockQuoteClientDotNet program

- g) Launch `StockQuoteClientDotNet` to test against the `StockQuote.asmx` service:
 - i) Press **F5**, click the green arrow in the action bar, or **Debug > Start Debugging (F5)**.
If the `StockQuoteDotNet` project is in the same solution, it starts automatically. Otherwise you need to start the service first.

The command window with the results opens behind the workspace. The `Console.ReadLine();` statement prevents it from closing until you press **Enter**.

Tip: Make sure `StockQuote.asmx` is the Start page in the `StockQuoteDotNet` project.

2. Modify `StockQuoteClientDotNet` to call the `StockQuote.asmx` service using WebSphere MQ transport for SOAP.

- a) Add the lines shown in bold to the client.

```
using System;
using StockQuoteClientDotNet.localhost;
namespace StockQuoteClientDotNet {
    class StockQuoteClientDotNet {
        static void Main(string[] args) {
            try {
                IBM.WMQSOAP.Register.Extension();
                StockQuote stockobj = new StockQuote();
                Console.WriteLine("http reply is: "
                    + stockobj.getNonInlineQuote("http request"));
                stockobj.Url = "jms:/queue?"
                + "initialContextFactory=com.ibm.mq.jms.Nojndi"
                + "&connectionFactory=()&destination=REQUESTDOTNET@QM1"
                + "&targetService=StockQuote.asmx";
                Console.WriteLine("jms reply is: "
                    + stockobj.getNonInlineQuote("jms request"));
            }
            catch (System.Exception e) {
                Console.WriteLine("Exception thrown: " + e.ToString());
            }
            Console.ReadLine();
        }
    }
}
```

Figure 191. Modified `StockQuoteClientDotNet` program

Alternatively, modify the default URL. Open **StockQuoteClientDotNet** > **Properties** > **Settings.settings** and change the value of the `StockQuoteClientDotNet_localhost_StockQuote` property to the WebSphere MQ transport for SOAP URL.

- b) Add a reference to `amqsoap.dll`
 - i) In the **StockQuoteClientDotNet** project in the **Solution Explorer**, right-click **References** > **Add Reference...** > Click the **Browse** tab > browse to `MQ_INSTALLATION_PATH\bin` > Select **amqsoap.dll** > **OK**. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.
3. Test the client with the `StockQuote.asmx` service using WebSphere MQ transport for SOAP.
 - a) Open a command window in the `StockQuoteDotNet` project directory: `.\StockQuoteDotNet\StockQuoteDotNet > Verify` that `.bin\StockQuoteDotNet.dll` exists. If it does not, rebuild the solution.
 - b) Type the command **amqwRegisterdotNet**.
You need only run **amqwRegisterdotNet** once per installation.
 - c) If you have run **amqwdeployWMQServer** with the `genAsmxWMQBits`, run the .NET SOAP Listener: `generated\server\startWMQNListener`
 - d) Alternatively run the listener directly:

```
amqwSOAPNETListener -u "jms:/queue?
destination=REQUESTDOTNET@QM1
&connectionFactory=()&initialContextFactory=com.ibm.mq.jms.Nojndi
&targetService=StockQuote.asmx&replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE"
-w C:\IBM\ID\StockQuoteDotNet\StockQuoteDotNet -n 10
```

4. In Visual Studio 2008, press **F5** to run `StockQuoteClientDotNet`.

.NET Framework 1 and .NET Framework 2 Web service clients

The sample .NET clients provided with the WebSphere MQ transport for SOAP use generated stubs to call the sample Axis and .NET services.

For .NET Framework 1 and .NET Framework 2 clients, WebSphere MQ provides access to web services using .NET clients. The **amqwdeloyWMQService** command has an option, **genProxiestoDotNet**, that generates .NET Framework 1 or .NET Framework 2 client stubs for a Web service. You can also use client stubs generated by the .NET **wsdl** tool, or by Microsoft Visual Studio 2005, or 2008.

The sample .NET Framework 1 and .NET Web service clients are installed in `MQ_INSTALLATION_PATH\tools\soap\samples\dotnet`. `MQ_INSTALLATION_PATH` is the directory where WebSphere MQ is installed.

SQVB2Axis.vb

`SQVB2Axis.vb`, [Figure 192 on page 954](#), is the Visual Basic client to call the **StockQuoteAxisService** service.

SQVB2DotNet.vb

`QVB2DotNet.vb`, [Figure 193 on page 954](#), is the Visual Basic client to call the **StockQuoteDotNet** service.

SQCS2Axis.cs

`SQCS2Axis.cs`, [Figure 194 on page 955](#), is the C# client to call the **StockQuoteAxisService** service. You can override the URL of the service by providing a URL on the command line.

SQCS2DotNet.cs

`SQCS2DotNet.cs`, [Figure 195 on page 955](#), is the C# client to call the **StockQuoteDotNet** service. You can override the URL of the service by providing a URL on the command line.

```
Module SQVB2Axis
    Function Main(ByVal CmdArgs() As String) As Integer
        IBM.WMQSOAP.Register.Extension()
        Dim obj As New StockQuoteAxisService()
        Dim res As Single = obj.getQuote("fromcs")
        System.Console.WriteLine("SQVB2Axis: reply is: '{0}'", res)
    End Function
End Module
```

Figure 192. SQVB2Axis

```
Module SQVB2DotNet
    Function Main(ByVal CmdArgs() As String) As Integer
        IBM.WMQSOAP.Register.Extension()
        Dim obj as new StockQuoteDotNet()
        Dim res as Single = obj.getQuote("fromcs")
        System.Console.WriteLine("SQVB2DotNet: reply is: '{0}'", res)
    End Function
End Module
```

Figure 193. SQVB2DotNet

```

using System;
class SQCS2Axis {
    [STAThread]
    static void Main(string[] args) {
        try {
            IBM.WMQSOAP.Register.Extension();
            StockQuoteAxisService stockobj = new StockQuoteAxisService();
            if (args.GetLength(0) >= 1)
                stockobj.Url = args[0];
            System.Single res = stockobj.getQuote("XXX");
            Console.WriteLine("SQCS2Axis RPC reply is: " + res);
        }
        catch (System.Exception e) {
            Console.WriteLine("\n>>> EXCEPTION WHILE RUNNING SQCS2Axis DEMO <<<\n"
                + e.ToString());
        }
    }
}

```

Figure 194. SQCS2Axis

```

using System;
class SQCS2DotNet {
    [STAThread]
    static void Main(string[] args) {
        try {
            IBM.WMQSOAP.Register.Extension();
            StockQuoteDotNet stockobj = new StockQuoteDotNet();
            if (args.GetLength(0) >= 1)
                stockobj.Url = args[0];
            System.Single res = stockobj.getQuote("XXX");
            Console.WriteLine("RPC reply is: " + res);
            if (args.GetLength(0) == 0) {
                res = stockobj.getQuoteDOC("XXX");
                Console.WriteLine("DOC reply is: " + res);
            }
        }
        catch (System.Exception e) {
            Console.WriteLine("\n>>> EXCEPTION WHILE RUNNING SQCS2DotNet DEMO <<<\n"
                + e.ToString());
        }
    }
}

```

Figure 195. SQCS2DotNet

Related tasks

[Developing a JAX-RPC client for WebSphere transport for SOAP using Eclipse](#)

Develop an Axis 1.4 Web service client to run using the WebSphere MQ transport for SOAP.

[Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse](#)

Develop an Axis2 Web service client to run using the WebSphere MQ transport for SOAP. The sample Axis2 clients provided with WebSphere MQ transport for SOAP are listed, and the **wsimport** command used to generate proxies.

Deploying Web services using the WebSphere MQ transport for SOAP

Deploy a Web service to one of a number of different server environments and connect to it using WebSphere MQ transport for SOAP.

Before you begin

Develop a Web service and test it using SOAP over HTTP in the target environment.

About this task

You can deploy a web service to run with WebSphere MQ transport for SOAP in a number of different SOAP run time environments. You can deploy a service to Axis 1.4 using only the software installed with WebSphere MQ. For the other run time environments, you must install additional software.

You are not restricted to running WebSphere MQ transport for SOAP to the servers for which there are deployment instructions. Use the instructions to deploy a service to one of the listed environments.

Note: Some integrated environments offer SOAP over JMS using the W3C recommended JMS SOAP binding, as well as the WebSphere MQ transport for SOAP binding. Releases of WebSphere MQ, up to and including 7.0.1.2, support only the WebSphere MQ transport for SOAP binding. From 7.0.1.3 onwards you can deploy Axis2 clients using a URI that conforms to the W3C candidate recommendation for SOAP over JMS. See the tutorial, [Develop a SOAP/JMS JAX-WS Web services application with WebSphere Application Server V7 and Rational Application Developer V7.5](#).

Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using `amqwdeployWMQService`

Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **`amqwdeployWMQService`** command, and starting the Axis 1.4 listener.

Before you begin

1. Follow the instructions for installing WebSphere MQ transport for SOAP
2. Verify the installation and your environment using the **`runivt`** command.
3. To redeploy a service:
 - a. Delete the `./generated` subdirectory, and all its subdirectories.
 - b. Remove requests from the destination queue and delete it.
 - c. Proceed with the instructions from step “2” on [page 956](#).

About this task

These instructions are to deploy an Axis 1.4 service for the first time. To restart an Axis 1.4 service, rerun the Axis 1.4 SOAP listener: step “11” on [page 957](#).

Use the following instructions to deploy a new Axis 1.4 service to WebSphere MQ transport for SOAP:

Procedure

1. Create a directory `deployDir` to hold the deployment files.

The deployment utility requires that each service is deployed from a separate directory.
2. Open a command window on Windows, or a command shell using X Window System on UNIX and Linux systems, in `deployDir` to run **`amqwdeployWMQService`**.
3. Run **`amqwsetcp`** to set the classpath.

The JRE and JDK must be in the classpath, at version 5.0 or later, and at the same version level.
4. Copy the class source, `className.java`, into `deployDir`
5. Copy all the Java source files in the same package as `className` into `deployDir/packageName`, where `packageName` is a directory tree corresponding to the package name.
6. Run **`javac packageName.className`**.

You might need to add a path to the current directory `."`, or to the `packageName` directory for **`javac`** to find the other classes.
7. Create the Axis WSDL for the service:

```
amqwdeployWMQService -f packageName.className.java -c genAxisWsd1  
-v -u "jms:/queue?destination=queueName"
```

```
&initialContextFactory=com.ibm.mq.jms.NoJndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

8. Create the WebSphere MQ resources for the service:

```
amqwdeployWMQService -f packageName.className.java -c genAxisWMQBits
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.NoJndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

Tip:

If you want to set up a new queue manager, and the resources it needs, to do development and testing, run **setupWMQSOAP**.

If you want set up the new queue manager as the default, take a copy of **setupWMQSOAP** from the *WMQ install directory\tools\soap\samples* directory, and add the -q parameter to the line

```
call :try -q crtmqm %QMGR%
```

9. Create the Axis listener and deploy the service:

```
amqwdeployWMQService -f packageName.className.java -c AxisDeploy
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.NoJndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))"
```

10. If you need to generate the WSDL for the service, generate client stubs, or client proxies, run **amqwdeployWMQService** with one of the following parameters:

- genAsmxWsd1
- genAxisWsd1
- genProxiesToDotNet
- genProxiestoAxis

Note: You must generate WSDL before generating the proxies. The AllAxis option fails if the CLASSPATH is not set up to find all the classes that are imported to compile *className.java*. If there are multiple Java files in the package containing *className.java*, you must compile them first using **javac**. **amqwdeployWMQService** -f *packageName.className.java* -c CompileJava compiles only *className.java*.

11. Start the generated Axis listener.

```
.\generated\server\startWMQJListener.cmd
```

Related tasks

[Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP](#)

Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

[Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

[Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

[Configuring WebSphere Application Server to use W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

[Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP

Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

Before you begin

1. Follow the instructions for installing WebSphere MQ transport for SOAP
2. Verify the installation and your environment using the **runivt** command.
3. The path to the .NET framework files `wsdl.exe` and `csc.exe` must be set. The copies of `wsdl.exe` and `csc.exe` identified by the `PATH` variable must be at the same level of the .NET framework. If you have multiple .NET frameworks installed, or are using Visual Studio, check the `PATH` variable carefully.
4. To redeploy a service:
 - a. Delete the `./generated` subdirectory, and all its subdirectories
 - b. Remove requests from the destination queue and delete it.
 - c. Proceed with the instructions from step “2” on page 958.

About this task

These instructions are to deploy a .NET service for the first time. To restart a .NET service, rerun the .NET SOAP listener, step “9” on page 959.

Use the following instructions to deploy a new .NET Framework 1 or .NET Framework 2 service to WebSphere MQ transport for SOAP:

Procedure

1. Create a directory *deployDir* to hold the deployment files.
The deployment utility requires that each service is deployed from a separate directory.
2. Open a command window in *deployDir* to run **amqwdeployWMQService**.

```
C:\IBM\ID\QuoteClient>
```

3. Run **amqwsetcp** to set the classpath.
A classpath is needed only for Axis clients.
4. Copy the .NET service, *className.asmx*, into *deployDir*
5. Build the service implementation into a library (.dll).

The inline service implementation is in *className.asmx*. The code-behind service implementation might be *className.asmx.cs*.

[Figure 196 on page 959](#) shows an example of a command to build a .NET Framework V2 service as a library.

```
c:\WINDOWS\Microsoft.NET\Framework\v3.5\Csc.exe /noconfig /nowarn:1701,1702
/errorreport:prompt /warn:4 /define:TRACE
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.configuration.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Data.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Drawing.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Web.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Web.Services.dll
/reference:c:\WINDOWS\Microsoft.NET\Framework\v2.0.50727\System.Xml.dll
/debug:pdbonly /filealign:512 /optimize+
/out:obj\Quote.dll /target:library Properties\AssemblyInfo.cs Quote.asmx.cs
```

Figure 196. Build command for .NET Framework V2 service

6. Copy *className.dll* into *deployDir\bin*.
7. Set up the WebSphere MQ resources, and create the listener required for the service:

```
amqwdeployWMQService -f className.asmx -c genAsmxWMQBits
-v -u "jms:/queue?destination=queueName
&initialContextFactory=com.ibm.mq.jms.Nojndi
&connectionFactory=(connectQueueManager(QmgrName)binding(auto))
&targetService=className.asmx"
```

8. If you need to generate the WSDL for the service, generate client stubs, or client proxies, run **amqwdeployWMQService** with one of the following parameters:

- genAsmxWsd1
- genAxisWsd1
- genProxiesToDotNet
- genProxiestoAxis

Note: You must generate WSDL before generating the proxies.

9. Start the generated .NET listener.

```
.\generated\server\startWMQListener.cmd
```

Related tasks

[Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService](#)
 Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **amqwdeployWMQService** command, and starting the Axis 1.4 listener.

[Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP](#)
 WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

[Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP](#)
 WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

[Configuring WebSphere Application Server to use W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

[Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP

WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

Before you begin

Use the same tools to develop for a client or service for WebSphere MQ, as you would to develop for HTTP. CICS has tools corresponding to **Java2wsdl** and **wsdl2Java**:

- **DFHWS2LS** takes a Web service description as a starting point. It uses the descriptions of the messages, and the data types used in those messages, to construct high-level language data structures. You can use in the structures in application programs written in different languages.
- **DFHLS2WS** takes a high-level language data structure as a starting point. It uses the structure to construct a Web services description that contains descriptions of messages. It also creates schemas for the messages from the language data structure.

Follow the instructions, [Creating a Web service](#) in the CICS product documentation, to create a Web service.

About this task

Follow the instructions, [Configuring CICS to use the WebSphere MQ transport](#) in the CICS product documentation. Using the instructions, you can deploy the Web service to WebSphere MQ transport for SOAP.

Related tasks

[Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService](#)

Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **amqwdeployWMQService** command, and starting the Axis 1.4 listener.

[Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP](#)

Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

[Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

[Configuring WebSphere Application Server to use W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

[Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP

WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

Before you begin

Use Rational Application Developer, WebSphere Integration Developer, or a Web services toolkit to develop the Web service.

About this task

Use the following instructions to deploy a service using WebSphere MQ transport for SOAP as a SOAP transport on WebSphere Application Server.

Procedure

1. Configure WebSphere MQ as the JMS messaging provider for the service integration bus on WebSphere Application Server.
2. Configure the WebSphere MQ resources required by the service.
3. Follow the instructions, [Configuring JMS resources for the synchronous SOAP over JMS endpoint listener](#), in the WebSphere Application Server Network Deployment product documentation.
There are corresponding instructions for other WebSphere Application Server platforms.
4. Modify the service URI to conform to the WebSphere MQ transport for SOAP URI.
5. Deploy the service to WebSphere Application Server.

What to do next

Deploy the service with HTTP as a transport so that clients can query the service and receive the WSDL in response.

Related tasks

[Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService](#)
Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **amqwdeployWMQService** command, and starting the Axis 1.4 listener.

[Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP](#)
Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

[Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP](#)
WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

[Configuring WebSphere Application Server to use W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

[Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

Configuring WebSphere Application Server to use W3C SOAP over JMS

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

Before you begin

The task requires WebSphere Application Server v7.0.0.9 and WebSphere MQ v7.0.1.3.

About this task

The task has two steps:

Procedure

1. [“Configure WebSphere MQ resources” on page 962](#)
2. [“Configure WebSphere Application Server resources” on page 963](#)

What to do next

[“Configure WebSphere MQ resources” on page 962](#)

Related tasks

[Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService](#)
Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **amqwdeployWMQService** command, and starting the Axis 1.4 listener.

[Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP](#)
Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

[Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP](#)
WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

[Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP](#)
WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

[Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP](#)
WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

Configure WebSphere MQ resources

Before you begin

For Axis2 support you required WebSphere MQ 7.0.1.3 or later.

About this task

For simplicity, the task assumes WebSphere MQ is installed on the same workstation as the other software, and uses bindings connections. The WebSphere Application Server and the Axis2 client configurations work with client connections. To run with task using client connections check that you can put and get messages to and from the request and reply queues from both the Axis2 client and the WebSphere Application Server computers.

Again, for simplicity, no security configuration is used. The user ID has full mqm authority.

Procedure

1. Create a default queue manager, QM1.

Use WebSphere MQ Explorer to create QM1 as a default queue manager. Configure it to start automatically and select the option to create a listener. Alternatively use the following commands:

```
crtmqm -q -sa QM1
strmqm
echo define listener (LISTENER.TCP) trptype(tcp) ipaddr(localhost) port(1414)
         control(qmgr) replace | runmqsc
echo start listener(LISTENER.TCP) | runmqsc
```

2. Define a request queue, REQUESTAXIS, and a reply queue, REPLYAXIS.

Use the Explorer, or the following commands:

```
echo define ql(REQUESTAXIS) replace | runmqsc
echo define ql(REPLYAXIS) replace | runmqsc
```

What to do next

[“Configure WebSphere Application Server resources” on page 963](#)

Configure WebSphere Application Server resources

Before you begin

For W3C SOAP over JMS support you require WebSphere Application Server v7. This configuration was performed on WebSphere Application Server Version 7.0 Test Environment v7.0.0.9 Update 1. WebSphere Application Server was supplied with Rational Software Architect for WebSphere Software 7.5.4. Rational Software Architect was updated to v7.5.5.1, by applying the latest updates that were available.

As part of the installation process, create a profile for WebSphere Application Server. In the task, administrative security is not enabled. The default profile name is was70profile1, and the server is server1.

About this task

Configure the WebSphere Application Server. You can either start the server from Rational Application Developer, and start the administrative console from the Servers view, or you can start the server using a command file and administer the server using a browser. The task uses the second method.

The server command files are in the folder, *Rational Installation Root\SDP\runtimes\base_v7\profiles\was70profile1\bin*. The log files you might want to inspect are in *Rational Installation Root\SDP\runtimes\base_v7\profiles\was70profile1\logs\server1*.

As a convention, all WebSphere MQ object names are uppercase, and all JNDI names referencing the WebSphere MQ objects are lowercase.

Procedure

1. Start the server.

```
startServer server1
```

2. Start a browser, open the administration console, and login.

```
http://localhost:9061/ibm/console/unsecureLogon.jsp
```

Type any string into the user ID field.

3. Create a connection factory, qm1
 - a) In the navigator, open **Resources > JMS > Connection factories**.
 - b) In the Connection factories window, select scope **Node=nodename**, click **New**.
 - c) Select **WebSphere MQ messaging provider > OK**.
 - d) Provide the queue manager connection information from [Table 143 on page 963](#) > **Next**.

Table 143. Queue manager connection information	
Field name	Value
Name	qm1
JNDI name	qm1

- e) Select **Enter all the required information into this wizard** as the connection method > **Next**.
- f) Type QM1 as the queue connection details > **Next**.
- g) Enter the connection details from [Table 144 on page 964](#) > **Next**.

Table 144. Connection details	
Field name	Value
Transport	Bindings, then client
Hostname	localhost
Port	1414
Server connection channel	SYSTEM.DEF.SVRCONN

- h) **Test connection > Next > Finish > Save.**
4. Create the JMS request queue, requestaxis.
- In the Navigator, open **Resources > JMS > Queues.**
 - In the Connection factories window, select scope **Node=nodename**, click **New.**
 - Select **WebSphere MQ messaging provider > OK.**
 - Enter the queue details from [Table 145 on page 964](#) > **OK > Save.**

Table 145. Queue details	
Field name	Value
Name	requestaxis
JNDI name	requestaxis
Queue name	REQUESTAXIS
Queue manager name	QM1

5. Repeat step “4” on page 964 to create the JMS reply queue, replyaxis.
6. Create an activation specification, qm1as.

The activation specification triggers the Web service router Message Driven Bean (MDB) when a message arrives on the request queue. The MDB is defined in the deployment descriptor of the Web service that is created by the Rational Application Developer Web service wizard.

- In the Navigator, open **Resources > JMS > Activation specifications.**
- In the Connection factories window, select scope **Node=nodename**, click **New.**
- Select **WebSphere MQ messaging provider > OK.**
- Enter the basic attributes of the activation specification from [Table 146 on page 964](#) > **Next.**

Table 146. Activation specification name	
Field name	Value
Name	qm1as
JNDI name	qm1as

- e) Specify its MDB information from [Table 147 on page 964](#) > **Next.**

Table 147. MDB information	
Field name	Value
Destination JNDI name	requestaxis
Message selector	Left blank
Destination type	Queue

- Select **Enter all the required information into this wizard** as the connection method > **Next.**
- Type QM1 as the queue connection details > **Next.**

- h) Enter the connection details from [Table 144 on page 964](#) > **Next**.

Table 148. Connection details	
Field name	Value
Transport	Bindings, then client
Hostname	localhost
Port	1414
Server connection channel	SYSTEM.DEF.SVRCONN

- i) **Test connection** > **Next** > **Finish** > **Save**.

7. Create a queue connection factory, `jms/WebServicesReplyQCF`, for the reply queue.

The Web services router uses a queue connection factory to access a reply queue. In the deployment descriptor of the Web service the queue connection factory is given the default JNDI name of `jms/WebServicesReplyQCF`. You can change the name in the deployment descriptor. In this task, add the default name to the JMS resource definitions.

- a) In the Navigator, open **Resources** > **JMS** > **Queue connection factories**.
b) In the Connection factories window, select scope **Node=nodename**, click **New**.
c) Select **WebSphere MQ messaging provider** > **OK**.
d) Enter the basic attributes of the queue connection factory from [Table 149 on page 965](#) > **Next**.

Table 149. Queue connection factory name	
Field name	Value
Name	WebServicesReplyQCF
JNDI name	jms/WebServicesReplyQCF

- e) Select **Enter all the required information into this wizard** as the connection method > **Next**.
f) Type QM1 as the queue connection details > **Next**.
g) Enter the connection details from [Table 144 on page 964](#) > **Next**.

Table 150. Connection details	
Field name	Value
Transport	Bindings, then client
Hostname	localhost
Port	1414
Server connection channel	SYSTEM.DEF.SVRCONN

- h) **Test connection** > **Next** > **Finish** > **Save**.

What to do next

[“Developing a JAX-WS EJB Web service for W3C SOAP over JMS” on page 925](#)

Deploying a service to WebSphere ESB and Process Server service endpoint to use WebSphere Transport for SOAP

WebSphere MQ transport for SOAP is not directly supported by WebSphere ESB and Process Server. You must configure a custom Export.

About this task

WebSphere Integration Developer provides a SOAP data transformation that you can bind to the WebSphere MQ JMS Export to create a custom WebSphere MQ JMS SOAP Export.

Follow the instructions to create a customized Export to receive SOAP requests over WebSphere MQ transport for SOAP.

Procedure

1. Read [Overview of imports and exports](#) and [How to connect to WebSphere MQ](#) in the WebSphere Process Server for Multiplatforms V6.2 product documentation.
2. Follow the task, [Generating an MQ JMS export binding](#) in the IBM Business Process Manager, Version 8.6 product documentation.

Use the SOAP data-binding described in [Prepackaged JMS data format transformations](#) to format the SOAP message.

Related tasks

[Deploying a service to Axis 1.4 to use for WebSphere transport for SOAP using amqwdeployWMQService](#)
Deploy an Axis 1.4 service to WebSphere MQ transport for SOAP by creating a deployment directory, running the **amqwdeployWMQService** command, and starting the Axis 1.4 listener.

[Deploying a service to .NET Framework 1 or 2 service to use WebSphere MQ transport for SOAP](#)

Deploy a .NET Framework 1 or 2 service to WebSphere MQ transport for SOAP. Create a deployment directory, run the **amqwdeployWMQService** command, and start the .NET listener.

[Deploying a service to CICS Transaction Server to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is integrated into CICS Transaction Server 4.1 Web services support.

[Deploying a service to WebSphere Application Server to use WebSphere Transport for SOAP](#)

WebSphere MQ transport for SOAP is integrated into the service integration bus on WebSphere Application Server.

[Configuring WebSphere Application Server to use W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 1 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application server using the W3C SOAP over JMS protocol. Configure the WebSphere MQ and WebSphere Application Server resources to develop and deploy Web service bound to W3C SOAP over JMS as a transport.

Deploying Web service clients to use WebSphere MQ transport for SOAP

Deploy a Web service client to one of a number of different client environments and connect to a service using WebSphere MQ transport for SOAP.

Before you begin

Develop the Web service and deploy it to use WebSphere MQ transport for SOAP.

About this task

You can deploy a web service client to run with WebSphere MQ transport for SOAP in a number of different client environments. You can deploy a Java client to Axis 1.4 using only the software installed with WebSphere MQ. For the other client environments, you must install additional software.

You are not restricted to running WebSphere transport for SOAP in the client environments for which there are deployment instructions. Use the instructions to deploy a client to one of the supported environments.

Note: Some integrated environments offer SOAP over JMS using the W3C recommended JMS SOAP binding, as well as the WebSphere MQ transport for SOAP binding. Releases of WebSphere MQ, up to and including 7.0.1.2, support only the WebSphere MQ transport for SOAP binding. From 7.0.1.3 onwards you can deploy Axis2 clients using a URI that conforms to the W3C candidate recommendation for SOAP over JMS. See the tutorial, [Develop a SOAP/JMS JAX-WS Web services application with WebSphere Application Server V7 and Rational Application Developer V7.5](#).

Deploying a Web service client to Axis 1.4 to use IBM WebSphere MQ transport for SOAP

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM WebSphere MQ queues and channels, start the service and test the client.

Before you begin

Tip: Deploy the service to HTTP, develop and test the client for HTTP, and then modify the client for IBM WebSphere MQ transport for SOAP:

1. Add the `Register.extension()` call to the client.
2. Change the static web service address in the client proxy locator class to use the URI for the IBM WebSphere MQ transport for SOAP.

About this task

Deploying an Axis 1.4 client to use IBM WebSphere MQ transport for SOAP requires one additional deployment step compared to an HTTP client. You must create a client deployment descriptor, `client-config.wsdd`, to map the `jms:transport` to the sender class `com.ibm.mq.soap.transport.jms.WMQSender`.

If you use the command **`amqwdeployMQService`** to generate client proxies, you can deploy the client using the directories the command generates.

Procedure

1. Create a directory *deployDir* to hold the client deployment files.
2. Open a command window on Windows systems, or a command shell using X Window System on UNIX and Linux systems, in *deployDir*.
3. Run the **`amqwsetcp.cmd`** command to set the CLASSPATH
4. Run the **`amqwclientconfig.cmd`** command to create an Axis 1.4 client deployment descriptor, `client-config.wsdd` in *deployDir*.
5. Make sure the classes in the client package, the client proxy classes, and the libraries the client uses, are in the CLASSPATH.

`amqwdeployMQService` places the .NET client proxies into `./generated/server/soap/client/remote/dotnetService` and the Axis 1.4 proxies into `./generated/server/soap/client/remote/client package`.

Example

The example shows the configuration and output, [Figure 199 on page 968](#), from an Axis 1.4 Java client. The client, [Figure 198 on page 968](#), calls a Web service that echoes its input parameter. The service definition, [Figure 197 on page 968](#), shows the URI taken from the service WSDL.

```

<wsdl:service name="QuoteSOAPImplService">
  wsdl:port binding="intf:org.example.www.QuoteSOAPImplBindingSoap"
            name="org.example.www.QuoteSOAPImpl_Wmq">
    <wsdlsoap:address location="jms:/queue?destination=REQUESTAXIS
      &connectionFactory=(connectQueueManager(QM1)binding(server))
      &initialContextFactory=com.ibm.mq.jms.NoJndi
      &targetService=org.example.www.QuoteSOAPImpl.java
      &replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE" />
  </wsdl:port>
</wsdl:service>

```

Figure 197. Service definition

```

package org.example.www;
import com.ibm.mq.soap.Register;
public class QuoteClient {
    public static void main(String[] args) {
        try {
            Register.extension();
            QuoteSOAPImplServiceLocator locator = new QuoteSOAPImplServiceLocator();
            System.out.println("Response = "
                + locator.getOrgExampleWwwQuoteSOAPImpl_Wmq().getQuote("IBM"));
        } catch (Exception e) {
            System.out.println("Exception = " + e.getMessage());
        }
    }
}

```

Figure 198. Axis 1.4 Java client

```

C:\IBM\ID\Test>dir /s /b
C:\IBM\ID\Test\client-config.wsdd
C:\IBM\ID\Test\org
C:\IBM\ID\Test\org\example
C:\IBM\ID\Test\org\example\www
C:\IBM\ID\Test\org\example\www\GetQuoteFaultMsg.class
C:\IBM\ID\Test\org\example\www\OrgExampleWwwQuoteSOAPImplBindingSoapStub.class
C:\IBM\ID\Test\org\example\www\QuoteClient.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImpl.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImplService.class
C:\IBM\ID\Test\org\example\www\QuoteSOAPImplServiceLocator.class

C:\IBM\ID\Test>amqwssetcp
C:\IBM\ID\Test>java org.example.www.QuoteClient.class
Response = IBM

```

Figure 199. Client configuration and output

What to do next

1. If you are deploying the client as an IBM WebSphere MQ client, configure the client and server connection channel.
2. If you are deploying the client to a different queue manager to the service, you must make the destination queue available to the client. Configure the destination queue on the service queue manager as a cluster queue, or on the client queue manager as a remote queue definition.

Related tasks

[Deploying a Web service client to Axis2 to use WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and Axis2 configuration file for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure WebSphere MQ queues and channels, start the service and test the client.

[Deploying to an Axis2 client using W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 4 of connecting an Axis2 Web service client and

a Web service deployed to WebSphere Application Server using the W3C SOAP over JMS protocol. Modify the URL in the Axis2 client developed for WebSphere MQ transport for SOAP to use the W3C candidate recommendation for SOAP over JMS.

Deploying a Web service client to .NET Framework 1 and 2 to use WebSphere MQ transport for SOAP
Prepare a deployment directory and deployment descriptor for the client. Provide the client proxy and client class. Configure WebSphere MQ queues and channels, start the service and test the client.

Deploying a Web service client to Axis2 to use WebSphere MQ transport for SOAP

Prepare a deployment directory and Axis2 configuration file for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure WebSphere MQ queues and channels, start the service and test the client.

Before you begin

Tip: Deploy the service to HTTP. Develop and test the client for HTTP, and then modify the URL to reference the service using WebSphere MQ transport for SOAP.

The task shows how to deploy an unmanaged Axis2 client to Java Standard Edition. You might want to deploy an Axis2 client to a Web container. In [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse”](#) on page 938, you developed a client in a Web container and deployed it to WebSphere Application Server Community Edition. As part of the server configuration, you enabled the Axis2 facet and included the facet in the configuration of the Web container. To configure Web containers on other application servers, refer to the Axis2 documentation, http://ws.apache.org/axis2/1_4_1/installationguide.html#servlet_container, or the documentation supplied with the Web server.

Note: Axis2 use the term, Servlet container. A Servlet container is the same as a Web container.

About this task

Deploying an Axis2 client to use WebSphere MQ transport for SOAP is like deploying an Axis2 client to use HTTP. Additional steps are required to provide a classpath to the WebSphere MQ JAR files, and to modify the Axis2 configuration file. The Axis2 configuration file requires an additional entry for JMS. The entry refers to the WebSphere MQ transport for SOAP JAR file that implements the JMS transportSender.

Axis2 provides a script, `axis2.bat` or `axis2.sh`, which simplifies client deployment; see the examples in [Figure 203 on page 971](#) and [Figure 204 on page 971](#).

Note:

1. `axis2.bat` has a bug that must be corrected. The string `-Djava.ext.dirs="%AXIS2_HOME%\lib\"` must be changed to `-Djava.ext.dirs="%AXIS2_HOME%\lib\\"`.
2. In `axis2.bat` and `axis2.sh`, `-Djava.ext.dirs` is used as a quick way to reference all the Axis2 JAR files, instead of adding them separately to the classpath. Unfortunately this approach is flawed, and only works with some JREs. It does not work with the IBM JREs.

The JVM parameter, `-Djava.ext.dirs="%AXIS2_HOME%\lib\\"`, makes the Axis JAR files available to the JVM. The JVM attempts to instantiate some of the Axis JAR files, and leads to an error, the details of which depend on the JVM. Typically, you might see one of the following lines in the stack trace:

```
org.apache.axiom.om.util.UUIDGenerator.getInitialUUID(UUIDGenerator.java:76)
```

```
or org.apache.axis2.deployment.DeploymentException:  
java.security.NoSuchAlgorithmException: MD5 MessageDigest not available
```

The correct way to run an unmanaged Axis2 client is to add the Axis2 JAR files to the classpath. The classpath is available only to the client application and not to the JVM.

The procedure describes the general steps to run an unmanaged Axis2 client without using the `axis2` script. The examples in [Figure 201 on page 971](#) and [Figure 202 on page 971](#) are scripts for Windows and Linux.

Procedure

1. Download Axis2 1.4.1 from http://ws.apache.org/axis2/download/1_4_1/download.cgi and unpack into a folder, Axis2-1.4.1.
2. Update axis2.xml in Axis2-1.4.1\conf.

- a) Update axis2.xml in Axis2-1.4.1\conf. Add WebSphere MQ transport for SOAP as a transportSender:

```
<transportSender name="jms"
class="com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender"/>
```

- b) If required, alter the size of the connection pool from the default of 10.

```
<transportSender name="jms"
class="com.ibm.mq.axis2.transport.jms.WMQJMSTransportSender">
<parameter name="ResourcePoolCapacity">20</parameter>
</transportSender>
```

ResourcePoolCapacity defines how many service endpoint entries are kept in the cache. The value must be at least 1. If the number of service endpoint entries exceeds the cache size, entries are deleted to make room for new entries. The size of an endpoint entry varies. Set a number that is large enough to avoid the cache thrashing.

See step 3 in [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse”](#) on page 938.

3. Create a directory *deployDir*. Under this directory copy the folder structure containing the client and client proxies. *deployDir* is equivalent to the *project\bin* folder in an Eclipse Java project.
4. Open a command window on Windows, or a command shell using X Window System on UNIX and Linux systems, in *deployDir*.
5. Update the classpath to include the current directory, Axis2 JAR files, *com.ibm.mqjms.jar* and *com.ibm.mq.axis2.jar*.
com.ibm.mqjms.jar references all the other WebSphere MQ JAR files that are required.
6. Use the **Java** command to start the client program.

Examples

Four example of running an Axis2 client are listed in [Figure 202 on page 971](#) to [Figure 204 on page 971](#). [Figure 200 on page 970](#) shows the output from running the asynchronous client listed in [Figure 183 on page 942](#).

```
cd C:\IBM\ID\Workspaces\Axis2docs\StockQuoteAxis2PojoClient\bin>
runpojo soap/client/SQA2AsyncClient

HTTP Sync: 55.25
Callback constructor
Waiting for HTTP callback
Result in Callback 55.25
HTTP poll: 55.25
JMS: Sync: 55.25
Callback constructor
Waiting for JMS callback
Result in Callback 55.25
JMS poll: 55.25
Press any key to continue . . .
```

Figure 200. Output from running SQA2AsyncClient

```

@echo off
set AXIS2_HOME=C:\OpenSource\axis2-1.4.1
set JAVA_HOME=C:\IBM\Java50
set WMQ_HOME=C:\IBM\MQ\java\lib

setlocal EnableDelayedExpansion
set CLASSPATH=
set AXIS2_CLASS_PATH=
FOR %%c in ("%AXIS2_HOME%\lib\*.jar") DO set AXIS2_CLASS_PATH=!AXIS2_CLASS_PATH!;%%c

"%JAVA_HOME%\bin\java" -Daxis2.repo="%AXIS2_HOME%\repository"
-Daxis2.xml="%AXIS2_HOME%\conf\axis2.xml" -cp
".;%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar;%AXIS2_CLASS_PATH%" %1

pause

```

Figure 201. runpojo.bat: Windows, using a classpath

```

export AXIS2_HOME=/home/OpenSource/axis2-1.4.1
export JAVA_HOME=/usr/lib/j2sdk1.5-ibm
# update classpath
AXIS2_CLASSPATH=""
for f in "$AXIS2_HOME"/lib/*.jar
do
    AXIS2_CLASSPATH="$AXIS2_CLASSPATH":$f
done
AXIS2_CLASSPATH="$AXIS2_HOME": "$JAVA_HOME/lib/tools.jar": "$AXIS2_CLASSPATH": "$CLASSPATH"
java -cp /home/alex/dev/sandbox/Soap/axis2:/opt/mqm/java/lib/com.ibm.mqjms.jar:
/opt/mqm/java/lib/com.ibm.mq.axis2.jar:$AXIS2_CLASSPATH
-Daxis2.xml=/home/alex/dev/sandbox/axis2-1.4.1/conf/axis2.xml %1

```

Figure 202. runpojo.sh: Linux, using a classpath.

```

@echo off
set AXIS2_HOME=C:\OpenSource\axis2-1.4.1
set JAVA_HOME=C:\IBM\Java50
set WMQ_HOME=C:\IBM\MQ\java\lib

%AXIS2_HOME%\bin\axis2 -cp .;%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar; %1
pause

```

Figure 203. runaxis2.bat: Windows, using axis2.bat

Note

```

export AXIS2_HOME=/home/OpenSource/axis2-1.4.1
export JAVA_HOME=/usr/lib/j2sdk1.5-ibm

%AXIS2_HOME%\bin\axis2 -cp .;%WMQ_HOME%\com.ibm.mqjms.jar;%WMQ_HOME%\com.ibm.mq.axis2.jar; %1

```

Figure 204. runaxis2.sh: Linux, using axis2.sh

Note

Related tasks

[Deploying a Web service client to Axis 1.4 to use IBM WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM WebSphere MQ queues and channels, start the service and test the client.

[Deploying to an Axis2 client using W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 4 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application Server using the W3C SOAP over JMS protocol. Modify

the URL in the Axis2 client developed for WebSphere MQ transport for SOAP to use the W3C candidate recommendation for SOAP over JMS.

Deploying a Web service client to .NET Framework 1 and 2 to use WebSphere MQ transport for SOAP
Prepare a deployment directory and deployment descriptor for the client. Provide the client proxy and client class. Configure WebSphere MQ queues and channels, start the service and test the client.

Deploying to an Axis2 client using W3C SOAP over JMS

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 4 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application Server using the W3C SOAP over JMS protocol. Modify the URL in the Axis2 client developed for WebSphere MQ transport for SOAP to use the W3C candidate recommendation for SOAP over JMS.

Before you begin

You must first complete the task, “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 938 to call **SimpleJavaListener** using an Axis2 client and the WebSphere MQ transport for SOAP protocol.

You must also have created the Web service and configured WebSphere MQ and WebSphere Application Server in the previous tasks:

1. “Configure WebSphere MQ resources” on page 962.
2. “Configure WebSphere Application Server resources” on page 963.
3. “Developing a JAX-WS EJB Web service for W3C SOAP over JMS” on page 925.

In the task, the client runs in Eclipse Galileo. You might run the client from the command line by modifying the `Axis2.bat` file shipped with Axis2.

About this task

The only change you must make to the existing `Axis2 StockQuoteAxis` static client to call the `StockQuoteAxis` service hosted by WebSphere Application Server is to change the URL passed to the client. Since the WSDL has not changed, you can use the same proxy classes in the `soap.server` package.

You have two approaches to defining the URL to pass to the client. You might use the same URL as in the generated `StockQuoteAxis.wsdl`. You must add the `jndiInitialContextFactory` and the `jndiURL` parameters to access the WebSphere Application Server JNDI directory. Another approach is to change the URL and give the client direct access to the `REQUESTAXIS` and `REPLYAXIS` queues on `QM1`, without using a JNDI lookup.

The connection parameters you define in the URL passed to the Axis2 client are used to connect to the WebSphere MQ queue manager and queues required to send and receive SOAP messages. The connection parameters passed to the Axis2 client are not necessarily used by the service. You can use the distributed queuing capabilities of WebSphere MQ to decouple the client and service from using the same queue manager, or the same name server.

Procedure

1. Save the URL from the generated `StockQuoteAxis.wsdl` and close down Rational Application Developer to save on memory.

If you did not change the server configuration, closing Rational Application Developer stops the application server. In which case start the server with the command:

```
startserver server1
```

2. Open Eclipse Galileo in the workspace with the Axis2 client project.
3. Open `SQA2StaticClient.java`.

See `SQA2StaticClient.java`.

4. Call the service using the queue variant of the URI.

a) Modify the URL.

The new URI is:

```
jms:queue:REQUESTAXIS
    ?replyToName=REPLYAXIS
    &connectionFactory=connectQueueManager(QM1)Bind(Server)
    &targetService=StockQuoteAxis;
```

Compare this to the URL from `StockQuoteAxis.wsdl`:

```
jms:jndi:requestaxis
    ?jndiConnectionFactoryName=qm1
    &targetService=StockQuoteAxis
```

Figure 205. URL from `StockQuoteAxis.wsdl`

- REQUESTAXIS is now uppercase as it is a queue name and not a JNDI name.
 - The connection to QM1 is straightforward.
 - The URI does not contain the name of the Reply to destination. The client must define the queue it expects replies on.
- b) Run `SQA2StaticClient.java` using the same **Run as...** configuration as you did in the task, “Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 938.
5. Call the service using the `jndi` variant of the URI, using WebSphere Application Server as the naming server.
- a) Use the URL from `StockQuoteAxis.wsdl`, Figure 205 on page 973, providing the missing parameters to use the naming service in WebSphere Application Server.

The missing parameters and values you must provide are:

Table 151. Additional JNDI parameters

Parameter	Value used in this example	Description
&jndiURL	iiop://localhost:2810 or corbaname:iiop:localhost:2810	URI of naming provider. For WebSphere Application Server the value defaults to 2809. It is also known as the port number of the RMI connector, and the bootstrap port. The value is listed in the <code>SystemOut.log</code> 00000000 NameServerImp A NMSV0018I: Name server available on bootstrap port 2810
&jndiInitialContextFactory	com.ibm.websphere.naming. WsnInitialContextFactory	The name of the initial context factory used by WebSphere Application Server.
&replyToName	replyaxis	JNDI name of REPLYAXIS queue.

```
jms:jndi:requestaxis?
    &jndiURL=iiop://localhost:2810
    &jndiConnectionFactoryName=qm1
    &jndiInitialContextFactory=com.ibm.websphere.naming.WsnInitialContextFactory
    &targetService=StockQuoteAxis
    &replyToName=replyaxis;
```

- b) Add the JAR files required by the JNDI lookup.

In this configuration, the following JAR files were added to the build path to run the task using the `jndi` variant of the JMS url:

- `com.ibm.jaxws.thinclient_7.0.0.jar` from *Rational install directory\SDP\runtimes\base_v7\runtimes*.
- `com.ibm.ws.runtime.jar` from *Rational install directory\SDP\runtimes\base_v7\plugins*

For a different JNDI provider you require different JAR files.

The other JAR files in the build path are:

- i) All the JAR files in *WebSphere MQ Install directory\java\lib*.
 - ii) All the JAR files in *Axis2-1.5.1\lib*.
 - iii) Java 6.0 JRE.
- c) Run `SQA2StaticClient.java` using the same **Run as...** configuration as you did in the task, [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 938](#).

Results

In both cases the reply from the service is displayed in the client console view.

Related tasks

[Deploying a Web service client to Axis 1.4 to use IBM WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM WebSphere MQ queues and channels, start the service and test the client.

[Deploying a Web service client to Axis2 to use WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and Axis2 configuration file for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure WebSphere MQ queues and channels, start the service and test the client.

[Deploying a Web service client to .NET Framework 1 and 2 to use WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxy and client class. Configure WebSphere MQ queues and channels, start the service and test the client.

Deploying a Web service client to .NET Framework 1 and 2 to use WebSphere MQ transport for SOAP

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxy and client class. Configure WebSphere MQ queues and channels, start the service and test the client.

Before you begin

Tip: Develop and test the service and the client using Visual Studio. Then modify the client for WebSphere MQ transport for SOAP.

1. If you are deploying a service using .NET Framework 1 or 2, build the service as a library (.dll). Deploy using WebSphere MQ transport for SOAP.
2. Add the `Register.Extension()` call to the client.
3. Add a reference to `amqsoap.dll`, which is located in *MQ_Install\bin*.
4. Change the static `Url` property in the client proxy class constructor to the `jms:/URI`, for WebSphere MQ transport for SOAP.

About this task

Deploying a Web service client for .NET Framework 1 or 2 to use WebSphere MQ transport for SOAP requires an additional deployment step. You need to register `amqsoap.dll` with the .NET Framework.

amqsoap.dll is automatically registered as part of installing WebSphere MQ transport for SOAP, but you might need to register it again.

If you use the command **amqwdeployMQService** to generate client proxies, you can deploy the client using the directories the command generates.

Procedure

1. Create a directory *deployDir* to hold the client deployment files.
2. Open a command window in *deployDir*.
3. Run **amqwsetcp** to set the CLASSPATH if the service is to run on Axis 1.4.
4. If necessary, run **amqwRegisterDotNet** to register amqsoap.dll with the .NET Framework.

Example

The example shows the configuration and output, [Figure 208 on page 975](#), from an .NET Framework V2 client. The client, [Figure 207 on page 975](#), calls a Web service that echoes its input parameter. The static Url definition, [Figure 206 on page 975](#), shows the constructor for the client proxy.

```
public Quote() {
    this.Url = "jms:/queue?destination=REQUESTDOTNET
               &connectionFactory=(connectQueueManager(QM1)binding(server))
               &initialContextFactory=com.ibm.mq.jms.Nojndi
               &targetService=Quote.asmx
               &replyDestination=SYSTEM.SOAP.RESPONSE.QUEUE";
}
```

Figure 206. Static client proxy constructor

```
using System;
namespace QuoteClientProgram {
    class QuoteMain {
        static void Main(string[] args) {
            try {
                IBM.WMQSOAP.Register.Extension();
                Quote q = new Quote();
                Console.WriteLine("Response is: " + q.getQuote("ibm"));
            } catch (Exception e) {
                Console.WriteLine("Exception is: " + e);
            }
        }
    }
}
```

Figure 207. Client program

```
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram>dir /s /b
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram\QuoteClientProgram.exe
C:\IBM\ID\DotNet\QuoteClientProgram\QuoteClientProgram>quoteclientprogram
Response is: IBM
```

Figure 208. Configuration and output

What to do next

1. If you are deploying the client as a WebSphere MQ MQI client, configure the client and server connection channel.

2. If you are deploying the client to a different queue manager to the service, you must make the destination queue available to the client. Configure the destination queue on the service queue manager as a cluster queue, or on the client queue manager as a remote queue definition.

Related tasks

[Deploying a Web service client to Axis 1.4 to use IBM WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and deployment descriptor for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure IBM WebSphere MQ queues and channels, start the service and test the client.

[Deploying a Web service client to Axis2 to use WebSphere MQ transport for SOAP](#)

Prepare a deployment directory and Axis2 configuration file for the client. Provide the client proxies and client class, and set up the CLASSPATH. Configure WebSphere MQ queues and channels, start the service and test the client.

[Deploying to an Axis2 client using W3C SOAP over JMS](#)

A Web service bound to the W3C candidate recommendation for SOAP over JMS must run in the EJB container of a Java EE application server. This task is step 4 of connecting an Axis2 Web service client and a Web service deployed to WebSphere Application Server using the W3C SOAP over JMS protocol. Modify the URL in the Axis2 client developed for WebSphere MQ transport for SOAP to use the W3C candidate recommendation for SOAP over JMS.

Connect an Axis2 client to a JAX-WS service using W3C SOAP over JMS and WebSphere Application Server

When you complete this task you will have called a JAX-WS Web service running in WebSphere Application Server from an Axis2 client. The Axis2 client and WebSphere Application Server use the W3C candidate recommendation for the SOAP over JMS protocol running on WebSphere MQ. Use Eclipse Galileo and Rational Application Developer to build the Web service client and Web service, respectively.

Before you begin

The task requires version 7 of Rational Software Development Environment and WebSphere Application Server. The task was created using the Rational Application Developer packaged with Rational Software Architect for WebSphere Software v7.5.5.1, and WebSphere Application Server Version 7.0 Test Environment v7.0.0.9 Update 1. You also require WebSphere MQ v7.0.1.3.

The task builds on two other tasks, [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse”](#) on page 919, and [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse”](#) on page 938. To complete these tasks your development environment already has Eclipse Galileo, WASCE, the Eclipse plugin for WASCE, and Axis2 1.4.1 installed. You do not require WASCE for this task.

Some of the steps are complex. The steps assume some familiarity with developing Web service applications for WebSphere Application Server using Rational Application Developer. The processor and memory demands of the task are large. The task was performed in a VMWare Windows XP SP3 virtual machine allocated 1.8GB of memory.

Install all the software before starting the task. The software takes about a day to download and a day to install, depending on bandwidth. The task takes at least half a day.

About this task

The scenario for this task is that you have developed a stock quotation Web service, StockQuoteAxis, using an open source tool, Eclipse Galileo. StockQuoteAxis is deployed using SOAP over HTTP running on an open source server, WASCE.

You want to bind the Web services you deploy to a standards-based messaging transport, such as SOAP over JMS, or to Web services reliable messaging, as well as to SOAP over HTTP. You want both the client and service to use standards-based interfaces. For this reason, although your future projects

development team have implemented a solution using WebSphere MQ transport for SOAP, you have not gone into production.

The Axis2 client has removed the problem that the SOAP client for the WebSphere MQ transport for SOAP required a change from the HTTP client. The problem still remained that the service connected by the IBM WebSphere MQ transport for SOAP is hosted by a special listener provided by WebSphere MQ: `SimpleJavaListener`.

With the W3C SOAP over JMS standard in candidate recommendation status, some vendors are providing support for W3C SOAP over JMS. The support enables you to deploy a Web service to an application server and connect to the same service using a variety of connectivity protocols. The support provided by WebSphere Application Server v7 removes problem of having to host the Web service separately in order to use a message-based SOAP transport. The use of a standards-based message transport interface, JMS, means you can develop solutions using tools from different vendors. You hope the Web services tools in Eclipse will include SOAP over JMS bindings in the future.

Most of the steps are performed using Eclipse, or the management tools provided with the WebSphere products. The steps are described for a Windows environment. With slight modifications to some commands, you can perform the steps on other platforms.

The preliminary steps creating the HTTP Web service, and connecting to it using Axis2 are listed. The client, and WSDL, from these steps are used to create the solution

Procedure

1. Connect to the StockQuoteAxis Web service using an Axis2 client and IBM WebSphere MQ transport for SOAP
 - a) [“Developing a JAX-RPC service for WebSphere MQ transport for SOAP using Eclipse” on page 919](#)
 - b) [“Developing a JAX-WS client for WebSphere transport for SOAP using Eclipse” on page 938](#)
 - c) [“Deploying a Web service client to Axis2 to use WebSphere MQ transport for SOAP” on page 969](#)
2. Connect to the StockQuoteAxis Web service using an Axis2 client and the W3C candidate recommendation for SOAP over JMS.
 - a) [“Configure WebSphere MQ resources” on page 962](#)
 - b) [“Configure WebSphere Application Server resources” on page 963](#)
 - c) [“Developing a JAX-WS EJB Web service for W3C SOAP over JMS” on page 925](#)
 - d) [“Deploying to an Axis2 client using W3C SOAP over JMS” on page 972](#)

WebSphere MQ bridge for HTTP

With WebSphere MQ bridge for HTTP, client applications can exchange messages with WebSphere MQ without the need to install a WebSphere MQ MQI client. You can call WebSphere MQ from any platform or language with HTTP capabilities.

Introduction to WebSphere MQ bridge for HTTP

The WebSphere MQ bridge for HTTP is a Java, Enterprise Environment (JEE) Web application. HTTP clients can send **POST**, **GET**, and **DELETE** requests to it to put, browse and delete messages from WebSphere MQ queues. The WebSphere MQ bridge for HTTP is not suitable for use with messages, if assured delivery is required.

Benefits

With the WebSphere MQ bridge for HTTP you can send and receive WebSphere MQ messages using HTTP from a wide variety of environments:

- Environments that support HTTP, but not WebSphere MQ.
- Environments that have insufficient storage space to install a WebSphere MQ MQI client.

- Environments that are too numerous to install the WebSphere MQ MQI client on each system that requires access to WebSphere MQ.
- Web-based applications from which you want to send or receive messages without coding your own bridge to WebSphere MQ.
- Web-based applications that you want to enhance, using asynchronous techniques such as AJAX. WebSphere MQ bridge for HTTP makes WebSphere MQ queues and topics available using Representation State Transfer (REST) over HTTP.

HTTP support can be used with both point-to-point and publish/subscribe messaging topologies.

How does HTTP support work?




Figure 209. WebSphere MQ bridge for HTTP

The WebSphere MQ bridge for HTTP Web application receives HTTP requests from one or more clients. It interacts with WebSphere MQ on their behalf, and returns HTTP responses to them.

The WebSphere MQ bridge for HTTP is a JEE servlet that is connected to WebSphere MQ using a resource adapter. The HTTP servlet handles three different types of HTTP requests: **POST**, **GET**, and **DELETE**.

Table 152. WebSphere MQ bridge for HTTP verbs	
HTTP Request	Result
POST	Puts a message on a queue or topic.
GET	Browses the first message on a queue. In line with the HTTP protocol, GET does not delete the message from the queue. Do not use GET with publish/subscribe messaging.
DELETE	Gets and deletes a message from a queue or topic.

HTTP POST example

HTTP **POST** puts a message to a queue, or a publication to a topic. The **HTTPPOST** Java sample is an example an HTTP **POST** request of a message to a queue. Instead of using Java, you could create an HTTP **POST** request using a browser form, or an AJAX toolkit instead.

Figure 210 on page 978 shows an HTTP request to put a message on a queue called myQueue. This request contains the HTTP header `x-msg-correlID` to set the correlation ID of the WebSphere MQ message.

```

POST /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
Content-Type: text/plain
x-msg-correlID: 1234567890
Content-Length: 50

Here is my message body that is posted on the queue.
  
```

Figure 210. Example of an HTTP **POST** request to a queue

Figure 211 on page 979 shows the response sent back to the client. There is no response content.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 0
```

Figure 211. Example of an HTTP POST response

HTTP DELETE example

HTTP **DELETE** gets a message from a queue and deletes the message, or retrieves and deletes a publication. The **HTTPDELETE** Java sample is an example an HTTP **DELETE** request reading a message from a queue. Instead of using Java, you could create an HTTP **DELETE** request using a browser form, or an AJAX toolkit instead.

Figure 212 on page 979 is an HTTP request to delete the next message on queue called myQueue. In response, the message body is returned to the client. In WebSphere MQ terms, HTTP **DELETE** is a destructive get.

The request contains the HTTP request header x-msg-wait, which instructs WebSphere MQ bridge for HTTP how long to wait for a message to arrive on the queue. The request also contains the x-msg-require-headers request header, which specifies that the client is to receive the message correlation ID in the response.

```
DELETE /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
x-msg-wait: 10
x-msg-require-headers: correlID
```

*Figure 212. Example of an HTTP **DELETE** request*

Figure 213 on page 979, is the response returned to the client. The correlation ID is returned to the client, as requested in x-msg-require-headers of the request.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 50
Content-Type: text/plain; charset=utf-8
x-msg-correlId: 1234567890

Here is my message body that is retrieved from the queue.
```

*Figure 213. Example of an HTTP **DELETE** response*

HTTP GET example

HTTP **GET** gets a message from a queue. The message remains on the queue. In WebSphere MQ terms, HTTP **GET** is a browse request. You could create an HTTP **GET** request using a Java client, a browser form, or an AJAX toolkit.

Figure 214 on page 980 is an HTTP request to browse the next message on queue called myQueue.

The request contains the HTTP request header x-msg-wait, which instructs WebSphere MQ bridge for HTTP how long to wait for a message to arrive on the queue. The request also contains the x-msg-require-headers request header, which specifies that the client is to receive the message correlation ID in the response.

```
GET /msg/queue/myQueue/ HTTP/1.1
Host: www.example.org
x-msg-wait: 10
x-msg-require-headers: correlID
```

*Figure 214. Example of an HTTP **GET** request*

Figure 215 on page 980 is the response returned to the client. The correlation ID is returned to the client, as requested in `x-msg-require-headers` of the request.

```
HTTP/1.1 200 OK
Date: Wed, 2 Jan 2007 22:38:34 GMT
Server: Apache-Coyote/1.1 WMQ-HTTP/1.1 JEE-Bridge/1.1
Content-Length: 50
Content-Type: text/plain; charset=utf-8
x-msg-correlId: 1234567890
```

Here is my message body that appears on the queue.

*Figure 215. Example of an HTTP **GET** response*

Installing, configuring, and verifying WebSphere MQ bridge for HTTP

Obtain WebSphere MQ bridge for HTTP by installing "Java Messaging and Web Services" from either the WebSphere MQ MQI client or server installation materials. Deploy WebSphere MQ bridge for HTTP to a suitable application server.

Before you begin

Check the prerequisite products at [System Requirements for IBM WebSphere MQ](#). The installation process does not check for the presence and availability of the prerequisite software for running WebSphere MQ bridge for HTTP. You must verify that the prerequisites are installed.

WebSphere MQ bridge for HTTP runs on any Java EE 1.4 compliant application server, by installing the WebSphere MQ resource adapter. You can also run WebSphere MQ bridge for HTTP on a WebSphere Application Server release earlier than version 6.0.2.1. Use the WebSphere Application Server Message Listener Port (MLP) to integrate WebSphere MQ as the JMS provider.

Support for WebSphere MQ bridge for HTTP is provided only for the following application servers:

- WebSphere Application Server 6.0.2.1 and later.
- WebSphere Application Server Community Edition Version 1.1 and later.

About this task

WebSphere MQ bridge for HTTP is supplied as a `.war` file, `WMQHTTP.war`.

- On UNIX platforms and Linux,
 - `WMQHTTP.war` is included as part of the "Java Messaging and Web Services" install option. The option is available in both the client and server installation materials.
 - `WMQHTTP.war` is installed to `<mqmtop>/java/http/WMQHTTP.war`. `<mqmtop>` is the directory where WebSphere MQ is installed.
 - `WMQHTTP.samples` is installed to `<mqmtop>/java/http/samples`. `<mqmtop>` is the directory where WebSphere MQ is installed.

Carry out the following installation steps to install WebSphere MQ bridge for HTTP, deploy, and configure it, and verify the configuration. The details of the configuration steps vary on different application servers.

Use [“Deploying and verifying WebSphere MQ bridge for HTTP on WebSphere Application Server V6.1.0.9”](#) on page 981 as a template for the steps to follow on your application server.

Procedure

1. Obtain WMQHTTP.war by installing either the WebSphere MQ MQI client or server.
2. Copy WMQHTTP.war to a server from which it can be deployed to an application server.
3. Deploy WMQHTTP.war to an application server.
4. If necessary, install WebSphere MQ as a resource adapter on your application server.
Find out if WebSphere MQ is already configured as a messaging provider on your application server. Use the administration or management tool supplied with your application server, to look for WebSphere MQ. WebSphere MQ might be found under the following path, **Resources > JMS > Messaging providers**.
5. Configure a connection factory on the application server to connect to a queue manager that uses the WebSphere MQ MQI client transport¹².
6. Configure the WMQHTTP.war Web application on the application server to use the connection factory
7. Verify the configuration.
 - a) Set up the queue manager named in the connection factory and a local queue.
 - b) Place a message on the local queue.
 - c) Create the server-connection channel named in the connection factory, with authority to read and write to the local queue.
 - d) Start the queue manager and the listener.
 - e) Start the application server and WMQHTTP.war, if they are not already running.
 - f) Open a browser and type `http://hostname:web_port/Context root/msg/queue/local queue`

Results

The browser window displays the message you placed on the local queue.

What to do next

1. Try the example, [“Deploying and verifying WebSphere MQ bridge for HTTP on WebSphere Application Server V6.1.0.9”](#) on page 981.
2. Run the sample HTTP Java applications.

Deploying and verifying WebSphere MQ bridge for HTTP on WebSphere Application Server V6.1.0.9

Use the following example to prepare a deployment of WebSphere MQ bridge for HTTP to run the sample HTTP Java programs. The deployment is on WebSphere Application Server V6.1.0.9.

Before you begin

1. Follow the instructions in [“Installing, configuring, and verifying WebSphere MQ bridge for HTTP”](#) on page 980, to copy WMQHTTP.war onto a server accessible to your installation of WebSphere Application Server.
2. Configure a queue manager, and a queue, to use to test the configuration:
 - In the example, the queue manager is configured as using the values in [Table 153](#) on page 982:

¹² Initially, at least, configure the client transport. Some application servers can connect to WebSphere MQ by using direct, or bindings mode connections.

Table 153. Queue manager configuration	
Object	Value
Host name	itso-01
Queue manager	QM1
Local queue	HTTPTESTQ
Server connection channel	MYSVRCON. Configure an MCA user ID with sufficient authority to read and write to HTTPTESTQ.
Listener port	1414

3. Start the queue manager and the listener
4. Place a test message onto HTTPTESTQ. For example:
 - a. Start WebSphere MQ Explorer.
 - b. In the list of local queues for QM1, right-click **HTTPTESTQ > Put test message > type First Message > Put message > Close**
5. Start the application server and sign on to the Integrated Solutions Console.

About this task

The example shows the steps to take if you are running WebSphere Application Server V6.1.0.9 as your application server. If you are running a different version of WebSphere Application Server, or running a different application server, the steps are different. WebSphere Application Server V6.1.0.9 is pre-configured with WebSphere MQ installed as a message provider, using the WebSphere MQ MQI client libraries. If WebSphere MQ is not pre-configured as a messaging provider, or if you want to use WebSphere MQ server bindings, you need to install and configure the WebSphere MQ resource adapter for JEE into your application server.

Follow the instructions to deploy WebSphere MQ bridge for HTTP onto WebSphere Application Server V6.1.0.9, and verify the deployment using a browser:

Procedure

1. In the navigation pane, click **Resources > JMS providers > WebSphere MQ messaging provider**.
You can configure at either Node, Cell, or Server level, depending on your WebSphere Application Server deployment. The example uses Server level deployment.
2. Under **Additional properties**, click **Connection factories > New**.
3. In the JMS providers form, provide the information in [Table 154 on page 982](#), or alternatives of your choosing, click **Apply > Save**.

Table 154. Set or modify the following fields	
Field	Value
Name	WMQHTTPBridge
JNDI Name	jms/WMQHTTPJCAConnectionFactory
Queue manager	QM1
Host	itso-01
Port	1414
Channel	MYSVRCON
Transport type	CLIENT

4. In the navigation pane, click **Applications > Install New Application**.
5. Insert the path to `WMQHTTP.war` into the form, and provide a Context root, click **Next**.
 - a) The Context root is optional. `mq` is the default Context root for the sample HTTP applications.
 - b) The Context root forms part of the URI identifying WebSphere MQ bridge for HTTP. You can omit the Context root, or change it later.
6. On the **Select installation options** page of the installation wizard, you do not have to change any of the defaults, click **Next**.
7. On the **Map modules to servers** page, select a Cluster or Server, check the Select box, click **Apply > Next**.
8. On the **Map resource references to resources** page, in the `javax.jms.ConnectionFactory` form, click **Browse...** on the IBM WebSphere MQ bridge for HTTP row.
9. On the **Enterprise Applications > Available resources** page, select **WMQHTTPBridge**, click **Apply**.
10. Back in the `javax.jms.ConnectionFactory` form, select the authentication method.
 - a) For the example, choose **None**, click **Apply**. The other options require additional configuration.
11. Check the **Select** check box for IBM WebSphere MQ bridge for HTTP, click **Next > Next > Finish > Save**.
12. In the navigation pane, click **Applications > Enterprise Applications**.
13. Check the selection box for `WMQHTTP.war`, click **Start**.
14. Open a browser window. Type `http://itso-01:9080/mq/msg/queue/HTTPTESTQ`, using the appropriate host name and port.

Results

The browser window displays `First Message`, if the configuration is successful.

What to do next

Run the sample HTTP Java applications.

Publish/subscribe using the WebSphere MQ bridge for HTTP

WebSphere MQ bridge for HTTP uses the WebSphere MQ classes for JMS publish/subscribe interface. HTTP **POST** creates a publication. HTTP **DELETE** creates a non-durable managed subscription. You must configure publish/subscribe for JMS before using the topic URI.

Publish/subscribe is fully integrated into WebSphere MQ in version 7. Before version 7, a separate publish/subscribe broker handled publications and subscriptions. It is called "queued" publish/subscribe, to distinguish it from the fully integrated publish/subscribe in version 7. Version 7 emulates queued publish/subscribe using integrated publish/subscribe. The emulation enables existing queued publish/subscribe applications to coexist with integrated applications running on the same queue manager. Queued publish/subscribe applications can also interoperate with integrated applications, sharing the same topics. In version 6, the broker was shipped with WebSphere MQ; before version 6 it was available as a SupportPack.

Configuration

The WebSphere MQ bridge for HTTP uses the JMS interface to publish and subscribe. In version 7, you can control whether the WebSphere MQ classes for JMS use queued or integrated publish/subscribe, using the `PROVIDERVERSION` JMS property.

An additional consideration is that you can use either WebSphere MQ MQI client libraries with WebSphere MQ bridge for HTTP, or server libraries. Version 6 client libraries only support queued publish/subscribe, whereas version 7 libraries support both queued and integrated publish/subscribe. Most Web or application servers that use WebSphere MQ as a messaging provider do so using client libraries. In order to use integrated publish/subscribe, both the WebSphere MQ MQI client and server libraries must be at

least at version 7. If either is running an earlier version of WebSphere than 7, then you must configure queued publish/subscribe; see [Table 155 on page 984](#). Check what libraries are installed or configured with the Web server or application server you are using.

Table 155. Publish/subscribe configuration modes		
	Client V6 or earlier	Client V7 or later
Server V6 or earlier	1. Run the <code>\java\bin\MQJMS_PSQ.mq</code> <code>sc</code> script	Not supported
Server V7 or later	1. Run the <code>\java\bin\MQJMS_PSQ.mq</code> <code>sc</code> script 2. Set the queue manager to <code>PSMODE=ENABLED</code>	1. If PROVIDERVERSION = 7 a. Set the queue manager to <code>PSMODE=ENABLED</code> or <code>PSMODE=COMPAT</code> 2. If PROVIDERVERSION = 6 a. Set the queue manager to <code>PSMODE=ENABLED</code>

Publish

Send an HTTP **POST** request with the URI:

```
http://hostname:port/context_root/msg/topic/topicString
```

The message contents are published using the topic string *topicString*.

Subscribe

Send an HTTP **DELETE** request with the URI:

```
http://hostname:port/context_root/msg/topic/topicString
```

WebSphere MQ bridge for HTTP creates a managed non-durable subscription to the topic string *topicString*. The subscription is deleted as soon as a publication is returned, or until the wait-interval set by the custom entity-header, `x-msg-wait`, expires.

Running the WebSphere MQ bridge for HTTP samples

The WebSphere MQ bridge for HTTP samples are available for use on only the Windows operating system. The samples show you how to submit HTTP **POST** and HTTP **DELETE** commands to WebSphere MQ bridge for HTTP from Java programs.

Before you begin

Verify your WebSphere MQ bridge for HTTP installation by running step “7” on [page 981](#) in “[Installing, configuring, and verifying WebSphere MQ bridge for HTTP](#)” on [page 980](#).

The HTTP samples are installed to the directories shown in [Table 156 on page 984](#). In each case, source code is installed to the `/src` subdirectory.

Table 156. Location of HTTP samples	
Platform	Location
Windows	<code>MQ_INSTALLATION_PATH/tools/http/samples</code>
All other platforms	<code>MQ_INSTALLATION_PATH/samp/http</code>


Table 156. Location of HTTP samples (continued)	
Platform	Location
<i>MQ_INSTALLATION_PATH</i> represents the directory where WebSphere MQ is installed.	

About this task

The samples simulate the WebSphere MQ AMQSPUT and AMQSGET sample applications. They illustrate the following functions in a point-to-point messaging environment:

- **HTTPPOST** - Sends HTTP **POST** requests in a Java application to put messages to a WebSphere MQ queue, using the WebSphere MQ bridge for HTTP and handles the responses.
- **HTTPDELETE** - Sends HTTP **DELETE** requests in a Java application to get messages from a WebSphere MQ queue, using the WebSphere MQ bridge for HTTP and handles the responses containing the WebSphere MQ message.

Parameters for HTTPPOST and HTTPDELETE



To run the **HTTPPOST** sample, complete the following steps:

Procedure

1. In a command window, navigate to the HTTP samples directory.
2. Run the **HTTPPOST** sample.

```
java -classpath . HTTPPOST [parameters]
```

When the **HTTPPOST** sample starts, the following output is displayed:

```
HTTP POST Sample start
Target server is 'hostName'
Target queue is 'queueName'
Target context-root is 'contextRoot'
```

3. In the command prompt, type the text that you want to form your message body.
4. Press enter to post the message to the WebSphere MQ queue.
 - a) If you want to send another message, enter some more text.
The texts forms the body of a second WebSphere MQ message.
 - b) Press enter to post the message to the WebSphere MQ queue.
5. Press enter twice to end **HTTPPOST**.

The following output is displayed:

```
HTTP POST Sample end
```

What to do next

The **HTTPDELETE** sample performs a destructive get of all the messages you placed on the WebSphere MQ queue.

Run the **HTTPDELETE** sample by completing the following steps:

1. In a command window, navigate to *MQ_INSTALLATION_PATH*/tools/samples.
MQ_INSTALLATION_PATH represents the directory where WebSphere MQ is installed.

2. Run the **HTTPDELETE** sample.

```
java -classpath . HTTPPOST [parameters]
```

When the **HTTPDELETE** sample starts, the following output is displayed:

```
HTTP DELETE Sample start
Target server is 'host:port'
Target queue is 'your queue name'
Target context-root is 'your context-root'
message
message
...
```

Security considerations for WebSphere bridge for HTTP

Standard Web security considerations apply to authenticating a Web browser client. Authorization to WebSphere MQ resources is at the level of the user running the WebSphere Bridge for HTTP servlet, and not the individual Web browser client. Standard WebSphere MQ security consideration apply to WebSphere MQ.

Data flowing from a Web browser to a WebSphere MQ application using WebSphere bridge for HTTP, and back, takes a three steps:

Client connection

From the browser to the WebSphere Bridge for HTTP over a TCP/IP connection using HTTP.

Resource adapter connection to WebSphere MQ

The connection is from the WebSphere Bridge for HTTP to a WebSphere MQ queue manager. The connection is either a client connection, over TCP/IP, or a local WebSphere MQ bindings connection. Once the connection is made, the HTTP request is placed on a standard local queue or a transmission queue.

From the WebSphere MQ local queue over one or more channels, to the target queue.

Apply standard techniques for securing queues, topics, queue managers, and channels.

The reply takes the steps in reverse.

Client connection

Secure connections between HTTP clients and the application server using the Web container. Use standard HTTP server techniques, such as using HTTPS. Refer to the documentation for your application server for information.

Resource adapter connection to WebSphere MQ

The connection between the resource adapter and queue manager is authorized using only a single user ID. Assign a single user ID to identify requests from the WebSphere Bridge for HTTP. The user ID must have restricted WebSphere MQ authorizations only to the resources external users must have access. You must authenticate the actual client separately, and establish trust for successive interactions with the client, using standard techniques for Web security.

Secure the connection between the resource adapter and the queue manager using the single user ID. Restrict the authorities the user ID has to no more than needed to read and write messages to queues and topics. The WebSphere Bridge for HTTP is a point of attack between the internet and your intranet.

How you secure the connection between your resource adapter and WebSphere MQ is dependent on your specific resource adapter. Refer to the documentation for the resource adapter.

Using the Component Object Model Interface (WebSphere MQ Automation Classes for ActiveX)

The WebSphere MQ Automation Classes for ActiveX (MQAX) are ActiveX components that provide classes that you can use in your application to access WebSphere MQ.

MQAX requires a WebSphere MQ environment and a corresponding WebSphere MQ application with which to communicate.

It gives your ActiveX application the ability to run transactions and access data on any of your enterprise systems that you can access through WebSphere MQ.

WebSphere MQ Automation Classes for ActiveX:

- Give you access to the functions and features of the WebSphere MQ API, permitting full interconnectivity to other WebSphere MQ platforms.
- Conform to the normal conventions expected of an ActiveX component.
- Conform to the WebSphere MQ object model, also available for .NET, C++, Java, and LotusScript®.

MQAX starter samples are provided. You can use these samples initially to check that your installation of the MQAX is successful and that you have the basic WebSphere MQ environment in place. The samples also demonstrate how MQAX can be used.

COM and ActiveX scripting

The Component Object Model (COM) is an object-based programming model defined by Microsoft. It specifies how software components can be provided in a way that allows them to locate and communicate with each other irrespective of the computer language in which they are written or their location.

ActiveX is a set of technologies, based on COM, that integrates application development, reusable components, and Internet technologies on the Microsoft Windows platforms. ActiveX components provide interfaces that can be accessed dynamically by applications. An ActiveX scripting client is an application, for example a compiler, that can build or execute a program or script that uses the interfaces provided by ActiveX (or COM) components.

WebSphere MQ environment support

WebSphere MQ Automation Classes for ActiveX can only be used with **32-bit** ActiveX scripting clients.

The COM component can only be used for **32-bit** applications. If you want to write 64-bit COM application, you can use the .NET interface.

To run the MQAX in a WebSphere MQ server environment you must have Windows 2000 or later installed on your system.

To run the MQAX in a WebSphere MQ MQI client environment you need WebSphere MQ MQI client on Windows 2000 or later installed on your system:

The WebSphere MQ MQI client requires access to at least one WebSphere MQ server. When both the WebSphere MQ MQI client and the WebSphere MQ server are installed on your system, MQAX applications always run against the server. The ActiveX interface to the MQAI is only available in WebSphere MQ server environments.

Designing and programming using WebSphere MQ Automation Classes for ActiveX

Designing MQAX applications that access non-ActiveX applications

The WebSphere MQ Automation Classes provide access to the functions of the WebSphere MQ API. You can therefore benefit from all the advantages that using WebSphere MQ can bring to your Windows application.

The overall design of your application is the same as for any WebSphere MQ application, so consider all the design aspects described in the [“Developing applications” on page 7](#) section.

To use the WebSphere MQ Automation Classes, you code the Windows programs in your application using a language that supports the creation and use of COM objects. For example, Visual Basic, Java, and other ActiveX scripting clients. The classes can then be easily integrated into your application because the WebSphere MQ objects you need can be coded using the native syntax of the implementation language.

Using WebSphere MQ Automation Classes for ActiveX

When designing an ActiveX application that uses WebSphere MQ Automation Classes for ActiveX, the most important item of information is the message that is sent or received from the remote WebSphere MQ system. Therefore you must know the format of the items that are inserted into the message. For an MQAX script to work, both it and the WebSphere MQ application that picks up or sends the message must know the message structure.

If you are sending a message with an MQAX application and you want to perform data conversion at the MQAX end, you must also know:

- The code page used by the remote system
- The encoding used by the remote system

To keep your code portable, it is good practice to set the code page and encoding, even if they are currently the same in both the sending and receiving systems.

When considering how to structure the implementation of the system you design, remember that your MQAX scripts run on the same machine as the one on which you have either the WebSphere MQ queue manager or the WebSphere MQ client installed.

Programming hints and tips

The following hints and tips are in no significant order. They are subjects that, if relevant to the work you are doing, might save you time.

Message Descriptor properties

If you manipulate message descriptor properties in a program, it can be better to use the hexadecimal equivalents of the fields.

The information in this section refers to the following properties:

- AccountingToken
- CorrelationId
- GroupId
- MessageId

Where a WebSphere MQ application is the originator of a message and WebSphere MQ generates these properties, it is better to use the AccountingTokenHex, CorrelationIdHex, GroupIdHex, and MessageIdHex properties if you want to look at their values, or manipulate them in any way, including passing them back in a message to WebSphere MQ. The reason for this is that WebSphere MQ generated values are strings of bytes that have any value from 0 through 255 inclusive, they are not strings of printable characters.

Where your MQAX script is the originator of a message you can use either the AccountingToken, CorrelationId, GroupId, and MessageId properties or their Hex equivalents.

WebSphere MQ constants

WebSphere MQ constants are provided as members of the enum WebSphere MQ in library MQAX200.

WebSphere MQ string constants

WebSphere MQ string constants and their corresponding character strings.

WebSphere MQ string constants are not available when using WebSphere MQ Automation Classes for ActiveX. You must use the explicit character string for those shown in the following list and any others you might need. The commands must be padded to eight characters using spaces:

MQFMT_NONE	" "
MQFMT_ADMIN	"MQADMIN "
MQFMT_CHANNEL_COMPLETED	"MQCHCOM "
MQFMT_CICS	"MQCICS "
MQFMT_COMMAND_1	"MQCMD1 "
MQFMT_COMMAND_2	"MQCMD2 "
MQFMT_DEAD_LETTER_HEADER	"MQDEAD "
MQFMT_DIST_HEADER	"MQHDIST "
MQFMT_EVENT	"MQEVENT "
MQFMT_IMS	"MQIMS "
MQFMT_IMS_VAR_STRING	"MQIMSVS "
MQFMT_MD_EXTENSION	"MQHMDE "
MQFMT_PCF	"MQPCF "
MQFMT_REF_MSG_HEADER	"MQHREF "
MQFMT_RF_HEADER	"MQHRF "
MQFMT_STRING	"MQSTR "
MQFMT_TRIGGER	"MQTRIG "
MQFMT_WORK_INFO_HEADER	"MQHWIH "
MQFMT_XMIT_Q_HEADER	"MQXMIT "

Null string constants

The WebSphere MQ constants, used for the initialization of four MQMessage properties, MQMI_NONE (24 NULL characters), MQCI_NONE (24 NULL characters), MQGI_NONE (24 NULL characters), and MQACT_NONE (32 NULL characters), are not supported by WebSphere MQ Automation Classes for ActiveX. Setting them to empty strings has the same effect.

For example, to set the various IDs of an MQMessage to these values: *mymessage*.**MessageId** = ""
mymessage.**CorrelationId** = "" *mymessage*.**AccountingToken** = ""

Receiving a message from WebSphere MQ

There are several ways of receiving a message from WebSphere MQ:

- Polling by issuing a GET followed by a Wait, using the Visual Basic TIMER function.

- Issuing a GET with the Wait option; you specify the wait duration by setting the WaitInterval property. Consider this when, even though you set your system up to run in multithreaded environment, the software running at the time might run only single threaded. This avoids your system locking up indefinitely.

Other threads operate unaffected. However, if your other threads require access to WebSphere MQ, they require a second connection to WebSphere MQ using additional MQAX queue manager and queue objects.

Issuing a GET with the Wait option and setting the WaitInterval to MQWI_UNLIMITED causes your system to lock up until the GET call completes, if the process is single threaded.

Using data conversion

Two forms of data conversion are supported by WebSphere MQ Automation Classes for ActiveX - numeric encoding, and character set conversion.

Numeric encoding

If you set the MQMessage Encoding property, the following methods convert between different numeric encoding systems:

- ReadDecimal2 method
- ReadDecimal4 method
- ReadDouble method
- ReadDouble4 method
- ReadFloat method
- ReadInt2 method
- ReadInt4 method
- ReadLong method
- ReadShort method
- ReadUInt2 method
- WriteDecimal2 method
- WriteDecimal4 method
- WriteDouble method
- WriteDouble4 method
- WriteFloat method
- WriteInt2 method
- WriteInt4 method
- WriteLong method
- WriteShort method
- WriteUInt2 method

The Encoding property can be set and interpreted using the supplied WebSphere MQ constants. [Figure 216 on page 991](#) shows an example of these:

```

/* Encodings for Binary Integers */
MQENC_INTEGER_UNDEFINED
MQENC_INTEGER_NORMAL
MQENC_INTEGER_REVERSED

/* Encodings for Decimals */
MQENC_DECIMAL_UNDEFINED
MQENC_DECIMAL_NORMAL
MQENC_DECIMAL_REVERSED

/* Encodings for Floating-Point Numbers */
MQENC_FLOAT_UNDEFINED
MQENC_FLOAT_IEEE_NORMAL
MQENC_FLOAT_IEEE_REVERSED
MQENC_FLOAT_S390

```

Figure 216. Supplied WebSphere MQ constants for encoding

For example, to send an integer from an Intel system to a System/390® operating system in System/390 encoding:

```

Dim msg As New MQMessage 'Define a WebSphere MQ message for our use..
Print msg.Encoding 'Currently 546 (or X'222')
                        'Set the encoding property
                        to 785 (or X'311')
msg.Encoding = MQENC_INTEGER_NORMAL OR MQENC_DECIMAL_NORMAL
                OR MQENC_FLOAT_S390
Print msg.Encoding 'Print it to see the change
Dim local_num As long 'Define a long integer
local_num = 1234 'Set it
msg.WriteLong(local_num) 'Write the number into the message

```

Character set conversion

Character set conversion is necessary when you send a message from one system to another system where the code pages are different. Code page conversion is used by:

- ReadString method
- ReadNullTerminatedString method
- WriteString method
- WriteNullTerminatedString method
- MessageData Property

You must set the MQMessage CharacterSet property to a supported character set value (CCSID).

WebSphere MQ Automation Classes for ActiveX uses conversion tables to perform character set conversion.

For example, to convert strings automatically to code page 437:

```

Dim msg As New MQMessage 'Define a WebSphere MQ message
msg.CharacterSet = 437 'Set code page required
msg.WriteString "A character string" 'Put character string in message

```

The WriteString method receives the string data ("A character string" in the example) as a Unicode string. It then converts this data from Unicode into code page 437 using the conversion table 34B001B5.TBL.

Characters in the Unicode string that are not supported by code page 437 are given the standard substitution character from code page 437.

Similarly, when you use the ReadString method, the incoming message has a character set established by the WebSphere MQ Message Descriptor (MQMD) value and there is a conversion from this code page into Unicode before it is passed back to your scripting language.

Threading

WebSphere MQ Automation Classes for ActiveX implement a free-threading model where objects can be used between threads.

While MQAX permits the use of MQQueue and MQQueueManager objects, WebSphere MQ does not currently permit the sharing of handles between different threads.

Attempts to use these on another thread result in an error and WebSphere MQ returns a return code of MQRC_HCONN_ERROR.

Note: There is only one MQSession object per process. Using the MQSession CompletionCode and ReasonCode is not recommended in multithreaded environments. The MQSession error values might be overwritten by a second thread between an error being raised and checked on the first thread. Threads are serialized for the duration of each method call or property access. So, issuing a Get with the Wait option causes other threads accessing MQAX objects to be suspended until the operation completes.

Error handling

This information describes MQAX object properties, how error handling works, rules describing how raising exceptions are handled, and getting a property.

Each MQAX object includes properties to hold error information and a method to reset or clear them. The properties are:

- CompletionCode
- ReasonCode
- ReasonName

The method is:

- ClearErrorCodes

How error handling works

Your MQAX script or application invokes an MQAX object's method, or accesses or updates a property of the MQAX object:

1. The ReasonCode and CompletionCode in the object concerned are updated.
2. The ReasonCode and CompletionCode in the MQSession object are also updated with the same information.

Note: See [“Threading” on page 992](#) for restrictions on the use of MQSession error codes in threaded applications.

If the CompletionCode is greater than or equal to the ExceptionThreshold property of MQSession, MQAX throws an exception (number 32000). Use this within your script using the On Error (or equivalent) statement to process it.

3. Use the Error function to retrieve the associated error string, which has the form:

```
MQAX: CompletionCode=xxx, ReasonCode=xxx, ReasonName=xxx
```

For more information about how to use the On Error statements, see the documentation for your ActiveX scripting language.

Using the CompletionCode and ReasonCode in the MQSession object is convenient for simple error handlers.

ReasonName property returns the WebSphere MQ symbolic name for the current value of the ReasonCode.

Raising exceptions

The following rules describe how raising exceptions are handled:

- Whenever a property or method sets the completion code to a value greater than or equal to the exception threshold (typically set to 2) an exception is raised.
- All method calls and property sets set the completion code.

Getting a property

This is a special case because the CompletionCode and ReasonCode are not always updated:

- If a property get succeeds, the object and MQSession object ReasonCode and CompletionCode remain unchanged.
- If a property get fails with a CompletionCode of warning, the ReasonCode and CompletionCode remain unchanged.
- If a property get fails with a CompletionCode of error, the ReasonCode and CompletionCode are updated to reflect the true values, and error processing proceeds as described.

The MQSession class has a method *ReasonCodeName* which might be used to replace a WebSphere MQ reason code with a symbolic name. This is especially useful while developing programs where unexpected errors might occur. However, the name is not ideal for presentation to users.

Each class also has a property *ReasonName*, that returns the symbolic name of the current reason code for that class.

WebSphere MQ Automation Classes for ActiveX reference

This section describes the classes of the WebSphere MQ Automation Classes for ActiveX (MQAX), developed for ActiveX. The classes enable you to write ActiveX applications that can access other applications running in your non-ActiveX environments, using WebSphere MQ.

WebSphere MQ Automation Classes for ActiveX interface

WebSphere MQ Automation Classes for ActiveX provides predefined numeric ActiveX constants (such as MQMT_REQUEST) needed to use the classes.

The ActiveX automation classes consist of the following:

- “MQSession Class” on page 995
- “MQQueueManager class” on page 998
- “MQQueue class” on page 1009
- “MQMessage class” on page 1023
- “MQPutMessageOptions class” on page 1044
- “MQGetMessageOptions class” on page 1047
- “MQDistributionList class” on page 1049
- “MQDistributionListItem class” on page 1053

In addition, WebSphere MQ Automation Classes for ActiveX provides predefined numeric ActiveX constants (such as MQMT_REQUEST) needed to use the classes. These are provided in the enum MQ in library MQAX200. The constants are a subset of those defined in the WebSphere MQ C header files (cmqc*.h) with some additional WebSphere MQ Automation Classes for ActiveX Reason codes.

About WebSphere MQ Automation Classes for ActiveX classes

Read this information alongside the reference topics under [Developing applications reference](#).

See [Features that can be used only with the primary installation on Windows](#) for important information.

The MQSession class provides a root object that contains the status of the last action performed on any of the MQAX objects. See [“Error handling”](#) on page 992 for more information.

The MQQueueManager and MQQueue classes provide access to the underlying WebSphere MQ objects. Methods or property accesses for these classes in general result in calls being made across the WebSphere MQ MQI.

The MQMessage, MQPutMessageOptions, and MQGetMessageOptions classes encapsulate the MQMD, MQPMO, and MQGMO data structures, and are used to help you send messages to queues and retrieve messages from them.

The MQDistributionList class encapsulates a collection of queues - local, remote, or alias for output. The MQDistributionListItem class encapsulates the MQOR, MQRR, and MQPMR structures and associates them with an owning distribution list.

Parameter passing

Parameters on method invocations are all passed by value, except where that parameter is an object, in which case it is a reference that is passed.

The class definitions provided list the Data Type for each parameter or property. For many ActiveX clients, such as Visual Basic, if the variable used is not of the required type, the value is automatically converted to or from the required type - providing such a conversion is possible. This follows standard rules of the client; MQAX provides no such conversion.

Many of the methods take fixed-length string parameters, or return a fixed-length character string. The conversion rules are as follows:

- If the user supplies a fixed-length string of the wrong length, as an input parameter or as a return value, the value is truncated or padded with trailing spaces as required.
- If the user supplies a variable-length string of the wrong length as an input parameter, the value is truncated or padded with trailing spaces.
- If the user supplies a variable-length string of the wrong length as a return value, the string is adjusted to the required length (because returning a value destroys the previous value in the string anyway).
- Strings provided as input parameters can contain embedded Nulls.

These classes can be found in the MQAX200 library.

Object access methods

These methods do not relate directly to any single WebSphere MQ call. Each of these methods creates an object in which reference information is then held, followed by connecting to or opening a WebSphere MQ object:

When a connection is made to a queue manager, it holds the 'connection handle' attribute generated by WebSphere MQ.

When a queue is opened, it holds the 'object handle' attribute generated by WebSphere MQ.

These WebSphere MQ attributes are not directly available to the MQAX program.

Errors

Syntactic errors on parameter passing can be detected at compile time and run time by the ActiveX client. Errors can be trapped using On Error in Visual Basic.

The WebSphere MQ ActiveX classes all contain two special read-only properties - ReasonCode and CompletionCode. These properties can be read at any time.

An attempt to access any other property, or to issue any method call might generate an error from WebSphere MQ.

If a property set or method invocation succeeds, the ReasonCode of the owning object is set to MQRC_NONE and CompletionCode is set to MQCC_OK.

If the property access or method invocation does not succeed, reason and completion codes are set in these fields.

MQSession Class

This is the root class for WebSphere MQ Automation Classes for ActiveX.

There is always only one MQSession object per ActiveX client process. An attempt to create a second object creates a second reference to the original object.

Creation

New creates a new MQSession object.

Syntax

Dim mqsess As New MQSession Set mqsess = New MQSession

Properties

- [“CompletionCode property” on page 995.](#)
- [“ExceptionThreshold property” on page 996.](#)
- [“ReasonCode property” on page 996.](#)
- [“ReasonName property” on page 996.](#)

Method

- [“AccessGetMessageOptions method” on page 997.](#)
- [“AccessMessage method” on page 997.](#)
- [“AccessPutMessageOptions method” on page 997.](#)
- [“AccessQueueManager method” on page 997.](#)
- [“ClearErrorCodes method” on page 997.](#)
- [“ReasonCodeName method” on page 997.](#)

CompletionCode property

Read-only. Returns the WebSphere MQ completion code set by the most recent method or property set issued against any WebSphere MQ object.

It is reset to MQCC_OK when a method or a property set is invoked successfully against any MQAX object.

An error event handler can inspect this property to diagnose the error, without having to know which object was involved.

Using the CompletionCode and ReasonCode in the MQSession object is very convenient for simple error handlers.

Note: See [“Threading” on page 992](#) for restrictions on the use of MQSession error codes in threaded applications.

Defined in: MQSession class

Data Type: Long

Values:

- MQCC_OK

- MQCC_WARNING
- MQCC_FAILED

Syntax:

To get: *completioncode*& = *MQSession.CompletionCode*

ExceptionThreshold property

Read-write. Defines the level of WebSphere MQ error for which MQAX will throw an exception. Defaults to MQCC_FAILED. A value greater than MQCC_FAILED effectively prevents exception processing, leaving the programmer to perform checks on the CompletionCode and ReasonCode.

Defined in: MQSession class

Data Type: Long

Values:

- Any, but consider MQCC_WARNING, MQCC_FAILED or greater.

Syntax:

To get: *ExceptionThreshold*& = *MQSession.ExceptionThreshold*

To set: *MQSession.ExceptionThreshold* = *ExceptionThreshold*\$

ReasonCode property

Read-only. Returns the reason code set by the most recent method or property set issued against any WebSphere MQ object.

An error event handler can inspect this property to diagnose the error, without having to know which object was involved.

Using the CompletionCode and ReasonCode in the MQSession object is very convenient for simple error handlers.

Note: See [“Threading” on page 992](#) for restrictions on the use of MQSession error codes in threaded applications.

Defined in: MQSession class

Data Type: Long

Values:

- See [Reason \(MQLONG\)](#) and the additional MQAX values listed under [“Reason codes” on page 1060](#).

Syntax: To get: *reasoncode*& = *MQSession.ReasonCode*

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE".

Note: See [“Threading” on page 992](#) for restrictions on the use of MQSession error codes in threaded applications.

Defined in: MQSession class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname*\$ = *MQSession.ReasonName*

AccessGetMessageOptions method

Creates a new MQGetMessageOptions object.

Defined in: MQSession class

Syntax: *gmo* = *MQSession*.AccessGetMessageOptions()

AccessMessage method

Creates a new MQMessage object.

Defined in: MQSession class

Syntax: *msg* = *MQSession*.AccessMessage()

AccessPutMessageOptions method

Creates a new MQPutMessageOptions object.

Defined in: MQSession class

Syntax: *pmo* = *MQSession*.AccessPutMessageOptions()

AccessQueueManager method

Creates a new MQQueueManager object and connects it to a real queue manager by means of the WebSphere MQ MQI client or WebSphere MQ server. As well as performing a connect, this method also performs an open for the queue manager object.

When both the WebSphere MQ MQI client and WebSphere MQ server are installed on your system, MQAX applications will run against the server by default. To run MQAX against the client, the client bindings library must be specified in the GMQ_MQ_LIB environment variable, for example, set GMQ_MQ_LIB=mqic.dll.

For a client only installation, it is not necessary to set the GMQ_MQ_LIB environment variable. When this variable is not set, WebSphere MQ attempts to load amqzst.dll. If this DLL is not present (as is the case in a client only installation), WebSphere MQ attempts to load mqic.dll.

If successful it sets the MQQueueManager's ConnectionStatus to TRUE.

A queue manager can be connected to at most one MQQueueManager object per ActiveX instance.

If the connection to the queue manager fails, an error event is raised, and the MQSession object's ReasonCode and CompletionCode are set.

Defined in: MQSession class

Syntax: *set qm* = *MQSession*.AccessQueueManager (*Name*\$)

Parameter: *Name*\$ String. Name of Queue Manager to be connected to.

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE.

Defined in: MQSession class

Syntax: Call *MQSession*.ClearErrorCodes()

ReasonCodeName method

Returns the name of the reason code with the given numeric value. It is useful to give clearer indications of error conditions to users. The name is still somewhat cryptic (for example, ReasonCodeName(2059) is **MQRC_Q_MGR_NOT_AVAILABLE**), so where possible errors should be caught and replaced with descriptive text appropriate to the application.

Defined in: MQSession class

Syntax: *errname\$* = *MQSession*.ReasonCodeName(*reasonCode*&)

Parameter: *reasoncode*& Long. The reason code for which the symbolic name is required.

MQQueueManager class

This class represents a connection to a queue manager. The queue manager can be running locally (a WebSphere MQ server) or remotely with access provided by the WebSphere MQ client. An application must create an object of this class and connect it to a queue manager. When an object of this class is destroyed it is automatically disconnected from its queue manager.

Containment

MQQueue class objects are associated with this class.

New creates a new MQQueueManager object and sets all the properties to initial values. Alternatively use the AccessQueueManager method of the MQSession class.

Creation

New creates a **new** MQQueueManager object and sets all the properties to initial values. Alternatively use the AccessQueueManager method of the MQSession class.

Syntax

Dim mgr As New MQQueueManager set mgr = New MQQueueManager

Properties

- [“AlternateUserId property” on page 999.](#)
- [“AuthorityEvent property” on page 1000.](#)
- [“BeginOptions property” on page 1000.](#)
- [“ChannelAutoDefinition property” on page 1000.](#)
- [“ChannelAutoDefinitionEvent property” on page 1000.](#)
- [“ChannelAutoDefinitionExit property” on page 1001.](#)
- [“CharacterSet property” on page 1001.](#)
- [“CloseOptions property” on page 1001.](#)
- [“CommandInputQueueName property” on page 1001.](#)
- [“CommandLevel property” on page 1001.](#)
- [“CompletionCode property” on page 1002.](#)
- [“ConnectionHandle property” on page 1002.](#)
- [“ConnectionStatus property” on page 1002.](#)
- [“ConnectOptions property” on page 1002.](#)
- [“DeadLetterQueueName property” on page 1003.](#)
- [“DefaultTransmissionQueueName property” on page 1003.](#)
- [“Description property” on page 1003.](#)
- [“DistributionLists property” on page 1003.](#)
- [“InhibitEvent property” on page 1003.](#)
- [“IsConnected property” on page 1003.](#)
- [“IsOpen property” on page 1004.](#)
- [“LocalEvent property” on page 1004.](#)

- [“MaximumHandles property” on page 1004.](#)
- [“MaximumMessageLength property” on page 1004.](#)
- [“MaximumPriority property” on page 1004.](#)
- [“MaximumUncommittedMessages property” on page 1005.](#)
- [“Name property” on page 1005.](#)
- [“ObjectHandle property” on page 1005.](#)
- [“PerformanceEvent property” on page 1005.](#)
- [“Platform property” on page 1005.](#)
- [“ReasonCode property” on page 1006.](#)
- [“ReasonName property” on page 1006.](#)
- [“RemoteEvent property” on page 1006.](#)
- [“StartStopEvent property” on page 1006.](#)
- [“SyncPointAvailability property” on page 1006.](#)
- [“TriggerInterval property” on page 1007.](#)

Methods

- [“AccessQueue method” on page 1007.](#)
- [“AddDistributionList method” on page 1007.](#)
- [“Backout method” on page 1008.](#)
- [“Begin method” on page 1008.](#)
- [“ClearErrorCodes method” on page 1008.](#)
- [“Commit method” on page 1008.](#)
- [“Connect method” on page 1008.](#)
- [“Disconnect method” on page 1008.](#)

Property Access

The following properties can be accessed at any time .

- [“AlternateUserId property” on page 999.](#)
- [“CompletionCode property” on page 1002.](#)
- [“ConnectionStatus property” on page 1002.](#)
- [“ReasonCode property” on page 1006.](#)

The remaining properties can be accessed only if the object is connected to a queue manager, and the user ID is authorized to inquire against that queue manager. If an alternate user ID is set and the current user ID is authorized to use it, the alternate user ID is checked for authorization for inquire instead

If these conditions do not apply, WebSphere MQ Automation Classes for ActiveX attempts to connect to the queue manager and open it for inquire automatically. If this is unsuccessful, the call sets a CompletionCode of MQCC_FAILED and one of the following ReasonCodes:

- MQRC_CONNECTION_BROKEN
- MQRC_NOT_AUTHORIZED
- MQRC_Q_MGR_NAME_ERROR
- MQRC_Q_MGR_NOT_AVAILABLE

AlternateUserId property

Read-write. The alternate user ID to be used to validate access to the queue manager attributes.

This property must not be set if IsConnected is TRUE.

This property cannot be set while the object is open.

Defined in: MQQueueManager class

Data Type: String of 12 characters

Syntax: To get: *altuser\$* = *MQQueueManager.AlternateUserId* To set:
MQQueueManager.AlternateUserId = *altuser\$*

AuthorityEvent property

Read-only. The MQI AuthorityEvent attribute.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *authevent* = *MQQueueManager.AuthorityEvent*

BeginOptions property

Read-write. These are the options that apply to the Begin method. Initially MQBO_NONE.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQBO_NONE

Syntax: To get: *beginoptions*&=*MQQueueManager.BeginOptions*

To set: *MQQueueManager.BeginOptions*=*beginoptions*&

ChannelAutoDefinition property

Read-only. This controls whether automatic channel definition is permitted.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQCHAD_DISABLED
- MQCHAD_ENABLED

Syntax: To get: *channelautodef*&= *MQQueueManager.ChannelAutoDefinition*

ChannelAutoDefinitionEvent property

Read-only. This controls whether automatic channel definition events are generated.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *channelautodefevent* = *MQQueueManager.ChannelAutoDefinitionEvent*

ChannelAutoDefinitionExit property

Read-only. The name of the user exit used for automatic channel definition.

Defined in:

MQQueueManager class

Data Type:

String

Syntax: To get: *channelautodefexit* = *MQQueueManager.ChannelAutoDefinitionExit*

CharacterSet property

Read-only. The MQI CodedCharSetId attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *characterset* = *MQQueueManager.CharacterSet*

CloseOptions property

Read-write. Options used to control what happens when the queue manager is closed. The initial value is MQCO_NONE.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQCO_NONE

Syntax: To get: *closeopt* = *MQQueueManager.CloseOptions*

To set: *MQQueueManager.CloseOptions* = *closeopt*

CommandInputQueueName property

Read-only. The MQI CommandInputQName attribute.

Defined in: MQQueueManager class

Data Type: String of 48 characters

Syntax: To get: *commandinputqname* = *MQQueueManager.CommandInputQueueName*

CommandLevel property

Read-only. Returns the version and level of the WebSphere MQ queue manager implementation (MQI CommandLevel attribute)

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *level*& = *MQQueueManager.CommandLevel*

CompletionCode property

Read-only. Returns the completion code set by the last method or property access issued against the object.

Defined in: *MQQueueManager* class

Data Type: Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode*& = *MQQueueManager.CompletionCode*

ConnectionHandle property

Read-only. The connection handle for the WebSphere MQ queue manager object.

Defined in:

MQQueueManager class

Data Type:

Long

Syntax: To get: *hconn*& = *MQQueueManager.ConnectionHandle*

ConnectionStatus property

Read-only. Indicates if the object is connected to its queue manager or not.

Defined in: *MQQueueManager* class

Data Type: Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *status* = *MQQueueManager.ConnectionStatus*

ConnectOptions property

Read-Write. These options apply to the Connect method. Initially MQCNO_NONE.

Defined in:

MQQueueManager class

Data Type:

Long

Values:

- MQCNO_STANDARD_BINDING
- MQCNO_FASTPATH_BINDING
- MQCNO_NONE

Syntax: To get: *connectoptions*&=*MQQueueManager.ConnectOptions*

To set: *MQQueueManager.ConnectOptions*=*connectoptions*&

DeadLetterQueueName property

Read-only. The MQI DeadLetterQName attribute.

Defined in: MQQueueManager class

Data Type: String of 48 characters

Syntax: to get: *dlqname\$* = *MQQueueManager.DeadLetterQueueName*

DefaultTransmissionQueueName property

Read-only. The MQI DefXmitQName attribute.

Defined in: MQQueueManager class

Data Type: String of 48 characters

Syntax: To get: *defxmitqname\$* = *MQQueueManager.DefaultTransmissionQueueName*

Description property

Read-only. The MQI QMgrDesc attribute.

Defined in: MQQueueManager class

Data Type: String of 64 characters

Syntax: To get: *description\$* = *MQQueueManager.Description*

DistributionLists property

Read-Only. This is the capability of the queue manager to support distribution lists.

Defined in:

MQQueueManager class

Data Type:

Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *distributionlists* = *MQQueueManager.DistributionLists*

InhibitEvent property

Read-only. The MQI InhibitEvent attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *inhibevent* = *MQQueueManager.InhibitEvent*

IsConnected property

A value that indicates whether the queue manager is currently connected.

Read-only.

Defined in: MQQueueManager class

Data Type: Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *isconnected* = *MQQueueManager.IsConnected*

IsOpen property

A value that indicates whether the queue manager is currently open for inquire.

Read-only.

Defined in:

MQQueueManager class

Data Type:

Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *IsOpen* = *MQQueueManager.IsOpen*

LocalEvent property

Read-only. The MQI LocalEvent attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *localevent*& = *MQQueueManager.LocalEvent*

MaximumHandles property

Read-only. The MQI MaxHandles attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *maxhandles*& = *MQQueueManager.MaximumHandles*

MaximumMessageLength property

Read-only. The MQI MaxMsgLength Queue Manager attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *maxmessagelength*& = *MQQueueManager.MaximumMessageLength*

MaximumPriority property

Read-only. The MQI MaxPriority attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *maxpriority*& = *MQQueueManager.MaximumPriority*

MaximumUncommittedMessages property

Read-only. The MQI MaxUncommittedMsgs attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *maxuncommitted&* = *MQQueueManager.MaximumUncommittedMessages*

Name property

Read-write. The MQI QMgrName attribute. This property cannot be written once the MQQueueManager is connected.

Defined in: MQQueueManager class

Data Type: String of 48 characters

Syntax: To get: *name\$* = *MQQueueManager.name*

To set: *MQQueueManager.name* = *name\$*

Note: Visual Basic reserves the "Name" property for use in the visual interface. Therefore, when using within Visual Basic use lower-case, that is, "name".

ObjectHandle property

Read-only. The object handle for the WebSphere MQ queue manager object.

Defined in:

MQQueueManager class

Data type

Long

Syntax: To get: *hobj&* = *MQQueueManager.ObjectHandle*

PerformanceEvent property

Read-only. The MQI PerformanceEvent attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *perfevent&* = *MQQueueManager.PerformanceEvent*

Platform property

Read-only. The MQI Platform attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQPL_WINDOWS_NT
- MQPL_WINDOWS

Syntax: To get: *platform&* = *MQQueueManager.Platform*

ReasonCode property

Read-only. Returns the reason code set by the last method or property access issued against the object.

Defined in: MQQueueManager class

Data Type: Long

Values:

- See [API reason codes](#).

Syntax: To get: *reasoncode*& = *MQQueueManager*.ReasonCode

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE".

Defined in: MQQueueManager class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname*\$ = *MQQueueManager*.ReasonName

RemoteEvent property

Read-only. The MQI RemoteEvent attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *remoteevent*& = *MQQueueManager*.RemoteEvent

StartStopEvent property

Read-only. The MQI StartStopEvent attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *strstpevent*& = *MQQueueManager*.StartStopEvent

SyncPointAvailability property

Read-only. The MQI SyncPoint attribute.

Defined in: MQQueueManager class

Data Type: Long

Values:

- MQSP_AVAILABLE
- MQSP_NOT_AVAILABLE

Syntax: To get: *syncpointavailability* = *MQQueueManager.SyncPointAvailability*

TriggerInterval property

Read-only. The MQI TriggerInterval attribute.

Defined in: MQQueueManager class

Data Type: Long

Syntax: To get: *trigint* = *MQQueueManager.TriggerInterval*

AccessQueue method

Creates an MQQueue object and associates it with this MQQueueManager object by setting the connection reference property of the queue. It sets the Name, OpenOptions, DynamicQueueName, and AlternateUserId properties of the MQQueue object to the values provided, and attempts to open it.

If the open is unsuccessful the call fails. An error event is raised against the object. The ReasonCode and CompletionCode, and the MQSession ReasonCode and CompletionCode of the object are set.

The DynamicQueueName, QueueManagerName, and AlternateUserId parameters are optional and default to "".

The OpenOption MQOO_INQUIRE should be specified in addition to other options if queue properties are to be read.

Do not set the QueueManagerName, or set it to "" if the queue to be opened is local. Otherwise, set it to the name of the remote queue manager that owns the queue, and an attempt is made to open a local definition of the remote queue. For more information about remote queue name resolution and queue manager aliasing, see [What are aliases?](#).

If the Name property is set to a model queue name, specify the name of the dynamic queue to be created in the DynamicQueueName\$ parameter. If the value provided in the DynamicQueueName\$ parameter is "", the value set into the queue object and used on the open call is "AMQ.*". See ["Creating dynamic queues"](#) on page 214 for more information about naming dynamic queues.

Definition

Defined in: MQQueueManager class.

Syntax

Syntax: set queue = MQQueueManager.**AccessQueue**(Name\$, OpenOptions&, QueueManagerName\$, DynamicQueueName\$, AlternateUserId\$)

Parameters

Name\$ String. Name of WebSphere MQ queue.

OpenOptions: Long. Options to be used when queue is opened. See [OpenOptions \(MQLONG\)](#).

QueueManagerName\$ String. Name of the queue manager that owns the queue to be opened. A value of "" implies the queue manager is local.

DynamicQueueName\$ String. The name assigned to the dynamic queue at the time the queue is opened when the Name\$ parameter specifies a model queue.

AlternateUserId\$ String. The alternate user ID used to validate access when opening the queue.

AddDistributionList method

Creates a new MQDistributionList object and sets its connection reference to the owning queue manager.

Defined in:

MQQueueManager class

Syntax: *set distributionlist* = **MQQueueManager.AddDistributionList**

Backout method

Backs out any uncommitted message puts and gets that have occurred as part of a unit of work since the last syncpoint.

Defined in: MQQueueManager class

Syntax: Call *MQQueueManager.Backout()*

Begin method

Begins a unit of work that is coordinated by the queue manager. The begin options affect the behavior of this method.

Defined in:

MQQueueManager class

Syntax: Call *MQQueueManager.Begin()*

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQQueueManager class and the MQSession class.

Defined in: MQQueueManager class

Syntax: Call *MQQueueManager.ClearErrorCodes()*

Commit method

Commits any message puts and gets that have occurred as part of a unit of work since the last syncpoint.

Defined in: MQQueueManager class

Syntax: Call *MQQueueManager.Commit()*

Connect method

Connects the MQQueueManager object to a real queue manager via the WebSphere MQ MQI client or server. As well as making the connection, this method also opens the queue manager object so that it can be queried.

Sets IsConnected to TRUE.

A maximum of one MQQueueManager object per ActiveX instance is allowed to connect to a queue manager.

Defined in: MQQueueManager class

Syntax: Call *MQQueueManager.Connect()*

Disconnect method

Disconnects the MQQueueManager object from the queue manager.

Sets IsConnected to FALSE.

All Queue objects associated with the MQQueueManager object are made unusable and cannot be re-opened.

Any uncommitted changes (message puts and gets) are committed.

Defined in: MQQueueManager class

Syntax: Call *MQQueueManager.Disconnect()*

MQQueue class

This class represents access to a WebSphere MQ queue. This connection is provided by an associated MQQueueManager object. When an object of this class is destroyed it is automatically closed.

Containment

The MQQueue class is contained by the MQQueueManager class.

Creation

New creates a new MQQueue object and sets all the properties to initial values. Alternatively, use the AccessQueue method of the MQQueueManager class.

Syntax

```
Dim que As New MQQueue Set que = New MQQueue
```

Properties

- [“AlternateUserId property” on page 1011.](#)
- [“BackoutRequeueName property” on page 1011.](#)
- [“BackoutThreshold property” on page 1012.](#)
- [“BaseQueueName property” on page 1012.](#)
- [“CloseOptions property” on page 1012.](#)
- [“CompletionCode property” on page 1012.](#)
- [“ConnectionReference property” on page 1013.](#)
- [“CreationDateTime property” on page 1013.](#)
- [“CurrentDepth property” on page 1013.](#)
- [“DefaultInputOpenOption property” on page 1013.](#)
- [“DefaultPersistence property” on page 1013.](#)
- [“DefaultPriority property” on page 1013.](#)
- [“DefinitionType property” on page 1014.](#)
- [“DepthHighEvent property” on page 1014.](#)
- [“DepthHighLimit property” on page 1014.](#)
- [“DepthLowEvent property” on page 1014.](#)
- [“DepthLowLimit property” on page 1014.](#)
- [“DepthMaximumEvent property” on page 1015.](#)
- [“DepthHighEvent property” on page 1014.](#)
- [“DepthHighLimit property” on page 1014.](#)
- [“DepthLowEvent property” on page 1014.](#)
- [“DepthLowLimit property” on page 1014.](#)
- [“DepthMaximumEvent property” on page 1015.](#)
- [“Description property” on page 1015.](#)
- [“DynamicQueueName property” on page 1015.](#)
- [“HardenGetBackout property” on page 1015.](#)
- [“InhibitGet property” on page 1015.](#)
- [“InhibitPut property” on page 1016.](#)

- [“InitiationQueueName property” on page 1016.](#)
- [“IsOpen property” on page 1016.](#)
- [“MaximumDepth property” on page 1016.](#)
- [“MaximumMessageLength property” on page 1016.](#)
- [“MessageDeliverySequence property” on page 1017.](#)
- [“ObjectHandle property” on page 1017.](#)
- [“OpenInputCount property” on page 1017.](#)
- [“OpenOptions property” on page 1017.](#)
- [“OpenOutputCount property” on page 1017.](#)
- [“OpenStatus property” on page 1018.](#)
- [“ProcessName property” on page 1018.](#)
- [“QueueManagerName property” on page 1018.](#)
- [“QueueType Property” on page 1018.](#)
- [“ReasonCode property” on page 1018.](#)
- [“ReasonName property” on page 1019.](#)
- [“RemoteQueueManagerName property” on page 1019.](#)
- [“RemoteQueueName property” on page 1019.](#)
- [“ResolvedQueueManagerName property” on page 1019.](#)
- [“ResolvedQueueName property” on page 1019.](#)
- [“RetentionInterval property” on page 1019.](#)
- [“Scope property” on page 1019.](#)
- [“ServiceInterval property” on page 1020.](#)
- [“ServiceIntervalEvent property” on page 1020.](#)
- [“Shareability property” on page 1020.](#)
- [“TransmissionQueueName property” on page 1020.](#)
- [“TriggerControl property” on page 1020.](#)
- [“TriggerData property” on page 1021.](#)
- [“TriggerDepth property” on page 1021.](#)
- [“TriggerMessagePriority property” on page 1021.](#)
- [“TriggerType property” on page 1021.](#)
- [“Usage property” on page 1021.](#)

Methods

- [“ClearErrorCodes method” on page 1022](#)
- [“Close method” on page 1022](#)
- [“Get method” on page 1022](#)
- [“Open method” on page 1023](#)
- [“Put method” on page 1023](#)

Property Access

If the queue object is not connected to a queue manager, you can read the following properties:

- [“CompletionCode property” on page 1012](#)
- [“OpenStatus property” on page 1018](#)

- “ReasonCode property” on page 1018

and you can read and write to:

- “AlternateUserId property” on page 1011
- “CloseOptions property” on page 1012
- “ConnectionReference property” on page 1013
- “Name property” on page 1017
- “OpenOptions property” on page 1017

If the queue object is connected to a queue manager, you can read all the properties.

Queue Attribute properties

Properties not listed in the previous section are all attributes of the underlying WebSphere MQ queue. They can be accessed only if the object is connected to a queue manager, and the user's user ID is authorized for Inquire or Set against that queue. If an alternate user ID is set and the current user ID is authorized to use it, the alternate user ID is checked for authorisation instead.

The property must be an appropriate property for the given QueueType. See [Attributes for queues](#) for more information.

If these conditions do not apply, the property access will set a CompletionCode of MQCC_FAILED and one of the following ReasonCodes:

- MQRC_CONNECTION_BROKEN
- MQRC_NOT_AUTHORIZED
- MQRC_Q_MGR_NAME_ERROR
- MQRC_Q_MGR_NOT_CONNECTED
- MQRC_SELECTOR_NOT_FOR_TYPE (CompletionCode is MQCC_WARNING)

Opening a queue

The only way to create an MQQueue object is by using the MQQueueManager AccessQueue method or by New. An open MQQueue object remains open (OpenStatus=TRUE) until it is closed or deleted or until the creating queue manager object is deleted or connection is lost to the queue manager. The value of the MQQueue CloseOptions property controls the behavior of the close operation that takes place when the MQQueue object is deleted.

The MQQueueManager AccessQueue method opens the queue using the OpenOptions parameter. The MQQueue.Open method opens the queue using the OpenOptions property. WebSphere MQ validates the OpenOptions against the user authorization as part of the open queue process.

AlternateUserId property

Read-write. The alternate user ID used to validate access to the queue when it is opened.

This property cannot be set while the object is open (that is, when IsOpen is TRUE).

Defined in: MQQueue class

Data Type: String of 12 characters

Syntax: To get: *altuser\$* = *MQQueue.AlternateUserId*

To set: *MQQueue.AlternateUserId* = *altuser\$*

BackoutRequeueName property

Read-only. The MQI BackOutRequeueQName attribute.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *backoutrequeuenam\$* = *MQQueue.BackoutRequeueName*

BackoutThreshold property

Read-only. The MQI BackoutThreshold attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- See BackoutThreshold (MQLONG).

Syntax: To get: *backoutthreshold&* = *MQQueue.BackoutThreshold*

BaseQueueName property

Read-only. The queue name to which the alias resolves.

Valid only for alias queues.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *baseqname\$* = *MQQueue.BaseQueueName*

CloseOptions property

Read-Write. Options used to control what happens when the queue is closed.

Defined in: MQQueue class

Data Type: Long

Values:

- MQCO_NONE
- MQCO_DELETE
- MQCO_DELETE_PURGE

MQCO_DELETE and MQCO_DELETE_PURGE are valid only for dynamic queues.

Syntax: To get: *closeopt&* = *MQQueue.CloseOptions*

To set: *MQQueue.CloseOptions* = *closeopt&*

CompletionCode property

Read-only. Returns the completion code set by the last method or property access issued against the object.

Defined in: MQQueue class

Data Type: Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode&* = *MQQueue.CompletionCode*

ConnectionReference property

Read-write. Defines the queue manager object to which a queue object belongs. The connection reference cannot be written while a queue is open.

Defined in: MQQueue class

Data Type: MQQueueManager

Values:

- A reference to an active WebSphere MQ Queue Manager object

Syntax: To set: *set MQQueue.ConnectionReference = ConnectionReference*

To get: *set ConnectionReference = MQQueue.ConnectionReference*

CreationDateTime property

Read-only. Date and time this queue was created.

Defined in: MQQueue class

Data Type: Variant of type 7 (date/time) or EMPTY

Syntax: To get: *datetime = MQQueue.CreationDateTime*

CurrentDepth property

Read-only. The number of messages currently on the queue.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *currentdepth& = MQQueue.CurrentDepth*

DefaultInputOpenOption property

Read-only. Controls the way that the queue is opened if the OpenOptions specify MQOO_INPUT_AS_Q_DEF.

Defined in: MQQueue class

Data Type: Long

Values:

- MQOO_INPUT_EXCLUSIVE
- MQOO_INPUT_SHARED

Syntax: To get: *defaultinop& = MQQueue.DefaultInputOpenOption*

DefaultPersistence property

Read-only. The default persistence for messages on a queue.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *defpersistence& = MQQueue.DefaultPersistence*

DefaultPriority property

Read-only. The default priority for messages on a queue.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *defpriority*& = *MQQueue.DefaultPriority*

DefinitionType property

Read-only. Queue definition type.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQDT_PREDEFINED
- MQQDT_PERMANENT_DYNAMIC
- MQQDT_TEMPORARY_DYNAMIC

Syntax: To get: *deftype*& = *MQQueue.DefinitionType*

DepthHighEvent property

Read-only. The MQI QDepthHighEvent attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *depthhighevent*& = *MQQueue.DepthHighEvent*

DepthHighLimit property

Read-only. The MQI QDepthHighLimit attribute.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *depthhighlimit*& = *MQQueue.DepthHighLimit*

DepthLowEvent property

Read-only. The MQI QDepthLowEvent attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *depthlowevent*& = *MQQueue.DepthLowEvent*

DepthLowLimit property

Read-only. The MQI QDepthLowLimit attribute.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *depthlowlimit*& = *MQQueue.DepthLowLimit*

DepthMaximumEvent property

Read-only. The MQI QDepthMaxEvent attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQEVR_DISABLED
- MQEVR_ENABLED

Syntax: To get: *depthmaximevent*& = *MQQueue.DepthMaximumEvent*

Description property

Read-only. A description of the queue.

Defined in: MQQueue class

Data Type: String of 64 characters

Syntax: To get: *description*\$ = *MQQueue.Description*

DynamicQueueName property

Read-write, read-only when the queue is open.

This controls the dynamic queue name used when a model queue is opened. It can be set with a wildcard by the user either as a property set (only when the queue is closed) or as a parameter to *MQQueueManager.AccessQueue()*.

The actual name of the dynamic queue is found by querying *QueueName*.

Defined in: MQQueue class

Data Type: String of 48 characters

Values:

- Any valid WebSphere MQ queue name.

Syntax: To set: *MQQueue.DynamicQueueName* = *dynamicqueuename*\$

To get: *dynamicqueuename*\$ = *MQQueue.DynamicQueueName*

HardenGetBackout property

Read-only. Whether to maintain an accurate back-out count.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQA_BACKOUT_HARDENED
- MQQA_BACKOUT_NOT HARDENED

Syntax: To get: *hardengetback*& = *MQQueue.HardenGetBackout*

InhibitGet property

Read-write. The MQI InhibitGet attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQA_GET_INHIBITED
- MQQA_GET_ALLOWED

Syntax: To get: *getstatus&* = *MQQueue.InhibitGet*

To set: *MQQueue.InhibitGet* = *getstatus&*

InhibitPut property

Read-write. The MQI InhibitPut attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQA_PUT_INHIBITED
- MQQA_PUT_ALLOWED

Syntax: To get: *putstatus&* = *MQQueue.InhibitPut*

To set: *MQQueue.InhibitPut* = *putstatus&*

InitiationQueueName property

Read-only. Name of initiation queue.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *initqname\$* = *MQQueue.InitiationQueueName*

IsOpen property

Returns whether the queue is open.

Read-only.

Defined in: MQQueue class

Data Type: Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *open* = *MQQueue.IsOpen*

MaximumDepth property

Read-only. Maximum queue depth.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *maxdepth&* = *MQQueue.MaximumDepth*

MaximumMessageLength property

Read-only. Maximum permitted message length in bytes for this queue.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *maxmlength&* = *MQQueue.MaximumMessageLength*

MessageDeliverySequence property

Read-only. Message delivery sequence.

Defined in: MQQueue class

Data Type: Long

Values:

- MQMDS_PRIORITY
- MQMDS_FIFO

Syntax: To get: *messdelseq*& = *MQQueue.MessageDeliverySequence*

Name property

Read-write. The MQI Queue attribute. This property cannot be written after the MQQueue is open.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *name*\$ = *MQQueue.name*

To set: *MQQueue.name* = *name*\$

Note: Visual Basic reserves the "Name" property for use in the visual interface. Therefore, when using within Visual Basic use lower-case, that is "name".

ObjectHandle property

Read-only. The object handle for the WebSphere MQ queue object.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *hobj*& = *MQQueue.ObjectHandle*

OpenInputCount property

Read-only. Number of opens for input.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *openincount*& = *MQQueue.OpenInputCount*

OpenOptions property

Read-write. Options to be used for opening the queue.

Defined in: MQQueue class

Data Type: Long

Values:

- See [OpenOptions \(MQLONG\)](#).

Syntax: To get: *openopt*& = *MQQueue.OpenOptions*

To set: *MQQueue.OpenOptions* = *openopt*&

OpenOutputCount property

Read-only. Number of opens for output.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *openoutcount*& = *MQQueue.OpenOutputCount*

OpenStatus property

Read-only. Indicates if the queue is opened or not. Initial value is TRUE after AccessQueue method or FALSE after New.

Defined in: MQQueue class

Data Type: Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *status*& = *MQQueue.OpenStatus*

ProcessName property

Read-only. The MQI ProcessName attribute.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *procname*\$ = *MQQueue.ProcessName*

QueueManagerName property

Read-write. The WebSphere MQ queue manager name.

Defined in: MQQueue class

Data Type: String

Syntax: To get: *QueueManagerName*\$ = *MQQueue.QueueManagerName*

To set: *MQQueue.QueueManagerName* = *QueueManagerName*\$

QueueType Property

Read-only. The MQI QType attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQT_ALIAS
- MQQT_LOCAL
- MQQT_MODEL
- MQQT_REMOTE

Syntax: To get: *queuetype*& = *MQQueue.QueueType*

ReasonCode property

Read-only. Returns the reason code set by the last method or property access issued against the object.

Defined in: MQQueue class

Data Type: Long

Values:

- See [API reason codes](#).

Syntax: To get: *reasoncode*& = *MQQueue.ReasonCode*

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE".

Defined in: MQQueue class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname*\$ = *MQQueue.ReasonName*

RemoteQueueManagerName property

Read-only. Name of remote queue manager. Valid for remote queues only.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *remqmanname*\$ = *MQQueue.RemoteQueueManagerName*

RemoteQueueName property

Read-only. The name of the queue as it is known on the remote queue manager. Valid for remote queues only.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *remqname*\$ = *MQQueue.RemoteQueueName*

ResolvedQueueManagerName property

Read-only. The name of the final destination queue manager as known to the local queue manager.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *resqmanname*\$ = *MQQueue.ResolvedQueueManagerName*

ResolvedQueueName property

Read-only. The name of the final destination queue as known to the local queue manager.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *resqname*\$ = *MQQueue.ResolvedQueueName*

RetentionInterval property

Read-only. The period of time for which the queue should be retained.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *retinterval*& = *MQQueue.RetentionInterval*

Scope property

Read-only. Controls whether an entry for this queue also exists in a cell directory.

Defined in: MQQueue class

Data Type: Long

Values:

- MQSCO_Q_MGR
- MQSCO_CELL

Syntax: To get: *scope*& = *MQQueue.Scope*

ServiceInterval property

Read-only. The MQI QServiceInterval attribute.

Defined in: MQQueue class

Data Type: Long

Syntax: To get: *serviceinterval*& = *MQQueue.ServiceInterval*

ServiceIntervalEvent property

Read-only. The MQI QServiceIntervalEvent attribute.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQSIE_HIGH
- MQQSIE_OK
- MQQSIE_NONE

Syntax: To get: *serviceintervalevent*& = *MQQueue.ServiceIntervalEvent*

Shareability property

Read-only. Queue shareability.

Defined in: MQQueue class

Data Type: Long

Values:

- MQQA_SHAREABLE
- MQQA_NOT_SHAREABLE

Syntax: To get: *shareability*& = *MQQueue.Shareability*

TransmissionQueueName property

Read-only. Transmission queue name. Valid for remote queues only.

Defined in: MQQueue class

Data Type: String of 48 characters

Syntax: To get: *transqname*\$ = *MQQueue.TransmissionQueueName*

TriggerControl property

Read-write. Trigger control.

Defined in: MQQueue class

Data Type: Long

Values:

- MQTC_OFF
- MQTC_ON

Syntax: To get: *trigcontrol*& = *MQQueue.TriggerControl*

To set: *MQQueue.TriggerControl* = *trigcontrol*&

TriggerData property

Read-write. Trigger data.

Defined in: *MQQueue* class

Data Type: String of 64 characters

Syntax: To get: *trigdata*\$ = *MQQueue.TriggerData*

To set: *MQQueue.TriggerData* = *trigdata*\$

TriggerDepth property

Read-write. The number of messages that have to be on the queue before a trigger message is written.

Defined in: *MQQueue* class

Data Type: Long

Syntax: To get: *trigdepth*& = *MQQueue.TriggerDepth*

To set: *MQQueue.TriggerDepth* = *trigdepth*&

TriggerMessagePriority property

Read-write. Threshold message priority for triggers.

Defined in: *MQQueue* class

Data Type: Long

Syntax: To get: *trigmesspriority*& = *MQQueue.TriggerMessagePriority*

To set: *MQQueue.TriggerMessagePriority* = *trigmesspriority*&

TriggerType property

Read-write. Trigger type.

Defined in: *MQQueue* class

Data Type: Long

Values:

- MQTT_NONE
- MQTT_FIRST
- MQTT EVERY
- MQTT_DEPTH

Syntax: To get: *trigtype*& = *MQQueue.TriggerType*

To set: *MQQueue.TriggerType* = *Trigtype*&

Usage property

Read-only. Indicates what the queue is used for.

Defined in: *MQQueue* class

Data Type: Long

Values:

- MQUS_NORMAL
- MQUS_TRANSMISSION

Syntax: To get: *usage* & = *MQQueue.Usage*

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQQueue class and the MQSession class.

Defined in: MQQueue class

Syntax: Call *MQQueue.ClearErrorCodes()*

Close method

Closes a queue using the current values of CloseOptions.

Defined in: MQQueue class

Syntax: Call *MQQueue.Close()*

Get method

Retrieves a message from the queue.

This method takes an MQMessage object as a parameter, using some of the fields in the MQMD of the object as input parameters. In particular, the MessageId and CorrelId fields are used, so it is important to ensure that these fields are set as required. For more information about these fields, see [MsgId \(MQBYTE24\)](#) and [CorrelId \(MQBYTE24\)](#).

If the method fails then the MQMessage object is unchanged. If it succeeds, the MQMD and Message Data portions of the MQMessage object are replaced by the MQMD and Message Data from the incoming message. The MQMessage control properties are set as follows

- **MessageLength** is set to length of the WebSphere MQ message
- **DataLength** is set to length of the WebSphere MQ message
- **DataOffset** is set to zero

Defined in:

MQQueue class

Syntax: Call *MQQueue.Get(Message, GetMsgOptions, GetMsgLength)*

Parameters

Message:

MQMessage Object representing message to be retrieved.

GetMsgOptions:

Optional MQGetMessageOptions object to control the get operation. If this parameter is not specified, default MQGetMessageOptions are used.

GetMsgLength:

Optional 2 or 4 byte length value to control the maximum length of the WebSphere MQ message that is retrieved from the queue.

If the MQGMO_ACCEPT_TRUNCATED_MSG option is specified, the GET succeeds with a completion code of MQCC_WARNING and a reason code of MQRC_TRUNCATED_MSG_ACCEPTED if the message size exceeds the specified length.

The MessageData holds the first GetMsgLength bytes of data.

If MQGMO_ACCEPT_TRUNCATED_MSG **is not** specified, and the message size exceeds the specified length, the completion code of MQCC_FAILED together with reason code MQRC_TRUNCATED_MESSAGE_FAILED is returned.

If the contents of the message buffer are undefined, the total message length is set to the full length of the message that would have been retrieved.

If the message length parameter is not specified, the length of the message buffer is automatically adjusted to at least the size of the incoming message.

Open method

Opens a queue using the current values of:

1. QueueName
2. QueueManagerName
3. AlternateUserId
4. DynamicQueueName

Defined in:

MQQueue class

Syntax: Call *MQQueue.Open()*

Put method

Places a message onto the queue.

This method takes an MQMessage object as a parameter. The Message Descriptor (MQMD) properties of this object might be altered as a result of this method. The values they have immediately after this method has run are the values that were put onto the WebSphere MQ.

Modifications to the MQMessage object after the Put has completed do not affect the actual message on the WebSphere MQ queue.

Defined in:

MQQueue class

Syntax: Call *MQQueue.Put(Message, PutMsgOptions)*

Parameters

Message

MQMessage object representing message to be put.

PutMsgOptions

MQPutMessageOptions object containing options to control the put operation. If these are not specified, default PutMessageOptions are used.

MQMessage class

This class represents a WebSphere MQ message. It includes properties to encapsulate the WebSphere MQ message descriptor (MQMD), and provides a buffer to hold the application-defined message data.

The class includes Write methods to copy data from an ActiveX application to an MQMessage object. Similarly, the class includes Read methods to copy data from an MQMessage object to an ActiveX application. The class manages the allocation and deallocation of memory for the buffer automatically. The application does not have to declare the size of the buffer when an MQMessage object is created because the buffer grows to accommodate data written to it.

You cannot place a message onto a WebSphere MQ queue if the buffer size exceeds the MaximumMessageLength property of that queue.

After it has been constructed, an MQMessage object can be Put onto a WebSphere MQ queue using the MQQueue.Put method. This method takes a copy of the MQMD and message data portions of the object and places that copy on the queue. The application can therefore modify or delete an MQMessage object after the Put, without affecting the message on the WebSphere MQ queue. The queue manager can adjust some of the fields in the MQMD when it copies the message on the WebSphere MQ queue.

An incoming message can be read into an MQMessage object using the MQQueue.Get method. This replaces any MQMD or message data that might already have been in the MQMessage object with values from the incoming message. It adjusts the size of the data buffer of the MQMessage object to match the size of the incoming message data.

Containment

Messages are contained by the MQSession class.

Creation

New creates an MQMessage object. Its Message Descriptor properties are initially set to default values, and its Message Data buffer is empty.

Syntax

```
Dim msg As New MQMessage or Set msg = New MQMessage
```

Properties

The control properties are:

- [“CompletionCode property” on page 1026](#)
- [“DataLength property” on page 1027](#)
- [“DataOffset property” on page 1027](#)
- [“MessageLength property” on page 1027](#)
- [“ReasonCode property” on page 1027](#)
- [“ReasonName property” on page 1028](#)

The Message Descriptor properties are:

- [“AccountingToken property” on page 1028](#)
- [“AccountingTokenHex property” on page 1028](#)
- [“ApplicationIdData property” on page 1028](#)
- [“ApplicationOriginData property” on page 1029](#)
- [“BackoutCount property” on page 1029](#)
- [“CharacterSet property” on page 1029](#)
- [“CorrelationId property” on page 1029](#)
- [“CorrelationIdHex property” on page 1030](#)
- [“Encoding property” on page 1030](#)
- [“Expiry property” on page 1031](#)
- [“Feedback property” on page 1031](#)
- [“Format property” on page 1031](#)
- [“GroupId property” on page 1031](#)
- [“GroupIdHex property” on page 1032](#)
- [“MessageData property” on page 1032](#)

- [“MessageFlags property” on page 1032](#)
- [“MessageId property” on page 1032](#)
- [“MessageIdHex property” on page 1033](#)
- [“MessageSequenceNumber property” on page 1033](#)
- [“MessageType property” on page 1033](#)
- [“Offset property” on page 1034](#)
- [“OriginalLength property” on page 1034](#)
- [“Persistence property” on page 1034](#)
- [“Priority property” on page 1034](#)
- [“PutApplicationName property” on page 1034](#)
- [“PutApplicationType property” on page 1035](#)
- [“PutDateTime property” on page 1035](#)
- [“ReplyToQueueManagerName property” on page 1035](#)
- [“ReplyToQueueName property” on page 1035](#)
- [“Report property” on page 1035](#)
- [“TotalMessageLength property” on page 1036](#)
- [“UserId property” on page 1036](#)

Methods

- [“ClearErrorCodes method” on page 1036](#)
- [“ClearMessage method” on page 1036](#)
- [“Read method” on page 1036](#)
- [“ReadBoolean method” on page 1037](#)
- [“ReadByte method” on page 1037](#)
- [“ReadDecimal2 method” on page 1037](#)
- [“ReadDecimal4 method” on page 1037](#)
- [“ReadDouble method” on page 1037](#)
- [“ReadDouble4 method” on page 1038](#)
- [“ReadFloat method” on page 1038](#)
- [“ReadInt2 method” on page 1038](#)
- [“ReadInt4 method” on page 1038](#)
- [“ReadLong method” on page 1038](#)
- [“ReadNullTerminatedString method” on page 1039](#)
- [“ReadShort method” on page 1039](#)
- [“ReadString method” on page 1039](#)
- [“ReadUInt2 method” on page 1039](#)
- [“ReadUnsignedByte method” on page 1040](#)
- [“ReadUTF method” on page 1040](#)
- [“ResizeBuffer method” on page 1040](#)
- [“Write method” on page 1040](#)
- [“WriteBoolean method” on page 1041](#)
- [“WriteByte method” on page 1041](#)
- [“WriteDecimal2 method” on page 1041](#)

- [“WriteDecimal4 method” on page 1041](#)
- [“WriteDouble method” on page 1041](#)
- [“WriteDouble4 method” on page 1042](#)
- [“WriteFloat method” on page 1042](#)
- [“WriteInt2 method” on page 1042](#)
- [“WriteInt4 method” on page 1042](#)
- [“WriteLong method” on page 1043](#)
- [“WriteNullTerminatedString method” on page 1043](#)
- [“WriteShort method” on page 1043](#)
- [“WriteString method” on page 1043](#)
- [“WriteUInt2 method” on page 1043](#)
- [“WriteUnsignedByte method” on page 1044](#)
- [“WriteUTF method” on page 1044](#)

Property access

All properties can be read at any time.

The control properties are read-only, except for DataOffset which is read-write. The Message Descriptor properties are all read-write, except BackoutCount and TotalMessageLength which are both read-only.

Note however that some of the MQMD properties might be modified by the queue manager when the message is put onto a WebSphere MQ queue. See the fields in [MQMD](#) for details of how they might be modified.

Data conversion

You can pass binary data to a WebSphere MQ message by setting the CharacterSet property to the Coded Character Set Identifier of the queue manager (MQCCSI_Q_MGR), and passing it a string. You can use the chr\$ function to set non-character data into the string.

The Read and Write methods perform data conversion. They convert between the ActiveX internal formats, and the WebSphere MQ message formats as defined by the Encoding and CharacterSet properties from the message descriptor. When writing a message, set values into Encoding and CharacterSet that match the characteristics of the recipient of the message before issuing a Write method. When reading a message, this step is not normally required because these values will have been set from those in the incoming MQMD.

This is an additional data conversion step that happens after any conversion performed by the MQQueue.Get method.

CompletionCode property

Read-only. Returns the WebSphere MQ completion code set by the most recent method or property access issued against this object.

Defined in: MQMessage class

Data Type: Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode* = MQMessage.CompletionCode

DataLength property

Read-only. This property returns the value:

```
MQMessage.MessageLength - MQMessage.DataOffset
```

It can be used before a Read method, to check that the expected number of characters are actually present in the buffer.

The initial value is zero.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *bytesleft*& = *MQMessage.DataLength*

DataOffset property

Read-write. The current position within the Message Data portion of the message object.

The value is expressed as a byte offset from the start of the message data buffer; the first character in the buffer corresponds to a DataOffset value of zero.

A read or write method commences its operation at the character referenced by DataOffset. These methods process data in the buffer sequentially from this position, and update DataOffset to point to the byte (if any) immediately following the last byte processed.

DataOffset can take only values in the range zero to MessageLength inclusive. When DataOffset = MessageLength it is pointing to the end, that is the first invalid character of the buffer. Write methods are permitted in this situation - they extend the data in the buffer and increase MessageLength by the number of bytes added. Reading beyond the end of the buffer is not valid.

The initial value is zero.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *currpos*& = *MQMessage.DataOffset*

To set: *MQMessage.DataOffset* = *currpos*&

MessageLength property

Read-only. Returns the total length of the Message Data portion of the message object in characters, irrespective of the value of DataOffset.

The initial value is zero. It is set to the incoming Message Length after a Get method invocation that referenced this message object. It is incremented if the application uses a Write method to add data to the object. It is unaffected by Read methods.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *msglength*& = *MQMessage.MessageLength*

ReasonCode property

Read-only. Returns the reason code set by the most recent method or property access issued against this object.

Defined in: MQMessage class

Data Type: Long

Values:

- See [API reason codes](#).

Syntax: To get: *reasoncode*& = *MQMessage.ReasonCode*

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE". **Defined in:** *MQMessage* class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname*\$ = *MQMessage.ReasonName*

AccountingToken property

Read-write. The MQMD AccountingToken - part of the message Identity Context.

Its initial value is all nulls.

Defined in: *MQMessage* class

Data Type: String of 32 characters

Syntax: To get: *actoken*\$ = *MQMessage.AccountingToken*

To set: *MQMessage.AccountingToken* = *actoken*\$

See “Message Descriptor properties” on page 988 for more information about when you must use AccountingTokenHex in place of the AccountingToken property.

AccountingTokenHex property

Read-write. The MQMD AccountingToken - part of the message Identity Context.

Every two characters represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 64 valid hexadecimal characters.

Its initial value is "0...0"

Defined in: *MQMessage* class

Data Type: String of 64 hexadecimal characters representing 32 ASCII characters

Syntax: To get: *actokenh*\$ = *MQMessage.AccountingTokenHex*

To set: *MQMessage.AccountingTokenHex* = *actokenh*\$

See “Message Descriptor properties” on page 988 for more information about when you must use AccountingTokenHex in place of the AccountingToken property.

ApplicationIdData property

Read-write. The MQMD ApplIdentityData - part of the message Identity Context.

Its initial value is all blanks.

Defined in: *MQMessage* class

Data Type: String of 32 characters

Syntax: To get: *applid*\$ = *MQMessage.ApplicationIdData*

To set: *MQMessage.ApplicationIdData* = *applid*\$

ApplicationOriginData property

Read-write. The MQMD ApplOriginData - part of the message origin context.

Its initial value is all blanks.

Defined in: MQMessage class

Data Type: String of 4 characters

Syntax: To get: *applor\$* = *MQMessage.ApplicationOriginData*

To set: *MQMessage.ApplicationOriginData* = *applor\$*

BackoutCount property

Read-only. The MQMD BackoutCount.

Its initial value is 0

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *backoutct&* = *MQMessage.BackoutCount*

CharacterSet property

Read-write. The MQMD CodedCharSetId.

Its initial value is the special value **MQCCSI_Q_MGR**.

If the CharacterSet is set to **MQCCSI_Q_MGR**, the code page for the current locale is used for character conversion in the WriteString method. For server applications, the code page used is the code page of the queue manager. For client applications, it is the default current locale code page.

For example:

```
msg.CharacterSet = MQCCSI_Q_MGR
msg.WriteString(chr$(n))
```

where 'n' is greater than or equal to zero and less than or equal to 255, results in a single byte of value of 'n' being written to the buffer.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *:30ccid&* = *MQMessage.CharacterSet*

To set: *MQMessage.CharacterSet* = *ccid&*

Example

If you want the string written out in code page 437, issue:

```
Message.CharacterSet = 437
Message.WriteString ("string to be written")
```

Set the value you want in the CharacterSet before issuing any WriteString calls.

CorrelationId property

Read-write. The CorrelationId to be included in the MQMD of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Its initial value is null.

Defined in: MQMessage class

Data Type: String of 24 characters

Syntax: To get: *correlid\$ = MQMessage.CorrelationId* To set: *MQMessage.CorrelationId = correlid\$*

See [“Message Descriptor properties”](#) on page 988 for more information about when you must use CorrelationIdHex in place of the CorrelationId property.

CorrelationIdHex property

Read-write. The CorrelationId to be included in the MQMD of a message when put on a queue. Also the CorrelationId to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

Defined in: MQMessage class

Data Type: String of 48 hexadecimal characters representing 24 ASCII characters

Syntax: To get: *correlidh\$ = MQMessage.CorrelationIdHex*

To set: *MQMessage.CorrelationIdHex = correlidh\$*

See [“Message Descriptor properties”](#) on page 988 for a discussion of when you must use CorrelationIdHex in place of the CorrelationId property.

Encoding property

Read-write. The MQMD field that identifies the representation used for numeric values in the application message data.

Its initial value is the special value MQENC_NATIVE, which varies by platform.

This property is used by the following methods:

- ReadDecimal2 method
- ReadDecimal4 method
- ReadDouble method
- ReadDouble4 method
- ReadFloat method
- ReadInt2 method
- ReadInt4 method
- ReadLong method
- ReadShort method
- ReadUInt2 method
- WriteDecimal2 method
- WriteDecimal4 method
- WriteDouble method
- WriteDouble4 method
- WriteFloat method
- WriteInt2 method
- WriteInt4 method
- WriteLong method
- WriteShort method

- WriteUInt2 method

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *encoding*& = MQMessage.**Encoding** To set: MQMessage.**Encoding** = *encoding*&

If you are preparing to write data to the message buffer, you should set this field to match the characteristics of the receiving queue manager platform if the receiving queue manager is incapable of performing its own data conversion.

Expiry property

Read-write. The MQMD expiry time field, expected in tenths of a second.

Its initial value is the special value MQEI_UNLIMITED

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *expiry*& = MQMessage.**Expiry**

To set: MQMessage.**Expiry** = *expiry*&

Feedback property

Read-write. The MQMD feedback field.

Its initial value is the special value MQFB_NONE.

Defined in: MQMessage class

Data Type: Long

Values:

- See [Feedback](#).

Syntax: To get: *feedback*& = MQMessage.**Feedback**

To set: MQMessage.**Feedback** = *feedback*&

Format property

Read-write. The MQMD format field. Gives the name of a built-in or user-defined format that describes the nature of the Message Data.

Its initial value is the special value MQFMT_NONE.

Defined in: MQMessage class

Data Type: String of 8 characters

Syntax: To get: *format*\$ = MQMessage.**Format**

To set: MQMessage.**Format** = *format*\$

GroupId property

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue. Also the ID to be matched against when getting a message from a queue. Its initial value is all nulls.

Defined in:

MQMessage class

Data Type:

String of 24 characters

Syntax: To get: *groupid*\$ = MQMessage.**GroupId**

To set: *MQMessage.GroupId* = *groupid*\$

See [“Message Descriptor properties” on page 988](#) for more information about when you must use *GroupIdHex* in place of the *GroupId* property.

GroupIdHex property

Read-write. The *GroupId* to be included in the MQPMR of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

Defined in:

MQMessage class

Data Type:

String of 48 hexadecimal characters representing 24 ASCII characters.

Syntax: To get: *groupidh*\$ = *MQMessage.GroupIdHex*

To set: *MQMessage.GroupIdHex* = *groupidh*\$

See [“Message Descriptor properties” on page 988](#) for more information about when you must use *GroupIdHex* in place of the *GroupId* property.

MessageData property

Read-write. Retrieves or sets the entire contents of a message as a character string.

Defined in: *MQMessage* class

Data Type: Variant

Note: The data type used by this property is Variant but MQAX expects this to be a variant type of String. If you pass in a variant of other than this type then the error MQRC_OBJECT_TYPE_ERROR will be returned.

Syntax: To get: *String*\$ = *MQMessage.MessageData*

To set: *MQMessage.MessageData* = *String*\$

MessageFlags property

Read-Write. Message flags specifying Segmentation control information. The initial value is 0.

Defined in:

MQMessage class

Data Type:

Long

Values:

See [MsgFlags \(MQLONG\)](#).

Syntax: To get: *messageflags*& = *MQMessage.MessageFlags*

To set: *MQMessage.MessageFlags* = *messageflags*&

MessageId property

Read-write. The *MessageId* to be included in the MQMD of a message when put on a queue. Also the ID to be matched against when getting a message from a queue.

Its initial value is all nulls.

Defined in: MQMessage class

Data Type: String of 24 characters

Syntax: To get: *messageid\$* = *MQMessage.MessageId*

To set: *MQMessage.MessageId* = *messageid\$*

See [“Message Descriptor properties”](#) on page 988 for more information about when you must use MessageIdHex in place of the MessageId property.

MessageIdHex property

Read-write. The MessageId to be included in the MQMD of a message when put on a queue. Also the MessageId to be matched against when getting a message from a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B", and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0...0".

Defined in: MQMessage class

Data Type: String of 48 hexadecimal characters representing 24 ASCII characters

Syntax: To get: *messageidh\$* = *MQMessage.MessageIdHex*

To set: *MQMessage.MessageIdHex* = *messageidh\$*

See [“Message Descriptor properties”](#) on page 988 for more information about when you must use MessageIdHex in place of the MessageId property.

MessageSequenceNumber property

Read-Write. Sequence information identifying a message within a group. The initial value is 1.

Defined in:

MQMessage class

Data Type:

Long

Values:

See [MsgSeqNumber \(MQLONG\)](#).

Syntax: To get: *sequencenumber&* = *MQMessage.SequenceNumber*

To set: *MQMessage.SequenceNumber* = *sequencenumber&*

MessageType property

Read-write. The MQMD MsgType field.

Its initial value is MQMT_DATAGRAM.

Defined in: MQMessage class

Data Type: Long

Values:

- See [MsgType \(MQLONG\)](#).

Syntax: To get: *msgtype&* = *MQMessage.MessageType*

To set: *MQMessage.MessageType* = *msgtype&*

Offset property

Read-Write. The offset in a segmented message. The initial value is 0.

Defined in:

MQMessage class

Data Type:

Long

Values:

See [Offset \(MQLONG\)](#).

Syntax: To get: *offset*& = *MQMessage.Offset*

To set: *MQMessage.Offset* = *offset*&

OriginalLength property

Read-Write. The original length of a segmented message. The initial value is MQOL_UNDEFINED.

Defined in:

MQMessage class

Data Type:

Long

Values:

See [OriginalLength \(MQLONG\)](#).

Syntax: To get: *originallength*& = *MQMessage.OriginalLength*

To set: *MQMessage.OriginalLength* = *originallength*&

Persistence property

Read-write. The message's persistence setting.

Its initial value is MQPER_PERSISTENCE_AS_Q_DEF.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *persist*& = *MQMessage.Persistence*

To set: *MQMessage.Persistence* = *persist*&

Priority property

Read-write. The message's priority.

Its initial value is the special value MQPRI_PRIORITY_AS_Q_DEF

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *priority*& = *MQMessage.Priority*

To set: *MQMessage.Priority* = *priority*&

PutApplicationName property

Read-write. The MQMD PutApplName - part of the Message Origin context.

Its initial value is all blanks.

Defined in: MQMessage class

Data Type: String of 28 characters

Syntax: To get: *putapplnm\$* = *MQMessage.PutApplicationName*

To set: *MQMessage.PutApplicationName* = *putapplnm\$*

PutApplicationType property

Read-write. The MQMD PutApplType - part of the Message Origin context.

Its initial value is MQAT_NO_CONTEXT

Defined in: MQMessage class

Data Type: Long

Values:

- See [PutApplType \(MQLONG\)](#).

Syntax: To get: *putappltp&* = *MQMessage.PutApplicationType*

To set: *MQMessage.PutApplicationType* = *putappltp&*

PutDateTime property

Read/write. This property combines the MQMD PutDate and PutTime fields. These are parts of the Message Origin context that indicate when the message was put.

The ActiveX Extension converts between ActiveX date/time format and the Date and Time formats used in a WebSphere MQ MQMD. If a message is received which has an invalid PutDate or PutTime, then the PutDateTime property after the get method is set to EMPTY.

Its initial value is EMPTY.

Defined in: MQMessage class

Data Type: Variant of type 7 (date/time) or EMPTY.

Syntax: To get: *datetime* = *MQMessage.PutDateTime*

To set: *MQMessage.PutDateTime* = *datetime*

ReplyToQueueManagerName property

Read-write. The MQMD ReplyToQMgr field.

Its initial value is all blanks

Defined in: MQMessage class

Data Type: String of 48 characters

Syntax: To get: *replytoqmgr\$* = *MQMessage.ReplyToQueueManagerName*

To set: *MQMessage.ReplyToQueueManagerName* = *replytoqmgr\$*

ReplyToQueueName property

Read-write. The MQMD ReplyToQ field.

Its initial value is all blanks

Defined in: MQMessage class

Data Type: String of 48 characters

Syntax: To get: *replytoq\$* = *MQMessage.ReplyToQueueName*

To set: *MQMessage.ReplyToQueueName* = *replytoq\$*

Report property

Read-write. The message's Report options.

Its initial value is MQRO_NONE.

Defined in: MQMessage class

Data Type: Long

Values:

- See [Report](#).

Syntax: To get: *report*& = MQMessage.**Report**

To set: MQMessage.**Report** = *report*&

TotalMessageLength property

Read-only. Retrieves the length of the last message received by MQGET. If the message has not been truncated, this value is equal to the value of the MessageLength property.

Defined in: MQMessage class

Data Type: Long

Syntax: To get: *totalmessagelength*& = MQMessage.**TotalMessageLength**

UserId property

Read-write. The MQMD UserIdentifier - part of the message Identity Context.

Its initial value is all blanks.

Defined in: MQMessage class

Data Type: String of 12 characters

Syntax: To get: *userid*\$ = MQMessage.**UserId**

To set: MQMessage.**UserId** = *userid*\$

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQMessage class and the MQSession class.

Defined in: MQMessage class

Syntax: Call MQMessage.**ClearErrorCodes()**

ClearMessage method

This method clears the data buffer portion of the MQMessage object. Any Message Data in the data buffer is lost, because MessageLength, DataLength, and DataOffset are all set to zero.

The Message Descriptor (MQMD) portion is unaffected; an application might need to modify some of the MQMD fields before reusing the MQMessage object. To set the MQMD fields back use New to replace the object with a new instance.

Defined in: MQMessage class

Syntax: Call MQMessage.**ClearMessage()**

Read method

Reads a sequence of bytes from the message buffer into a byte array. DataOffset is incremented and DataLength decremented by the number of bytes read.

Defined in:

MQMessage class

Syntax: Data = MQMessage.**Read**(len&)

Parameters:

len&: Long. Length of data in bytes to be read.

ReadBoolean method

Reads a 1-byte Boolean value from the current position in the message buffer and returns a 2-byte Boolean TRUE(-1)/FALSE(0) value. DataOffset is incremented by one and Data Length is decremented by one.

Defined in:

MQMessage class

Syntax: *value* = MQMessage.**ReadBoolean**

ReadByte method

This method reads 1 byte from the Message Data buffer, starting with the character referred to by DataOffset and returns it as an Integer (signed 2-byte) integer value in the range -128 to 127.

The method fails if MQMessage.DataLength is less than 1 when it is issued.

DataOffset is incremented by 1 and DataLength is decremented by 1 if the method succeeds.

The byte of message data is assumed to be a signed binary integer.

Defined in: MQMessage class

Syntax: *integerv%* = MQMessage.**ReadByte**

ReadDecimal2 method

Reads a 2-byte packed decimal number and returns it as a signed 2-byte integer value. DataOffset is incremented by two and Data Length is decremented by two.

Defined in:

MQMessage class

Syntax: *value%* = MQMessage.**ReadDecimal2**

ReadDecimal4 method

Reads a 4-byte packed decimal number and returns it as a signed 4-byte integer value. DataOffset is incremented by four and Data Length is decremented by four.

Defined in:

MQMessage class

Syntax: Call *value&* = MQMessage.**ReadDecimal4**

ReadDouble method

This method reads 8 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as a Double (signed 8-byte) floating point value.

The method fails if MQMessage.DataLength is less than 8 when it is issued.

DataOffset is incremented by 8 and DataLength is decremented by 8 if the method succeeds.

The 8 characters of message data are assumed to be a binary floating point number. The encoding is specified by the MQMessage.Encoding property. Note that conversion from System/360 format is not supported.

Defined in: MQMessage class

Syntax: *doublev#* = MQMessage.**ReadDouble**

ReadDouble4 method

The ReadDouble4 and WriteDouble4 methods are alternatives to ReadFloat and WriteFloat. This is because they support 4-byte System/390 floating point message values that are too large to convert to 4-byte IEEE floating point format.

This method reads 4 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as a Double (signed 8-byte) floating point value.

The method fails if MQMessage.DataLength is less than 4 when it is issued.

DataOffset is incremented by 4 and DataLength is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a binary floating point number. The encoding is specified by the MQMessage.Encoding property. Note that conversion from System/360 format is not supported.

Defined in: MQMessage class

Syntax: *doublev#* = MQMessage.**ReadDouble4**

ReadFloat method

This method reads 4 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as a Single (signed 4-byte) floating point value.

The method fails if MQMessage.DataLength is less than 4 when it is issued.

DataOffset is incremented by 4 and DataLength is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a floating point number. The encoding is specified by the MQMessage.Encoding property. Note that conversion from System/360 format is not supported.

Defined in: MQMessage class

Syntax: *singlev!* = MQMessage.**ReadFloat**

ReadInt2 method

The method is identical to the ReadShort method.

Syntax: *integerv%* = MQMessage.**ReadInt2**

ReadInt4 method

This method is identical to the ReadLong method.

Syntax: *bigint&* = MQMessage.**ReadInt4**

ReadLong method

This method reads 4 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as a Long (signed 4-byte) integer value.

The method fails if MQMessage.DataLength is less than 4 when it is issued.

DataOffset is incremented by 4 and DataLength is decremented by 4 if the method succeeds.

The 4 characters of message data are assumed to be a binary integer. The encoding is specified by the MQMessage.Encoding property.

Defined in: MQMessage class

Syntax: *bigint&* = MQMessage.**ReadLong**

ReadNullTerminatedString method

This method is for use in place of ReadString if the string might contain embedded null characters.

This method reads the specified number of bytes from the message data buffer starting with the byte referred to by DataOffset and returns it as an ActiveX string. If the string contains an embedded null before the end then the length of the returned string is reduced to reflect only those characters before the null.

DataOffset is incremented and DataLength is decremented by the value specified regardless of whether the string contains embedded null characters.

The characters in the message data are assumed to be a string in the code page that is specified by the MQMessage.CharacterSet property. Conversion to ActiveX representation is performed for the application.

Defined in:

MQMessage class

Syntax: *string\$* = MQMessage.**ReadNullTerminatedString**(length&)

Parameters:

length& Long. Length of string field in bytes.

ReadShort method

This method reads 2 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as an Integer (signed 2-byte) value.

The method fails if MQMessage.DataLength is less than 2 when it is issued.

DataOffset is incremented by 2 and DataLength is decremented by 2 if the method succeeds.

The 2 characters of message data are assumed to be a binary integer. The encoding is specified by the MQMessage.Encoding property.

Defined in: MQMessage class

Syntax: *integerv%* = MQMessage.**ReadShort**

ReadString method

This method reads n bytes from the Message Data buffer starting with the byte referred to by DataOffset and returns it as an ActiveX string.

The method fails if MQMessage.DataLength is less than n when it is issued.

DataOffset is incremented by n and DataLength is decremented by n if the method succeeds.

The n characters of message data are assumed to be a string in the code page that is specified by the MQMessage.CharacterSet property. Conversion to ActiveX representation is performed for the application.

Defined in: MQMessage class

Syntax: *stringv\$* = MQMessage.**ReadString**(length&)

Parameter

length& Long. Length of string field in bytes.

ReadUInt2 method

This method reads 2 bytes from the Message Data buffer, starting with the byte referred to by DataOffset and returns it as a Long (signed 4-byte) integer value.

The method fails if MQMessage.DataLength is less than 2 when it is issued.

DataOffset is incremented by 2 and DataLength is decremented by 2 if the method succeeds.

The 2 bytes of message data are assumed to be an unsigned binary integer. The encoding is specified by the MQMessage.Encoding property.

Defined in: MQMessage class

Syntax: *bigint*& = *MQMessage*.**ReadUInt2**

ReadUnsignedByte method

This method reads 1 byte from the Message Data buffer, starting with the byte referred to by **DataOffset** and returns it as an Integer (signed 2-byte) integer value in the range 0 to 255.

The method fails if *MQMessage.DataLength* is less than 1 when it is issued.

DataOffset is incremented by 1 and **DataLength** is decremented by 1 if the method succeeds.

The 1 character of message data is assumed to be an unsigned binary integer.

Defined in: *MQMessage* class

Syntax: *integerv%* = *MQMessage*.**ReadUnsignedByte**

ReadUTF method

This method reads a UTF format string from the message starting with the byte referred to by **DataOffset** and returns it as an ActiveX string. The string in the message consists of a 2-byte length followed by the character data.

The method fails if *MQMessage.DataLength* is less than the string length when it is issued.

DataOffset is incremented by the string length and **DataLength** is decremented by the string length if the method succeeds.

Defined in:

MQMessage class

Syntax: *value\$* = *MQMessage*.**ReadUTF**

ResizeBuffer method

This method alters the amount of storage currently allocated internally to hold the Message Data buffer. It gives the application some control over the automatic buffer management, in that if the application knows that it is going to deal with a large message, it can ensure that a sufficiently large buffer is allocated. The application does not need to use this call - if it does not, the automatic buffer management code will grow the buffer size to fit.

If you resize the buffer to be smaller than the current **MessageLength**, you risk losing data. If you do lose data, the method returns a **CompletionCode** of **MQCC_WARNING** and a **ReasonCode** of **MQRC_DATA_TRUNCATED**.

If you resize the buffer to be smaller than the value of the **DataOffset** property the:

- **DataOffset** property is changed to point to the end of the new buffer
- **DataLength** property is set to zero
- **MessageLength** property is changed to the new buffer size

Defined in:

MQMessage class

Syntax: *MQMessage*.**ResizeBuffer**(*Length*&)

Parameter:

Length& Long. Size required in characters.

Write method

Writes a sequence of bytes to the message buffer from a byte array at the position referred to by **Data Offset**. If necessary the length of the buffer (*MQMessage.MQMessageLength*) is extended to accommodate the full length of the byte array. **DataOffset** is incremented by the number of bytes written if the method succeeds.

Defined in:

MQMessage class

Syntax: Call *MQMessage*.**Write**(value)

Parameters:

data: a byte array or a variant reference to a byte array

WriteBoolean method

Writes a 1-byte Boolean value at the current position in the message buffer from a 2-byte Boolean value. DataOffset is incremented by one.

Defined in:

MQMessage class

Syntax: Call *MQMessage*.**WriteBoolean**(value)

Parameter:

value: Boolean (2-bytes). Value to be written.

WriteByte method

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 1-byte binary number at the position referred to by DataOffset. It replaces any data already at the position in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by one if the method succeeds.

The value specified should be in the range -128 to 127. If it is not, the method returns with CompletionCode MQCC_FAILED and ReasonCode MQRC_WRITE_VALUE_ERROR.

Defined in: MQMessage class

Syntax: Call *MQMessage*.**WriteByte**(value%)

Parameter: *value%* Integer. Value to be written.

WriteDecimal2 method

Writes a signed 2-byte integer as a 2-byte packed decimal number. DataOffset is incremented by two.

Defined in:

MQMessage class

Syntax: Call *MQMessage*.**WriteDecimal2**(value%)

Parameter:

value% Integer. Value to be written.

WriteDecimal4 method

Writes a signed 4-byte integer as a 4-byte packed decimal number. DataOffset is incremented by four.

Defined in:

MQMessage class

Syntax: Call *MQMessage*.**WriteDecimal4**(value&)

Parameter:

value& Long. Value to be written.

WriteDouble method

This method takes a signed 8-byte floating point value and writes it into the Message Data buffer as an 8-byte floating point number starting at the position referred to by DataOffset. It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 8 if the method succeeds.

The method converts to the floating point representation specified by the MQMessage.Encoding property.
Conversion to System/360 format is not supported.

Defined in: MQMessage class

Syntax: Call *MQMessage*.**WriteDouble**(value#)

Parameter:

value# Double. Value to be written.

WriteDouble4 method

See “ReadDouble4 method” on page 1038 for a description of when ReadDouble4 and WriteDouble4 should be used in place of ReadFloat and WriteFloat.

This method takes a signed 8-byte floating point value and writes it into the Message Data buffer as a 4-byte floating number starting at the position referred to by DataOffset.

DataOffset is incremented by 4 if the method succeeds.

It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

The method converts to the floating point representation specified by the MQMessage.Encoding property.
Conversion to System/360 format is not supported.

Defined in: MQMessage class

Syntax: Call *MQMessage*.**WriteDouble4**(value#)

Parameter: value# Double. Value to be written.

WriteFloat method

This method takes a signed 4-byte floating point value and writes it into the Message Data buffer as a 4-byte floating point number starting at the character referred to by DataOffset. It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 4 if the method succeeds.

The method converts to the binary representation specified by the MQMessage.Encoding property.
Conversion to System/360 format is not supported.

Defined in: MQMessage class

Syntax: Call *MQMessage*.**WriteFloat**(value!)

Parameter value! Float. Value to be written.

WriteInt2 method

This method is identical to the WriteShort method.

Syntax: Call *MQMessage*.**WriteInt2**(value%)

Parameter value% Integer. Value to be written.

WriteInt4 method

This method is identical to the WriteLong method.

Syntax: Call *MQMessage*.**WriteInt4**(value&)

Parameter value& Long. Value to be written.

WriteLong method

This method takes a signed 4-byte integer value and writes it into the Message Data buffer as a 4-byte binary number starting at the byte referred to by DataOffset. It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 4 if the method succeeds.

The method converts to the binary representation specified by the MQMessage.Encoding property.

Defined in: MQMessage class

Syntax: Call *MQMessage.WriteLong*(value&)

Parameter value& Long. Value to be written.

WriteNullTerminatedString method

This method performs a normal WriteString and pads any remaining bytes up to the specified length with null. If the number of bytes written by the initial write string is equal to the specified length then no nulls are written. If the number of bytes exceeds the specified length then an error (reason code MQRC_WRITE_VALUE_ERROR) is set.

DataOffset is incremented by the specified length if the method succeeds.

Defined in: MQMessage class

Syntax: Call *MQMessage.WriteNullTerminatedString*(value\$, length&)

Parameters:

value\$ String. Value to be written.

length& Long. Length of string field in bytes.

WriteShort method

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 2-byte binary number starting at the byte referred to by DataOffset. It replaces any data already at these positions in the buffer, and will extend the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 2 if the method succeeds.

The method converts to the binary representation specified by the MQMessage.Encoding property.

Defined in: MQMessage class

Syntax: Call *MQMessage.WriteShort*(value%)

Parameter value% Integer. Value to be written.

WriteString method

This method takes an ActiveX string and writes it into the Message Data buffer starting at the byte referred to by DataOffset. It replaces any data already at these positions in the buffer, and will extend the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by the length of the string in bytes if the method succeeds.

The method converts characters into the code page specified by the MQMessage.CharacterSet property.

Defined in: MQMessage class

Syntax: Call *MQMessage.WriteString*(value\$)

Parameter value\$ String. Value to be written.

WriteUInt2 method

This method takes a signed 4-byte integer value and writes it into the Message Data buffer as a 2-byte unsigned binary number starting at the byte referred to by DataOffset. It replaces any data already

at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 2 if the method succeeds.

The method converts to the binary representation specified by the MQMessage.Encoding property. The value specified should be in the range 0 to 2**16-1. If it is not the method returns with CompletionCode MQCC_FAILED and ReasonCode MQRC_WRITE_VALUE_ERROR.

Defined in: MQMessage class

Syntax: Call *MQMessage*.WriteUInt2(*value*&)

Parameter *value*& Long. Value to be written.

WriteUnsignedByte method

This method takes a signed 2-byte integer value and writes it into the Message Data buffer as a 1-byte unsigned binary number starting at the character referred to by DataOffset. It replaces any data already at these positions in the buffer, and extends the length of the buffer (MQMessage.MessageLength) if necessary.

DataOffset is incremented by 1 if the method succeeds.

The value specified should be in the range 0 to 255. If it is not the method returns with CompletionCode MQCC_FAILED and ReasonCode MQRC_WRITE_VALUE_ERROR.

Defined in:
MQMessage class

Syntax: Call *MQMessage*.WriteUnsignedByte(*value*%)

Parameter *value*% Integer. Value to be written.

WriteUTF method

This method takes an ActiveX string and writes it into the message data buffer at the current position in UTF format. The data written consists of a 2-byte length followed by the character data. DataOffset is incremented by the length of the string if the method succeeds.

Defined in:
MQMessage class

Syntax: Call *MQMessage*.WriteUTF(*value*\$)

Parameter:
value\$ String. Value to be written.

MQPutMessageOptions class

This class encapsulates the various options that control the action of putting a message onto a WebSphere MQ Queue.

Containment

The MQPutMessageOptions class is contained by the MQSession class.

Creation

New creates a new MQPutMessageOptions object and sets all its properties to initial values.

Alternatively, use the AccessPutMessageOptions method of the MQSession class.

Syntax

Dim *pmo* **As New** MQPutMessageOptions or

Set *pmo* = New MQPutMessageOptions

Properties

- [“CompletionCode property” on page 1045.](#)
- [“Options property” on page 1045.](#)
- [“ReasonCode property” on page 1045.](#)
- [“ReasonName property” on page 1046.](#)
- [“RecordFields property” on page 1046.](#)
- [“ResolvedQueueManagerName property” on page 1046.](#)
- [“ResolvedQueueName property” on page 1046.](#)

Methods

- [“ClearErrorCodes method” on page 1046.](#)

CompletionCode property

Read-only. Returns the completion code set by the last method or property access issued against the object.

Defined in: MQPutMessageOptions class

Data Type: Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode*& = *PutOpts.CompletionCode*

Options property

Read-write. The MQPMO Options field. The initial value of this field is MQPMO_NONE. For more information, see [MQPMO options](#).

Defined in: MQPutMessageOptions Class.

Data Type: Long

Syntax: To get: *options*& = *PutOpts.Options*

To set: *PutOpts.Options* = *options*&

The MQPMO_PASS_IDENTITY_CONTEXT and MQPMO_PASS_ALL_CONTEXT options are not supported.

ReasonCode property

Read-only. Returns the reason code set by the last method or property access issued against the object.

Defined in: MQPutMessageOptions class

Data Type: Long

Values:

- See [API reason codes](#).

Syntax: To get: *reasoncode*& = *PutOpts.ReasonCode*

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE".

Defined in: MQPutMessageOptions class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname\$* = *PutOpts.ReasonName*

RecordFields property

Read-write. Flags indicating which fields are to be customized on a per-queue basis when putting a message to a distribution list. The initial value is zero.

This property corresponds to the PutMsgRecFields flags in the MQI MQPMO structure. In the MQI, these flags control which fields (in the MQPMR structure) are present and used by the MQPUT. In an MQPutMessageOptions object these fields are always present and the flags therefore only affect which fields are used by the Put. See the *WebSphere MQ Application Programming Reference* for further details.

Defined in:

MQPutMessageOptions class

Data Type:

Long

Syntax: To get: *recordfields&* = *PutOpts.RecordFields*

To set: *PutOpts.RecordFields* = *recordfields&*

ResolvedQueueManagerName property

Read-only. The MQPMO ResolvedQMGrName field. See [ResolvedQMGrName \(MQCHAR48\)](#) for details. The initial value is all blanks.

Defined in: MQPutMessageOptions class

Data Type: String of 48 characters

Syntax: To get: *qmgr\$* = *PutOpts.ResolvedQueueManagerName*

ResolvedQueueName property

Read-only. The MQPMO ResolvedQName field. See [ResolvedQName \(MQCHAR48\)](#) for details. The initial value is all blanks.

Defined in: MQPutMessageOptions class

Data Type: String of 48 characters

Syntax: To get: *qname\$* = *PutOpts.ResolvedQueueName*

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQPutMessageOptions class and the MQSession class.

Defined in: MQPutMessageOptions class

Syntax: Call *PutOpts.ClearErrorCodes()*

MQGetMessageOptions class

This class encapsulates the various options that control the action of getting a message from a WebSphere MQ queue.

Containment

The MQGetMessageOptions class is contained by the MQSession class.

Creation

New creates a new MQGetMessageOptions object and sets all its properties to initial values.

Alternatively, use the AccessGetMessageOptions method of the MQSession class.

Properties

- [“CompletionCode property” on page 1047](#)
- [“MatchOptions property” on page 1047](#)
- [“Options property” on page 1048](#)
- [“ReasonCode property” on page 1048](#)
- [“ReasonName property” on page 1048](#)
- [“ResolvedQueueName property” on page 1048](#)
- [“WaitInterval property” on page 1048](#)

Methods

- [“ClearErrorCodes method” on page 1049](#)

Syntax

Dim *gmo* **As New** MQGetMessageOptions or

Set *gmo* = **New** MQGetMessageOptions

CompletionCode property

Read-only. Returns the completion code set by the last method or property access issued against the object.

Defined in: MQGetMessageOptions Class.

Data Type: Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode*& = *GetOpts.CompletionCode*

MatchOptions property

Read-write. Options controlling selection criteria used for MQGET. The initial value is MQMO_MATCH_MSG_ID + MQMO_MATCH_CORREL_ID.

Defined in:

MQGetMessageOptions class

Data Type:

Long

Values:

See [MatchOptions \(MQLONG\)](#).

Syntax: To get: *matchoptions*& = *GetOpts.MatchOptions*

To set: *GetOpts.MatchOptions* = *matchoptions*&

Options property

Read-write. The MQGMO Options field. See [Options](#) for details. Initial value is MQGMO_NO_WAIT.

Defined in: MQGetMessageOptions Class.

Data Type: Long

Syntax: To get: *options*& = *GetOpts.Options* To set: *GetOpts.Options* = *options*&

ReasonCode property

Read-only. Returns the reason code set by the last method or property access issued against the object.

Defined in: MQGetMessageOptions class

Data Type: Long

Values:

- See [API reason codes](#).

Syntax: To get: *reasoncode*& = *GetOpts.ReasonCode*

ReasonName property

Read-only. Returns the symbolic name of the latest reason code. For example, "MQRC_QMGR_NOT_AVAILABLE". **Defined in:** MQGetMessageOptions class

Data Type: String

Values:

- See [API reason codes](#).

Syntax: To get: *reasonname*\$ = *MQGetMessageOptions.ReasonName*

ResolvedQueueName property

Read-only. The MQGMO ResolvedQName field. See [ResolvedQName \(MQCHAR48\)](#) for details. The initial value is all blanks.

Defined in: MQGetMessageOptions class

Data Type: String of 48 characters

Syntax: To get: *qname*\$ = *GetOpts.ResolvedQueueName*

WaitInterval property

Read/write. The MQGMO WaitInterval field. The maximum time, in milliseconds, that the Get waits for a suitable message to arrive - if wait action has been requested by the Options property. This field has an initial value of 0. For details of MQGMO options, see [MQGMO](#).

Defined in: MQGetMessageOptions class

Data Type: Long

Syntax: To get: *wait*& = *GetOpts.WaitInterval*

To set: *GetOpts.WaitInterval* = *wait*&

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQGetMessageOptions class and the MQSession class.

Defined in: MQGetMessageOptions class

Syntax: Call *GetOpts.ClearErrorCodes()*

MQDistributionList class

This class encapsulates a collection of queues - local, remote, or alias for output.

Creation

new creates a new MQDistributionList object.

Alternatively, use the AddDistributionList method of the MQQueueManager class

Properties

- [“AlternateUserId property” on page 1049](#)
- [“CloseOptions property” on page 1050](#)
- [“CompletionCode property” on page 1050](#)
- [“ConnectionReference property” on page 1050](#)
- [“FirstDistributionListItem property” on page 1050](#)
- [“IsOpen property” on page 1050](#)
- [“OpenOptions property” on page 1051](#)
- [“ReasonCode property” on page 1051](#)
- [“ReasonName property” on page 1051](#)

Method

- [“AddDistributionListItem method” on page 1051](#)
- [“ClearErrorCodes method” on page 1052](#)
- [“Close method” on page 1052](#)
- [“Open method” on page 1052](#)
- [“Put method” on page 1052](#)

Syntax

Dim *distlist*.**As New** MQDistributionList or **Set** *distlist* = **New** MQDistributionList

AlternateUserId property

Read-write. The alternate user ID used to validate access to the list of queues when they are opened.

Defined in:

MQDistributionList class

Data Type:

String of 12 characters

Syntax: To get: *altuser\$* = *MQDistributionList.AlternateUserId*

To set: *MQDistributionList.AlternateUserId* = *altuser\$*

CloseOptions property

Read-write. Options used to control what happens when the distribution list is closed. The initial value is MQCO_NONE.

Defined in:

MQDistributionList class

Data Type:

Long

Values:

- MQCO_NONE
- MQCO_DELETE
- MQCO_DELETE_PURGE

Syntax: To get: *closeopt*& = *MQDistributionList.CloseOptions*

To set: *MQDistributionList.CloseOptions* = *closeopt*&

CompletionCode property

Read-only. The completion code set by the last method or property access issued against the object.

Defined in:

MQDistributionList class

Data Type:

Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode*& = *MQDistributionList.CompletionCode*

ConnectionReference property

Read-write. The queue manager to which the distribution list belongs.

Defined in:

MQDistributionList class

Data Type:

MQQueueManager

Syntax: To get: *set queuemanager* = *MQDistributionList.ConnectionReference*

To set: *set MQDistributionList.ConnectionReference* = *queuemanager*

FirstDistributionListItem property

Read-only. The first distribution list item object associated with the distribution list.

Defined in:

MQDistributionList class

Data Type:

MQDistributionListItem

Values:

Syntax: To get: *set distributionlistitem* = *MQDistributionList.FirstDistributionListItem*

IsOpen property

Read-only.

Defined in:

MQDistributionList class

Data Type:

Boolean

Values:

- TRUE (-1)
- FALSE (0)

Syntax: To get: *IsOpen* = *MQDistributionList.IsOpen*

OpenOptions property

Read-write. Options to be used when the distribution list is opened.

Defined in:

MQDistributionList class

Data Type:

Long

Values:

See [MQPMO options](#).

Syntax: To get: *openopt&* = *MQDistributionList.OpenOptions*

To set: *MQDistributionList.OpenOptions* = *openopt&*

ReasonCode property

Read-only. The reason code set by the last method or property access issued against the object.

Defined in:

MQDistributionList class

Data Type:

Long

Values:

See [API reason codes](#).

Syntax: To get: *reasoncode&* = *MQDistributionList.ReasonCode*

ReasonName property

Read-only. The symbolic name for the ReasonCode. For example "MQRC_QMGR_NOT_AVAILABLE".

Defined in:

MQDistributionList class

Data Type:

String

Values:

See [API reason codes](#).

Syntax: To get: *reasonname\$* = *MQDistributionList.ReasonName*

AddDistributionListItem method

Creates a new MQDistributionListItem object and associates it with the distribution list object. The queue name parameter is mandatory.

The DistributionList property of the distribution list item is set to the owning distribution list and the FirstDistributionListItem property of the distribution list is set to reference this new distribution list item.

For the new distribution list item, the PreviousDistributionListItem property is set to nothing and the NextDistributionListItem property is set to reference any distribution list item that was previously first, or nothing if there was none previously (that is, the new one is inserted in front of those that exist already).

This will return an error if the distribution list is open.

Defined in:

MQDistributionList class

Syntax: set distributionlistitem = *MQDistributionList.AddDistributionListItem* (QName\$, QMgrName\$)

Parameters:

QName\$ String. Name of the WebSphere MQ queue.

QMgrName\$ String. Name of the WebSphere MQ queue manager.

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQDistributionList class and the MQSession class.

Defined in:

MQDistributionList class

Syntax: Call *MQDistributionList.ClearErrorCodes*()

Close method

Closes a distribution list using the current value of Close options.

Defined in:

MQDistributionList class

Syntax: Call *MQDistributionList.Close*()

Open method

Opens each of the queues specified by the QueueName and (where appropriate) QueueManagerName properties of the distribution list items associated with the current object using the current value of AlternateUserId.

Defined in:

MQDistributionList class

Syntax: Call *MQDistributionList.Open*()

Put method

Places a message on each of the queues identified by the distribution list items associated with the distribution list.

Defined in:

MQDistributionList class

Syntax

Call MQDistributionList.**Put**(Message, PutMsgOptions&)

Parameters

Message MQMessage object representing the message to be put.

PutMsgOptions MQPutMessageOptions object containing options to control the put operation. If not specified, default PutMessageOptions are used.

This method takes an MQMessage object as a parameter. The following distribution list item properties can be altered as a result of this method:

- CompletionCode
- ReasonCode
- ReasonName
- MessageId
- MessageIdHex
- CorrelationId
- CorrelationIdHex
- GroupId
- GroupIdHex
- Feedback
- AccountingToken
- AccountingTokenHex

MQDistributionListItem class

This class encapsulates the MQOR, MQRR, and MQPMR structures and associates them with an owning distribution list.

Creation

Use the AddDistributionListItem Method of the MQDistributionList class

Properties

Methods

- [“AccountingToken property” on page 1054.](#)
- [“AccountingTokenHex property” on page 1054.](#)
- [“CompletionCode property” on page 1055.](#)
- [“CorrelationId property” on page 1055.](#)
- [“CorrelationIdHex property” on page 1055.](#)
- [“DistributionList property” on page 1055.](#)
- [“Feedback property” on page 1056.](#)
- [“GroupId property” on page 1056.](#)
- [“GroupIdHex property” on page 1056.](#)
- [“MessageId property” on page 1056.](#)
- [“MessageIdHex property” on page 1057.](#)
- [“NextDistributionListItem property” on page 1057.](#)
- [“PreviousDistributionListItem property” on page 1057.](#)
- [“QueueManagerName property” on page 1057.](#)
- [“QueueName property” on page 1057.](#)
- [“ReasonCode property” on page 1058.](#)
- [“ReasonName property” on page 1058.](#)
- [“ClearErrorCodes method” on page 1058.](#)

Properties:

- AccountingToken property
- AccountingTokenHex property
- CompletionCode property
- CorrelationId property
- CorrelationIdHex property
- DistributionList property
- Feedback property
- GroupId property
- GroupIdHex property
- MessageId property
- MessageIdHex property
- NextDistributionListItem property
- PreviousDistributionListItem property
- QueueManagerName property
- QueueName property
- ReasonCode property
- ReasonName property

Methods:

- ClearErrorCodes method

Creation:

Use the AddDistributionListItem Method of the MQDistributionList class

AccountingToken property

Read-write. The AccountingToken to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

Defined in:

MQDistributionListItem class

Data Type:

String of 32 characters

Syntax: To get: *accountingtoken\$* = *MQDistributionListItem.AccountingToken*

To set: *MQDistributionListItem.AccountingToken* = *accountingtoken\$*

AccountingTokenHex property

Read-write. The AccountingToken to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 64 valid hexadecimal characters.

Its initial value is "0...0".

Defined in:

MQDistributionListItem class

Data Type:

String of 64 hexadecimal characters representing 32 ASCII characters.

Syntax: To get: *accountingtokenh\$* = *MQDistributionListItem.AccountingTokenHex*

To set: *MQDistributionListItem.AccountingTokenHex* = *accountingtokenh\$*

CompletionCode property

Read-only. The completion code set by the last open or put request issued against the owning distribution list object.

Defined in:

MQDistributionListItem class

Data Type:

Long

Values:

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *completioncode\$* = *MQDistributionListItem.CompletionCode*

CorrelationId property

Read-write. The CorrelId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

Defined in:

MQDistributionListItem class

Data Type:

String of 24 characters

Syntax: To get: *correlid\$* = *MQDistributionListItem.CorrelationId*

To set: *MQDistributionListItem.CorrelationId* = *correlid\$*

CorrelationIdHex property

Read-write. The CorrelId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

Defined in:

MQDistributionListItem class

Data Type:

String of 48 hexadecimal characters representing 24 ASCII characters.

Syntax: To get: *correlidh\$* = *MQDistributionListItem.CorrelationIdHex*

To set: *MQDistributionListItem.CorrelationIdHex* = *correlidh\$*

DistributionList property

Read-only. The distribution list with which this distribution list item is associated.

Defined in:

MQDistributionListItem class

Data Type:

MQDistributionList

Syntax: To get: *set distributionlist = MQDistributionListItem.DistributionList*

Feedback property

Read-write. The Feedback value to be included in the MQPMR of a message when put on a queue.

Defined in:

MQDistributionListItem class

Data Type:

Long

Values:

See [Feedback \(MQLONG\)](#).

Syntax: To get: *feedback& = MQDistributionListItem.Feedback*

To set: *MQDistributionListItem.Feedback = feedback&*

GroupId property

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

Defined in:

MQDistributionListItem class

Data Type:

String of 24 characters

Syntax: To get: *groupid\$ = MQDistributionListItem.GroupId*

To set: *MQDistributionListItem.GroupId = groupid\$*

GroupIdHex property

Read-write. The GroupId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

Defined in:

MQDistributionListItem class

Data Type:

String of 48 hexadecimal characters representing 24 ASCII characters.

Syntax: To get: *groupidh\$ = MQDistributionListItem.GroupIdHex*

To set: *MQDistributionListItem.GroupIdHex = groupidh\$*

MessageId property

Read-write. The MessageId to be included in the MQPMR of a message when put on a queue. Its initial value is all nulls.

Defined in:

MQDistributionListItem class

Data Type:

String of 24 characters

Syntax: To get: *messageid\$ = MQDistributionListItem.MessageId*

To set: *MQDistributionListItem*.**MessageId** = *messageid*\$

MessageIdHex property

Read-write. The MessageId to be included in the MQPMR of a message when put on a queue.

Every two characters of the string represent the hexadecimal equivalent of a single ASCII character. For example, the pair of characters "6" and "1" represent the single character "A", the pair of characters "6" and "2" represent the single character "B" and so on.

You must supply 48 valid hexadecimal characters.

Its initial value is "0..0".

Defined in:

MQDistributionListItem class

Data Type:

String of 48 hexadecimal characters representing 24 ASCII characters.

Syntax: To get: *messageidh*\$ = *MQDistributionListItem*.**MessageIdHex**

To set: *MQDistributionListItem*.**MessageIdHex** = *messageidh*\$

NextDistributionListItem property

Read-only. The next distribution list item object associated with the same distribution list.

Defined in:

MQDistributionListItem class

Data Type:

MQDistributionListItem

Syntax: To get: *set distributionlistitem* = *MQDistributionListItem*.**NextDistributionListItem**

PreviousDistributionListItem property

Read-only. The previous distribution list item object associated with the same distribution list.

Defined in:

MQDistributionListItem class

Data Type:

MQDistributionListItem

Syntax: To get: *set distributionlistitem* = *MQDistributionListItem*.**PreviousDistributionListItem**

QueueManagerName property

Read-write. The WebSphere MQ queue manager name.

Defined in:

MQDistributionListItem class

Data Type:

String of 48 characters.

Syntax: To get: *qmname*\$ = *MQDistributionListItem*.**QueueManagerName**

To set: *MQDistributionListItem*.**QueueManagerName** = *qmname*\$

QueueName property

Read-write. The WebSphere MQ queue name.

Defined in:

MQDistributionListItem class

Data Type:

String of 48 characters.

Syntax: To get: *qname\$ = MQDistributionListItem.QueueName*

To set: *MQDistributionListItem.QueueName = qname\$*

ReasonCode property

Read-only. The completion code set by the last open or put issued to the owning distribution list object.

Defined in:

MQDistributionListItem class

Data Type:

Long

Values:

See [API reason codes](#).

- MQCC_OK
- MQCC_WARNING
- MQCC_FAILED

Syntax: To get: *reasoncode& = MQDistributionListItem.ReasonCode*

ReasonName property

Read-only. The symbolic name for the ReasonCode. For example "MQRC_QMGR_NOT_AVAILABLE".

Defined in:

MQDistributionListItem class

Data Type:

String

Values:

See [API reason codes](#).

Syntax: To get: *reasonname\$ = MQDistributionListItem.ReasonName*

ClearErrorCodes method

Resets the CompletionCode to MQCC_OK and the ReasonCode to MQRC_NONE for both the MQDistributionListItem class and the MQSession class.

Defined in:

MQDistributionListItem class

Syntax: Call *MQDistributionListItem.ClearErrorCodes*

Troubleshooting

Information on the trace facility provided, common pitfalls and help on how to avoid them.

This following section explains the trace facility provided and details common pitfalls, with help to avoid them:

- [“Using trace” on page 1058](#)
- [“When your WebSphere MQ Automation Classes for ActiveX script fails” on page 1060](#)
- [“Reason codes” on page 1060](#)
- [“Code level tool” on page 1062](#)

Using trace

MQAX includes a trace facility to help the service organization identify what is happening when you have a problem. It shows the paths taken when you run your MQAX script. Unless you have a problem, run with tracing set off to avoid any unnecessary use of system resources.

There are three environment variables that you set to control trace:

- OMQ_TRACE
- OMQ_TRACE_PATH
- OMQ_TRACE_LEVEL

Note that specifying *any* value for OMQ_TRACE switches the trace facility on. Even if you set OMQ_TRACE to OFF, trace is still active.

To switch trace off, do not specify a value for OMQ_TRACE.

1. Click **Start**
2. Click **Control Panel**
3. Double-click **System**
4. Click **Advanced**
5. Click **Environment**
6. In the section titled "User variables for (username)", click **New**
7. Enter the variable name and a valid value in the appropriate fields and click **OK**
8. Click **OK** to close the Environment Variables window
9. Click **OK** to close the System Properties window
10. Close the Control Panel window

When deciding where you want the trace files written, ensure that you have sufficient authority to write to, not just read from, the disk.

With tracing switched on, it slows down the running of the MQAX, but it does not affect the performance of your ActiveX or WebSphere MQ environments. When you no longer need a trace file, you can delete it.

You must stop MQAX running to change the status of the OMQ_TRACE variable.

Trace file name and directory

The trace file name takes the form OMQnnnnn.trc, where nnnnn is the id of the ActiveX process running at the time.

Table 157. Commands and their effects	
Command	Effect
SET OMQ_TRACE_PATH = drive:\directory	Sets the trace directory where the trace file will be written.
SET OMQ_TRACE_PATH =	Removes the OMQ_PATH environment variable the current working directory (when ActiveX is started) is used.
ECHO %OMQ_TRACE_PATH%	Displays the current setting of the trace directory on Windows.
SET OMQ_TRACE = xxxxxxxx	This sets tracing ON. You switch tracing on by putting one or more characters after the '=' sign. For example: SET OMQ_TRACE=yes SET OMQ_TRACE=no. In both of these examples, tracing will be set ON. This is only effective for a single window/session
SET OMQ_TRACE=	Sets tracing OFF
ECHO %OMQ_TRACE%	Displays the contents of the environment variable on Windows.
SET	Displays the contents of all the environment variables on Windows.
SET OMQ_TRACE_LEVEL = 9	Sets the trace level to 9. Values greater than 9 do not produce any additional information in the trace file.

When your WebSphere MQ Automation Classes for ActiveX script fails

If your WebSphere MQ Automation Classes for ActiveX script fails, there are a number of sources of information.

First failure symptom report

Independently of the trace facility, for unexpected and internal errors, a First failure symptom report might be produced.

This report is found in a file named OMQnnnnn.fdc, where nnnnn is the number of the ActiveX process that is running at the time. You find this file in the working directory from which you started ActiveX or in the path specified in the OMQ_PATH environment variable.

Other sources of information

WebSphere MQ provides various error logs and trace information, depending on the platform involved. See your Windows NT application event log.

Reason codes

The following reason codes can occur in addition to those documented for the WebSphere MQ MQI. For other codes, refer to your WebSphere MQ application event log.

Table 158. Reason codes and what they mean	
Reason code	Explanation
MQRC_LIBRARY_LOAD_ERROR (6000)	One or more of the WebSphere MQ libraries could not be loaded. Check that all WebSphere MQ libraries are in the correct search path on the system you are using. For example, make sure that the directories containing the WebSphere MQ libraries are in PATH.
MQRC_CLASS_LIBRARY_ERROR (6001)	One of the WebSphere MQ classlibrary calls returned an unexpected ReasonCode/CompletionCode. Check the First Failure Symptom Report for details. Take note of the last method/property and class being used and inform IBM Support of the problem.
MQRC_STRING_LENGTH_TOO_BIG (6002)	An attempt has been made to write a UTF format string with a length greater than 65,535 bytes to the message buffer.
MQRC_WRITE_VALUE_ERROR (6003)	A value is used that is out of range; for example msg.WriteByte (240).
MQRC_PACKED_DECIMAL_ERROR (6004)	An attempt has been made to read a packed decimal number from the message buffer but the data at the data pointer is not in a valid packed data format.
MQRC_FLOAT_CONVERSION_ERROR (6005)	An attempt has been made to read a single or double floating point number from the message buffer but the data at the data pointer is not in an appropriate floating point format.
MQRC_REOPEN_EXCL_INPUT_ERROR (6100)	An open object does not have the correct OpenOptions and requires one or more additional options. An implicit reopen is required but closure has been prevented. Set the OpenOptions explicitly to cover all eventualities so that implicit reopening is not required. Closure has been prevented because the queue is open for exclusive input and closure would present a window of opportunity for others potentially to gain access to the queue.
MQRC_REOPEN_INQUIRE_ERROR (6101)	An open object does not have the correct OpenOptions and requires one or more additional options. An implicit reopen is required but closure has been prevented. Set the OpenOptions explicitly to include MQOO_INQUIRE. Closure has been prevented because one or more characteristics of the object need to be checked dynamically prior to closure, and the OpenOptions do not already include MQOO_INQUIRE.

Table 158. Reason codes and what they mean (continued)

Reason code	Explanation
MQRC_REOPEN_SAVED_CONTEXT_ERR (6102)	An open object does not have the correct OpenOptions and requires one or more additional options. An implicit reopen is required but closure has been prevented. Set the OpenOptions explicitly to cover all eventualities so that implicit reopening is not required. Closure has been prevented because the queue is open with MQOO_SAVE_ALL_CONTEXT, and a destructive Get has been performed previously. This has caused retained state information to be associated with the open queue and this information would be destroyed by closure.
MQRC_REOPEN_TEMPORARY_Q_ERROR (6103)	An open object does not have the correct OpenOptions and requires one or more additional options. An implicit reopen is required, but closure has been prevented. Set the OpenOptions explicitly to cover all eventualities so that implicit reopening is not required. Closure has been prevented because the queue is a local queue of the definition type MQQDT_TEMPORARY_DYNAMIC, which would be destroyed by closure.
MQRC_ATTRIBUTE_LOCKED (6104)	An attempt has been made to change the value or attribute of an object while that object is open. Certain attributes, such as AlternateUserId , cannot be changed while an object is open.
MQRC_CURSOR_NOT_VALID (6105)	The browse cursor for an open queue has been invalidated since it was last used by an implicit reopen. Set the OpenOptions explicitly to cover all eventualities so that implicit reopening is not required.
MQRC_ENCODING_ERROR (6106)	The encoding of the next message item needs to be MQENC_NATIVE for reading.
MQRC_STRUCID_ERROR (6107)	The structure of the id for the next message item, which is derived from the 4 characters beginning at the data pointer, is either missing or is inconsistent with the type of variable into which the item is being read.
MQRC_NULL_POINTER (6108)	A null pointer has been supplied where a non-null pointer is either required or implied. This may be caused by using explicit declarations for WebSphere MQ objects used from VBA as parameters to calls (for example dim msg as Object is ok, dim msg as MqMessage can cause problems). For example, in Excel, with q defined and set dim msg as MqMessageq.put msg gives reasonCode MQRC_NULL_POINTER. It operates correctly from VisualBasic.
MQRC_NO_CONNECTION_REFERENCE (6109)	The MQQueue object has lost its connection to the MQQueueManager . This will occur if the MQQueueManager is disconnected. Delete the MQQueue object.
MQRC_NO_BUFFER (6110)	No buffer is available. For an MQMessage object, one cannot be allocated, denoting an internal inconsistency in the object state that should not occur.
MQRC_BINARY_DATA_LENGTH_ERROR (6111)	The length of the binary data is inconsistent with the length of the target attribute. Zero is a correct length for all attributes. 24 is a correct length for a CorrelationId and for a MessageId 32 is a correct length for an AccountingToken
MQRC_BUFFER_NOT_AUTOMATIC (6112)	A user-defined and managed buffer cannot be resized. Because message buffers are system managed, this indicates an internal inconsistency.
MQRC_INSUFFICIENT_BUFFER (6113)	There is insufficient buffer space available after the data pointer to accommodate the request. This could be because the buffer cannot be resized.
MQRC_INSUFFICIENT_DATA (6114)	There is insufficient data after the data pointer to accommodate the read request. Reduce the buffer to the correct size and read the data again.
MQRC_DATA_TRUNCATED (6115)	Data has been truncated when copying from one buffer to another. This could be because the target buffer cannot be resized, or because there is a problem addressing one or other buffer, or because a buffer is being downsized with a smaller replacement.

Table 158. Reason codes and what they mean (continued)

Reason code	Explanation
MQRC_ZERO_LENGTH (6116)	A zero length has been supplied where a positive length is either required or implied.
MQRC_NEGATIVE_LENGTH (6117)	A negative length has been supplied where a zero or positive length is required.
MQRC_NEGATIVE_OFFSET (6118)	A negative offset has been supplied where a zero or positive offset is required.
MQRC_INCONSISTENT_FORMAT (6119)	The format of the next message item is inconsistent with the type of variable into which the item is being read.
MQRC_INCONSISTENT_OBJECT_STATE (6120)	There is an inconsistency between this object, which is open, and the referenced MQQueueManager object, which is not connected.
MQRC_CONTEXT_OBJECT_NOT_VALID (6121)	The MQPutMessageOptions context reference does not reference a valid MQQueue object. The object has been previously destroyed.
MQRC_CONTEXT_OPEN_ERROR (6122)	The MQPutMessageOptions context reference references an MQQueue object that could not be opened to establish a context. This may be because the MQQueue object has inappropriate open options. Inspect the referenced object reason code to establish the cause.
MQRC_STRUC_LENGTH_ERROR (6123)	The length of an internal data structure is inconsistent with its content. For an MQRMH, the length is insufficient to contain the fixed fields and all offset data.
MQRC_NOT_CONNECTED (6124)	A method failed because a required connection to a queue manager was not available, and a connection cannot be established implicitly.
MQRC_NOT_OPEN (6125)	A method failed because a WebSphere MQ object was not open, and opening cannot be accomplished implicitly.
MQRC_DISTRIBUTION_LIST_EMPTY (6126)	An MQDistributionList failed to open because there are no MQDistributionList Item objects in the distribution list. Corrective action: Add at least one MQDistributionListItem object to the distribution list.
MQRC_INCONSISTENT_OPEN_OPTIONS (6127)	A method failed because the object is open, and the open options are inconsistent with the required operation. Corrective action: Open the object with appropriate open options and retry.
MQRC_WRONG_VERSION (6128)	A method failed because a version number specified or encountered is either incorrect or not supported.

Code level tool

You might be asked by the IBM Service Team which level of code you have installed.

To find this out, run the 'MQAXLEV' utility program.

From the command prompt, change to the directory containing the MQAX200.dll or add the full path length and enter:

```
MQAXLev MQAX200.dll > MQAXLEV.OUT
```

where MQAXLEV.OUT is the name of the output file.

If you do not specify an output file, the detail is displayed on the screen.

An example output file from code level tool is detailed in the following example:

Example output file from code level tool

```
5639-B43 (C) Copyright IBM Corp. 1996, 2024. ALL RIGHTS RESERVED.
***** Code Level is 5.1 ***** lib/mqole/mqole.cpp, mqole, p000, L981119 1.8 98/08/21
lib/mqlsx/gmqdyn0a.c, mqlsx, p000, p000 L990212 1.6 99/02/11 16:40:24
lib/mqlsx/pc/gmqdyn1p.c, mqlsx, p000, p000 L990212 1.6 99/02/11 16:44:14
lib/mqlsx/xmqcsa.c, mqole, p000, p000 L990216 1.3 99/02/15 13:24:34
lib/mqlsx/xmqfdca.c, mqlsx, p000, p000 L990212 1.3 99/02/11 16:40:35
lib/mqlsx/xmqtrca.c, mqlsx, p000, p000 L990212 1.5 99/02/11 16:12:02
lib/mqlsx/xmqutila.c, mqlsx, p000, p000 L990212 1.3 99/02/11 16:40:40
lib/mqlsx/xmqutl1a.c, mqlsx, p000, p000 L990212 1.4 99/02/11 16:40:30
lib/mqlsx/xmqcnv1a.c, mqlsx, p000, p000 L990212 1.9 99/02/11 16:40:56
lib/mqlsx/xmqmsg.c, mqole, p000, p000 L990219 1.11 99/02/18 12:12:59
```

ActiveX interface to the MQAI

For a brief overview of COM interfaces and their use in the MQAI, see [“Using the Component Object Model Interface \(WebSphere MQ Automation Classes for ActiveX\)”](#) on page 987.

The MQAI enables applications to build and send Programmable Command Format (PCF) commands without directly obtaining and formatting the variable length buffers required for PCF. For more information about the MQAI, see [Introduction to the WebSphere MQ Administration Interface \(MQAI\)](#). The MQAI ActiveX MQBag class encapsulates the data bags supported by the MQAI in a way that is possible to use in any language that supports the creation of COM objects; for example, Visual Basic, C++, Java, and other ActiveX scripting clients.

The MQAI ActiveX interface is for use with the MQAX classes that provide a COM interface to the MQI. For more information about the MQAX classes, see [“Designing MQAX applications that access non-ActiveX applications”](#) on page 988.

The ActiveX interface provides a single class called MQBag. This class is used to create MQAI data bags and its properties and methods are used to create and work with data items within each bag. The MQBag Execute method sends the bag data to a WebSphere MQ queue manager as a PCF message and collects the replies.

For more information about MQBag class, its properties and methods, see [“The MQBag class”](#) on page 1063.

The PCF message is sent to the queue manager object specified, optionally using specified request and reply queues. Replies are returned in a new MQBag object. The full set of commands and replies is described in [Definitions of the Programmable Command Formats](#). Commands can be sent to any queue manager in the WebSphere MQ network by selecting the appropriate request and reply queues.

The MQBag class

The class, MQBag, is used to create MQBag objects as required. When instantiated, the MQBag class returns a new MQBag object reference.

Create a MQBag object in Visual Basic as follows:

```
Dim mqbag As MQBag
Set mqbag = New MQBag
```

MQBag Property

The properties of MQBag objects are explained in the following list:

- [“Item property”](#) on page 1064.
- [“Count property”](#) on page 1065.
- [“Options property”](#) on page 1066.

MQBag methods

The methods of the MQBag objects are explained in the following list:

- [“Add method” on page 1066.](#)
- [“AddInquiry method” on page 1067.](#)
- [“Clear method” on page 1067.](#)
- [“Execute method” on page 1068.](#)
- [“FromMessage method” on page 1068.](#)
- [“ItemType method” on page 1069.](#)
- [“Remove method” on page 1069.](#)
- [“Selector method” on page 1070.](#)
- [“ToMessage method” on page 1070.](#)
- [“Truncate method” on page 1071.](#)

Error Handling

If an error is detected during an operation on an MQBag object, including those errors returned to the bag by an underlying MQAX or MQAI object, an error exception is raised. The MQBag class supports the COM ISupportErrorInfo interface so the following information is available to your error handling routine:

- Error number: composed of the WebSphere MQ reason code for the error detected and a COM facility code. The facility field, as standard for COM, indicates the area of responsibility for the error. For errors detected by WebSphere MQ it is always FACILITY_ITF.
- Error source: identifies the type and version of the object that detected the error. For errors detected during MQBag operations, the error source is always MQBag.MQBag1.
- Error description: the string giving the symbolic name for the WebSphere MQ reason code.

How you access the error information depends on your scripting language; for example, in Visual Basic the information is returned in the Err object and the WebSphere MQ reason code is obtained by subtracting the constant vbObjectError from Err.Number.

ReasonCode = Err.Number - vbObjectError

If the MQBag Execute message sends a PCF message and a reply is received, the operation is considered successful although the command sent might have failed. In this case, the reply bag itself contains the completion and error reason codes as described in [Definitions of the Programmable Command Formats](#).

Item property

Purpose

The Item property represents an item in a bag. It is used to set or inquire about the value of an item. Use of this property corresponds to the following MQAI calls:

- "mqSetString"
- "mqSetInteger"
- "mqInquireInteger"
- "mqInquireString"
- "mqInquireBag"

in the [Programmable command formats reference](#).

Format

Item (Selector, ItemIndex, Value)

Parameters

Selector (VARIANT) - input

Selector of the item to be set or inquired.

When inquiring about an item, MQSEL_ANY_USER_SELECTOR is the default. When setting an item, MQIA_LIST or MQCA_LIST is the default.

If the Selector is not of type long, MQRC_SELECTOR_TYPE_ERROR results.

This parameter is optional.

ItemIndex (LONG) - input

This value identifies the occurrence of the item of the specified selector that is to be set or inquired on. MQIND_NONE is the default.

This parameter is optional.

Value (VARIANT) - input/output

The value returned or the value to be set. When inquiring about an item, the return value can be of type long, string, or MQBag. However, when setting an item, the value must be of type long or string; if not, MQRC_ITEM_VALUE_ERROR results.

Visual Basic Language Invocation

When inquiring about a value of an item within a bag:

```
Value = mqbag[.Item]([Selector],  
[ItemIndex])
```

For MQBag references:

```
Set abag = mqbag[.Item]([Selector].  
[ItemIndex])
```

To set the value of an item in a bag:

```
mqbag[.Item]([Selector],  
[ItemIndex]) = Value
```

Count property

Purpose

The Count property represents the number of data items within a bag. This property corresponds to the MQAI call, "mqCountItems," in the [Programmable command formats reference](#).

Format

Count (Selector, Value)

Parameters

Selector (VARIANT) - input

Selector of the data items to be included in the count.

MQSEL_ALL_USER_SELECTORS is the default.

If the Selector is not of type long, MQRC_SELECTOR_TYPE_ERROR is returned.

Value (LONG) - output

The number of items in the bag included by the *Selector*.

Visual Basic Language Invocation

To return the number of items in a bag:

```
ItemCount = mqbag.Count([Selector])
```

Options property

Purpose

The Options property sets options for the use of a bag. This property corresponds to the Options parameter of the MQAI call, "mqCreateBag," in the [Programmable command formats reference](#).

Format

Options (*Options*)

Parameters

Options (LONG) - input/output

The bag options.

Note: The bag options must be set *before* data items are added to or set within the bag. If the options are changed when the bag is not empty, MQRC_OPTIONS_ERROR results. This applies even if the bag is subsequently cleared.

Visual Basic Language Invocation

When inquiring about the options of an item within a bag:

```
Options = mqbag.Options
```

To set an option of an item in a bag:

```
mqbag.Options = Options
```

MQBag methods

The methods of the MQBag objects are explained over the following pages.

Add method

Purpose

The Add method adds a data item to a bag. This method corresponds to the MQAI calls, "mqAddInteger" and "mqAddString," in the [Programmable command formats reference](#).

Format

Add (*Value*, *Selector*)

Parameters

Value (VARIANT) - input

Integer or string value of the data item.

Selector (VARIANT) - input

Selector identifying the item to be added.

Depending on the type of Value, MQIA_LIST or MQCA_LIST is the default. If the Selector parameter is not of type long, MQRC_SELECTOR_TYPE_ERROR results.

Visual Basic Language Invocation

To add an item to a bag:

```
mqbag.Add(Value,[Selector])
```

AddInquiry method

Purpose

The AddInquiry method adds a selector specifying the attribute to be returned when an administration bag is sent to execute an INQUIRE command. This method corresponds to the MQAI call, "mqAddInquiry," in the [Programmable command formats reference](#).

Format

AddInquiry (Inquiry)

Parameters

Inquiry (LONG) - input

Selector of the WebSphere MQ attribute to be returned by the INQUIRE administration command.

Visual Basic Language Invocation

To use the AddInquiry method:

```
mqbag.AddInquiry(Inquiry)
```

Clear method

Purpose

The Clear method deletes all data items from a bag. This method corresponds to the MQAI call, "mqClearBag," in the [Programmable command formats reference](#).

Format

Clear

Visual Basic Language Invocation

To delete all data itmes from a bag:

```
mqbag.Clear
```

Execute method

Purpose

The Execute method sends an administration command message to the command server and waits for any reply messages. This method corresponds to the MQAI call, "mqExecute," in the [Programmable command formats reference](#).

Format

Execute (*QueueManager*, *Command*, *OptionsBag*, *RequestQ*, *ReplyQ*, *ReplyBag*)

Parameters

***QueueManager* (MQQueueManager) - input**

The queue manager to which the application is connected.

***Command* (LONG) - input**

The command to be executed.

***OptionsBag* (MQBag) - input**

The bag containing options that affect the processing of the call.

***RequestQ* (MQQueue) - input**

The queue on which the administration command message will be placed.

***ReplyQ* (MQQueue) - input**

The queue on which any reply messages are received.

***ReplyBag* (MQBag) - output**

A bag reference containing data from reply messages.

Visual Basic Language Invocation

To send an administration command message and wait for any reply messages:

```
Set ReplyBag = mqbag.Execute(QueueManager, Command,  
[OptionsBag],[RequestQ],[ReplyQ])
```

FromMessage method

Purpose

The FromMessage method loads data from a message into a bag. This method corresponds to the MQAI call, "mqBufferToBag," in the [Programmable command formats reference](#).

Format

FromMessage (*Message*, *OptionsBag*)

Parameters

***Message* (MQMessage) - input**

The message containing the data to be converted.

***OptionsBag* (MQBag) - input**

Options to control the processing of the call.

Visual Basic Language Invocation

To load data from a message into a bag:

```
mqbag.FromMessage(Message,[OptionsBag])
```

ItemType method

Purpose

The ItemType method returns the type of the value in a specified item in a bag. This method corresponds to the MQAI call, "mqInquireItemInfo," in the [Programmable command formats reference](#).

Format

ItemType (Selector, ItemIndex, ItemType)

Parameters

Selector (VARIANT) - input

Selector identifying the item to be inquired.

MQSEL_ANY_USER_SELECTOR is the default. If the Selector parameter is not of type long, MQRC_SELECTOR_TYPE_ERROR results.

ItemIndex (LONG) - input

Index of items to be inquired.

MQIND_NONE is the default.

ItemType (LONG) - output

Data type of the specified item.

Note: Either the Selector parameter, ItemIndex parameter, or both must be specified. If neither parameter is present, MQRC_PARAMETER_MISSING results.

Visual Basic Language Invocation

To return the type of a value:

```
ItemType = mqbag.ItemType([Selector],  
[ItemIndex])
```

Remove method

Purpose

The Remove method deletes an item from a bag. This method corresponds to the MQAI call, "mqDeleteItem," in the [Programmable command formats reference](#).

Format

Remove (Selector, ItemIndex)

Parameters

Selector (VARIANT) - input

Selector identifying the item to be deleted.

MQSEL_ANY_USER_SELECTOR is the default. If the Selector parameter is not of type long, MQRC_SELECTOR_TYPE_ERROR results.

ItemIndex (LONG) - input

Index of the item to be deleted.

MQIND_NONE is the default.

Note: Either the Selector parameter, ItemIndex parameter, or both must be specified. If neither parameter is present, MQRC_PARAMETER_MISSING results.

Visual Basic Language Invocation

To delete an item from a bag:

```
mqbag.Remove([Selector],[ItemIndex])
```

Selector method

Purpose

The Selector method returns the selector of a specified item within a bag. This method corresponds to the MQAI call, "mqInquireItemInfo," in the [Programmable command formats reference](#).

Format

Selector (Selector, ItemIndex, OutSelector)

Parameters

Selector (VARIANT) - input

Selector identifying the item to be inquired.

MQSEL_ANY_USER_SELECTOR is the default. If the Selector parameter is not of type long, MQRC_SELECTOR_TYPE_ERROR results.

ItemIndex (LONG) - input

Index of the item to be inquired.

MQIND_NONE is the default.

OutSelector (VARIANT) - output

Selector of the specified item.

Note: Either the Selector parameter, ItemIndex parameter, or both must be specified. If neither parameter is present, MQRC_PARAMETER_MISSING results.

Visual Basic Language Invocation

To return the selector of an item:

```
OutSelector = mqbag.Selector([Selector],  
[ItemIndex])
```

ToMessage method

Purpose

The ToMessage method returns a reference to an MQMessage object. The reference contains data from a bag. This method corresponds to the MQAI call, "mqBagToBuffer," in the [Programmable command formats reference](#).

Format

ToMessage (OptionsBag, Message)

Parameters

OptionsBag (MQBag) - input

A bag containing options that control the processing of the method.

Message (MQMessage) - output

An MQMessage object reference containing data from the bag.

Visual Basic Language Invocation

To use the ToMessage Method:

```
Set Message = mqbag.ToMessage([OptionsBag])
```

Truncate method

Purpose

The Truncate method reduces the number of user items in a bag. This method corresponds to the MQAI call, "mqTruncateBag," in the [Programmable command formats reference](#).

Format

Truncate (ItemCount)

Parameters

ItemCount (LONG) - input

The number of user items to remain in the bag after truncation has occurred.

Visual Basic Language Invocation

To reduce the number of user items in a bag:

```
mqbag.Truncate(ItemCount)
```

About the WebSphere MQ Automation Classes for ActiveX Starter samples

This appendix describes the WebSphere MQ Automation Classes for ActiveX Starter samples, and explains how to use them.

WebSphere MQ for Windows provides the following Visual Basic sample programs:

- MQAXTRIV.VBP
- MQAXBSRV.VBP
- MQAXDLST.VBP
- MQAXCLSS.VBP

These samples run on Visual Basic 4 or Visual Basic 5. You will find them in the directory ...\\tools\\mqax\\samples\\vb.

In the same directory you will also find samples for Microsoft Excel and html. These are:

- MQAX.XLS
- MQAXTRIV.XLS
- MQAXTRIV.HTM

Note: If using Visual Basic 5, you **must** select and install Visual Basic component grid32.ocx.

What is demonstrated in the samples

The samples demonstrate how to use WebSphere MQ Automation Classes for ActiveX to:

- Connect to a queue manager
- Access a queue
- Put a message on a queue
- Get a message from a queue

The central part of the Visual Basic sample is shown on the following pages.

[“Preparing to run the samples” on page 1072](#) and

[“Error handling in the samples” on page 1073](#)

Running the ActiveX Starter samples

Before you run the WebSphere MQ Automation Classes for ActiveX Starter samples check that you have a default queue manager running and that you have created the required queue definitions. For details of creating and running a queue manager and creating a queue, refer to [Administering](#). The sample uses the queue SYSTEM.DEFAULT.LOCAL.QUEUE which should be defined on any normally set up WebSphere MQ server.

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

- Connect to a queue manager
- Access a queue
- Put a message on a queue
- Get a message from a queue

For information on the MQAX starter samples for Microsoft Basic Version 4 or later, see [“Running the MQAXTRIV sample” on page 1073](#)

For information on a sample that allow you to browse properties and methods of queue managers and queue objects, see [“Starting the MQAXCLSS sample” on page 1074](#)

For information on the MQAXDLST sample, [“The MQAXDLST sample” on page 1075](#)

For information on running the MQAX starter sample for Microsoft Excel 95 or later, MQAXTRIV.XLS, see [“Running the MQAXTRIV.XLS sample” on page 1075](#).

For information on running the Bank demonstration with MQAX.XLS, see [“Running the Bank demonstration with MQAX.XLS” on page 1075](#)

For information on starter sample using an ActiveX compatible WWW browser, see [“Starter sample using an ActiveX compatible WWW browser” on page 1075](#)

Preparing to run the samples

To run any of the samples you need one of the following depending on which of the samples you intend to run.

- Microsoft Visual Basic Version 4 (or later)
- Microsoft Excel 95 (or later)
- A web browser

You also need:

- A WebSphere MQ queue manager running.
- A WebSphere MQ queue already defined.

Error handling in the samples

Most of the samples provided in the WebSphere MQ Automation Classes for ActiveX package exhibit little or no error handling. For more information about error handling, see [“Error handling” on page 992](#).

Running the MQAXTRIV sample

1. Start the queue manager.
2. In Windows Explorer or File Manager, select the icon for the sample, MQAXTRIV.VBP (Visual Basic Project file) and open the file.

The Visual Basic program starts and opens the file, MQAXTRIV.VBP.

3. In Visual Basic, press function key 5 (F5) to run the sample.
4. Click anywhere in the window form, "MQAX trivial tester".

If everything is working correctly, the window background should change to green. If there is a problem with your setup, the window background should change to red and error information will be displayed.

The following figure shows the central part of the Visual Basic sample.

```
Option Explicit

Private Sub Form_Click()

'*****
'* This simple example illustrates how to put and get a WebSphere MQ message to
'* and from a WebSphere MQ message queue. The data from the message returned by the
'*get is read and compared with that from the original message.
'*****
Dim MQSess As MQSession          '* session object
Dim QMgr As MQQueueManager       '* queue manager object
Dim Queue As MQQueue             '* queue object
Dim PutMsg As MQMessage          '* message object for put
Dim GetMsg As MQMessage          '* message object for get
Dim PutOptions As MQPutMessageOptions '* get message option

Dim GetOptions As MQGetMessageOptions '* put message options
Dim PutMsgStr As String           '* put message data string
Dim GetMsgStr As String           '* get message data string
'*****
'* Handle errors
'*****
On Error GoTo HandleError

'*****
'* Initialize the current position for the form
'*****
CurrentX = 0
CurrentY = 0

'*****
'* Create the MQSession object and access the MQQueueManager and (local) MQQueue
'*****
Set MQSess = New MQSession
Set QMgr = MQSess.AccessQueueManager("")
Set Queue = QMgr.AccessQueue("SYSTEM.DEFAULT.LOCAL.QUEUE", _
                             MQOO_OUTPUT Or MQOO_INPUT_AS_Q_DEF)

'*****
'* Create a new MQMessage object for use with put, add some data then create an
'* MQPutMessageOptions object and put the message
'*****
Set PutMsg = MQSess.AccessMessage()
PutMsgStr = "12345678 " & Time
PutMsg.MessageData = PutMsgStr
Set PutOptions = MQSess.AccessPutMessageOptions()
Queue.Put PutMsg, PutOptions

'*****
'* Create a new MQMessage object for use with get, set the MessageId (to that of
'* the message that was put), create an MQGetMessageOptions object and get the
'* message.
```

```

'*
'* Note: Setting the MessageId ensures that the get returns the MQMessage
'* that was put earlier.
'*****

Set GetMsg = MQSess.AccessMessage()
GetMsg.MessageId = PutMsg.MessageId
Set GetOptions = MQSess.AccessGetMessageOptions()
Queue.Get GetMsg, GetOptions
'*****
'* Read the data from the message returned by the get, compare it with
'* that from the original message and output a suitable message.
'*****
GetMsgStr = GetMsg.MessageData
Cls
If GetMsgStr = PutMsgStr Then
    BackColor = RGB(127, 255, 127) '* set to green for ok
    Print
    Print "Message data comparison was successful."
    Print "Message data: "" & GetMsgStr & """"

Else
    BackColor = RGB(255, 255, 127) '* set to amber for compare error
    Print "Compare error: "
    Print "The message data returned by the get did not match the " & _
    "input data from the original message that was put."
    Print
    Print "Input message data: "" & PutMsgStr & """"
    Print "Returned message data: "" & GetMsgStr & """"
End If

Exit Sub
'*****
'* Handle errors
'*****
HandleError:
Dim ErrMsg As String
Dim StrPos As Integer

Cls
BackColor = RGB(255, 0, 0) '* set to red for error
Print "An error occurred as follows:"
Print ""
If MQSess.CompletionCode <> MQCC_OK Then
    ErrMsg = Err.Description
    StrPos = InStr(ErrMsg, " ") '* search for first blank
    If StrPos > 0 Then
        Print Left(ErrMsg, StrPos) '* print offending MQAX object name
    Else
        Print Error(Err) '* print complete error object
    End If
    Print ""
    Print "WebSphere MQ Completion Code = " & MQSess.CompletionCode
    Print "WebSphere MQ Reason Code = " & MQSess.ReasonCode
    Print "(" & MQSess.ReasonName & ")"
Else
    Print "Visual Basic error: " & Err
    Print Error(Err)
End If

Exit Sub

End Sub

```

Starting the MQAXCLSS sample

This sample allows you to browse properties and methods of queue managers and queue objects.

1. Start the queue manager.
2. Open the file, MQAXCLSS.VBP, by double clicking on the document icon in Windows Explorer or by clicking File - Open from the file menu in Visual Basic.
3. Start the sample.
4. Enter the appropriate queue manager and queue names then click the corresponding buttons.

The MQAXDLST sample

The Visual Basic MQAXDLST sample demonstrates the use of a distribution list to send the same message to two queues with one put. To run the sample, do the same as for the MQAXCLSS sample.

MQAX Starter sample for Microsoft Excel 95 or later

This section explains how to run the MQAX starter sample for Microsoft Excel 95 or later, MQAXTRIV.XLS.

Running the MQAXTRIV.XLS sample

1. Start the queue manager.
2. In Explorer or File Manager, select the icon for the MQAX sample MQAXTRIV.XLS.
3. Click the button in the spreadsheet.
4. The screen is updated with a success (or failure) message.

Running the Bank demonstration with MQAX.XLS

Follow these steps to run the Bank demonstration.

1. Start the queue manager.
2. Run the IBM WebSphere MQ MQSC command file, BANK.TST. This sets up the necessary IBM WebSphere MQ queue definitions.

To find out how to use an MQSC command file, refer to [Script \(MQSC\) Commands](#).

3. Run MQAXBRSRV.VBP. This sample program is the server simulating a back-end application, and it has to run with Microsoft Excel.
4. Run MQAX.XLS. This sample is the client IBM WebSphere MQ demonstration.
5. Select a customer from the list.
6. Click **Submit**.

After a short pause (approximately 3 seconds), the fields populate with values and a bar chart is displayed.

Starter sample using an ActiveX compatible WWW browser

Note: To run this sample, you must be running an ActiveX compatible Web browser. Microsoft Internet Explorer (but not Netscape Navigator) is a compatible Web browser.

Running the HTML sample

This sample demonstrates how you can invoke MQAX from both VBScript and JavaScript.

1. Start the queue manager.
2. Open the file, "MQAXTRIV.HTM", in your ActiveX compatible Web browser.

You can do this either by double-clicking the file icon in Windows Explorer or you can choose File - Open from the File menu of your ActiveX compatible Web browser.

3. Follow the instructions on the screen.

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 IBM 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.



Part Number:

(1P) P/N: