WebSphere MQ for z/OS

Concepts and Planning Guide

Version 5 Release 3.1

GC34-6051-01

Note!
Before using this information and the product it supports, be sure to read the general information under Appendix C, "Notices", on page 181.

Third edition (March 2003)

This edition applies to WebSphere MQ for z/OS Version 5 Release 3.1, and to all subsequent releases and modifications until otherwise indicated in new editions.

Editions of this book before WebSphere MQ for z/OS Version 5 Release 3 were entitled *MQSeries for OS/390 Concepts and Planning Guide*, GC34-5650-00.

Contents

Figures

Tables

About this book

This book describes the concepts of IBM® WebSphere® MQ for z/OS™; it does not describe the concepts of WebSphere MQ messaging and queuing. If you are unfamiliar with these concepts, you should read *MQSeries : An Introduction to Messaging and Queuing*. This book also describes how to plan your WebSphere MQ for z/OS systems.

Changes to the information in this book since the last edition are marked with vertical bars in the left-hand margin.

Who this book is for

This book is for:

- Planners of z/OS systems using WebSphere MQ message queuing techniques.
- System programmers who install, customize, and operate WebSphere MQ for z/OS.
- Users of WebSphere Application Server who want to know what functions of WebSphere MQ for z/OS are unavailable in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.

What you need to know to understand this book

This book assumes that you are familiar with the basic concepts of:

- CICS®
- DB2® (if you are going to use queue-sharing groups)
- IMS™
- WebSphere MQ
- z/OS
- The zSeries™ Coupling Facility (if you are going to use queue-sharing groups)

Conventions used in this book

- z/OS means any release of z/OS or OS/390® that supports the current version of WebSphere MQ.
- The examples in this book are taken from a queue manager with a command prefix string (CPF) of +CSQ1. The commands are shown in UPPERCASE.
- CICS means both CICS Transaction Server for OS/390 (z/OS) and CICS for MVS/ESA™ unless otherwise stated. IMS means IMS/ESA® unless otherwise stated.
- Throughout this book, the default value `thlqual` is used to indicate the target library high-level qualifier for WebSphere MQ data sets in your installation.
- Throughout this book, the term *distributed queuing* refers to the distributed queuing feature (also known as the *non-CICS mover*). The term *distributed queuing using CICS ISC* is used to refer to the optional CICS distributed queuing feature (also known as the *CICS mover*).

What's new for this release

If you are already familiar with previous versions of WebSphere MQ for z/OS, "What's new in WebSphere MQ for z/OS" on page 167 summarizes the new functions that have been added for Version 5.3.1, and explains where to find more information about them.

Summary of changes

This section describes changes in this edition of *WebSphere MQ for z/OS Concepts and Planning Guide*. Changes since the previous edition of the book are marked by vertical lines to the left of the changes.

Changes for this edition (GC34–6051–01)

The main change in this edition is the addition of information for users of the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.

WebSphere MQ for z/OS Version 5 Release 3.1 provides JMS support for WebSphere Application Server asynchronous ***embedded messaging*** as part of the WebSphere JMS provider. Embedded messaging is implemented by WebSphere Application Server using either the supplied WebSphere MQ for z/OS Version 5 Release 3.1 reduced function queue manager, or a WebSphere MQ for z/OS Version 5 Release 3.1 full function queue manager. The reduced function queue manager provides simple point-to-point messaging. For complete information on installing, configuring, managing, and using the WebSphere JMS provider, see the WebSphere Application Server InfoCenter.

In Chapter 14, "WebSphere MQ and WebSphere Application Server", on page 113, you'll find a general introduction to the functions that are not available in the reduced function form of WebSphere MQ for z/OS. Throughout the rest of this book, and all other books providing information for WebSphere MQ for z/OS users, we identify the unavailable features, commands, API verbs, and so on, as part of their descriptions.

In this book, you will mainly see a sentence of the following form in a note at the start of chapters, sections, or paragraphs:

- This is **not** available when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

Changes for the previous edition (GC34–6051–00)

- WebSphere MQ is the new name for MQSeries®.
- The minimum level of z/OS required for WebSphere MQ for z/OS Version 5 Release 3 is OS/390 V2.9.
- WebSphere MQ is now fully integrated with the Secure Sockets Layer (SSL) protocol. A new WebSphere MQ object called an *authentication information object* has been introduced for use with SSL. Full details of the SSL implementation on WebSphere MQ can be found in the *WebSphere MQ Security* book.
- You can now use persistent messages with queue-sharing groups, as they can be logged in the queue manager logs. To support the recovery of shared queue messages, a new object called a *CF structure* (CFSTRUCT), and two new WebSphere MQ commands, BACKUP CFSTRUCT and RECOVER CFSTRUCT, have been introduced. For details of these commands, see the *WebSphere MQ Script (MQSC) Command Reference*.

Changes

- Some queue manager parameters can now be changed while the queue manager is active, using the SET LOG, SET SYSTEM or SET ARCHIVE commands. See the *WebSphere MQ Script (MQSC) Command Reference* for details of these commands.
- A new event queue SYSTEM.ADMIN.CONFIG.EVENT has been introduced, which can be used to hold configuration event messages.
- A new utility, CSQJUFMT, has been introduced, which is an active log preformat utility. This utility is used to format active log data sets before they are used by a queue manager. If the active log data sets are preformatted by the utility, log write performance is improved on the queue manager's first pass through the active logs.
- A new PAGEINFO function has been added to the CSQUTIL utility, which is used to extract information from page sets.
- A new queue manager attribute, EXPRYINT, has been introduced, which enables you to control how often queues are scanned for expired messages, which are then deleted.
- A new queue manager attribute, MAXUMSGS, has been introduced, which enables you to limit the number of messages that can be put within a single syncpoint, replacing the DEFINE MAXUMSGS command.
- You can now use up to 16 buffer pools.
- Two new DB2 tables have been created: CSQ.ADMIN_B_STRBACKUP and CSQ.OBJ_B_AUTHINFO.
- Two of your load libraries, thlqual.SCSQMVR1 and thlqual.SCSQMVR2, must now be in PDS-E format.
- Context security can now be implemented for each individual queue, or globally using generic profiles.
- You no longer need to define a SYSTEM.CHANNEL.REPLY.INFO queue for distributed queueing.
- This book now includes a brief definition of messages, queues, and other WebSphere MQ objects.

For full details of the new function that has been added in WebSphere MQ for z/OS, see "What's new in WebSphere MQ for z/OS" on page 167.

Part 1. Introduction

Chapter 1. Introduction

This book describes things that you need to know about WebSphere MQ for z/OS before you can install and use it on your z/OS system. It explains the concepts of WebSphere MQ, and gives you information to help you plan your WebSphere MQ subsystems. If you need to find out about WebSphere MQ on all platforms, see *MQSeries : An Introduction to Messaging and Queuing*.

This chapter introduces the concepts you need to understand, and directs you to more detailed explanations later in the book. It contains the following sections:

What is a message?

A *message* is a string of bytes that is meaningful to the applications that use it. Messages are used to transfer information from one application program to another (or to different parts of the same application). The applications can be running on the same platform, or on different platforms.

WebSphere MQ messages have two parts:

- The *application data*. The content and structure of the application data is defined by the application programs that use them.
- A *message descriptor*. The message descriptor identifies the message and contains additional control information, such as the type of message and the priority assigned to the message by the sending application.

What is a queue?

A *queue* is a data structure used to store messages until they are retrieved by an application.

Queues are managed by a *queue manager*. The queue manager is responsible for maintaining the queues it owns, storing all the messages it receives onto the appropriate queues, and retrieving the messages in response to application requests. The messages might be put on the queues by application programs, or by a queue manager as part of its operation.

What is a WebSphere MQ object?

Many of the tasks you carry out when using WebSphere MQ involve manipulating WebSphere MQ *objects*. In WebSphere MQ for z/OS the object types are queues, processes, authentication information objects, namelists, storage classes, Coupling Facility structures, channels, clusters, system objects, and queue manager security objects.

The manipulation or administration of objects includes:

- Starting and stopping queue managers
- Creating objects, particularly queues, for applications
- Working with channels to create communication paths to queue managers on other (remote) systems

The properties of an object are defined by its attributes. Some you can specify, others you can only view.

What is a queue manager?

Before you can let your application programs use WebSphere MQ on your z/OS system, you must install the WebSphere MQ for z/OS product and start a queue manager. The queue manager owns and manages the set of resources that are used by WebSphere MQ. These resources include:

- Page sets that hold the WebSphere MQ object definitions and message data
- Logs that are used to recover messages and objects in the event of queue manager failure
- Processor storage
- Connections through which different application environments (CICS, IMS, and Batch) can access the WebSphere MQ API
- The WebSphere MQ channel initiator, which allows communication between WebSphere MQ on your z/OS system and other systems

The queue manager has a name, and applications can connect to it using this name.

Figure 1 on page 5 illustrates a queue manager, showing connections to different application environments, and the channel initiator.

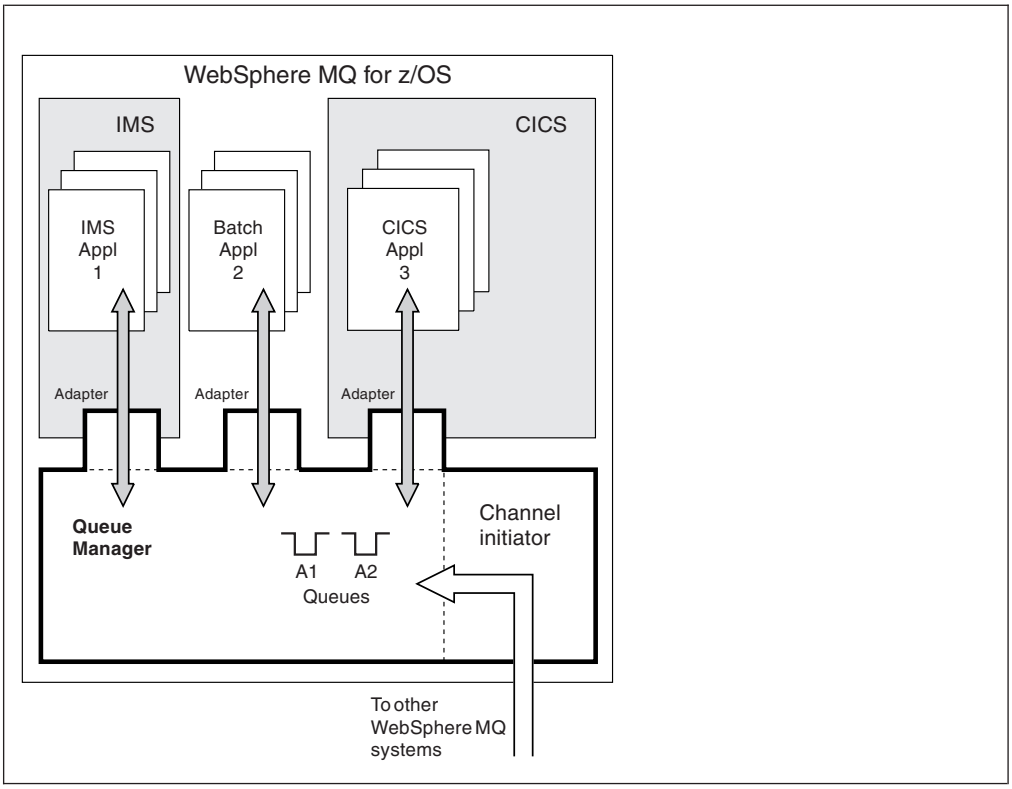

Figure 1. Overview of WebSphere MQ for z/OS

The queue manager subsystem

On z/OS, WebSphere MQ runs as a z/OS subsystem that is started at IPL time. Within the subsystem, the queue manager is started by executing a JCL procedure that specifies the z/OS data sets that contain information about the logs, and that hold object definitions and message data (the page sets). The subsystem and the queue manager have the same name, of up to four characters. All queue managers in your network must have unique names, even if they are on different systems, sysplexes, or platforms.

Shared queues

Queues can be *non-shared*, owned by and accessible to only one queue manager, or *shared*, owned by a *queue-sharing group*. (Queue-sharing groups are not available in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.) A queue-sharing group consists of a number of queue managers, running within a single z/OS sysplex, that can access the same WebSphere MQ object definitions and message data concurrently. Within a queue-sharing group, the shareable object definitions are stored in a shared DB2 database, and the messages are held inside one or more Coupling Facility structures (CF structures). The shared DB2 database and the Coupling Facility structures are resources that are owned by several queue managers.

Shared queues are discussed in Chapter 2, "Shared queues and queue-sharing groups", on page 13.

Page sets and buffer pools

When a message is put on to a non-shared queue, the queue manager stores the data on a page set in such a way that it can be retrieved when a subsequent operation gets a message from the same queue. If the message is removed from the queue, space in the page set that holds the data is subsequently freed for reuse. As the number of messages held on a queue increases, so the amount of space in the page set holding message data increases, and as the number of messages on a queue reduces, the space used in the page set reduces.

To reduce the overhead of writing data to and reading data from the page sets, the queue manager buffers the updates into processor storage. The amount of storage used to buffer the page set access is controlled through WebSphere MQ objects called *buffer pools*.

Page sets and buffer pools are discussed in Chapter 3, “Storage management”, on page 27.

Logging

Any changes to objects held on page sets, and operations on persistent messages, are recorded as log records. These log records are written to a log data set called the *active log*. The name and size of the active log data set is held in a data set called the *bootstrap data set* (BSDS).

When the active log data set fills up, the queue manager switches to another log data set so that logging can continue, and copies the content of the full active log data set to an *archive log* data set. Information about these actions, including the name of the archive log data set, is held in the bootstrap data set. Conceptually, there is a ring of active log data sets that the queue manager cycles through; when an active log is filled, the log data is off-loaded to an archive log, and the active log data set is available for reuse.

The log and bootstrap data sets are discussed in Chapter 4, “Logging”, on page 31.

Tailoring the queue manager environment

When the queue manager is started, a set of initialization parameters that control how the queue manager operates are read. In addition, data sets containing WebSphere MQ commands are read, and the commands they contain are executed. Typically, these data sets contain definitions of the system objects required for WebSphere MQ to run, and you can tailor these to define or initialize the WebSphere MQ objects necessary for your operating environment. When these data sets have been read, any objects defined by them are stored, either on a page set or in DB2.

Initialization parameters and system objects are discussed in Chapter 5, “Defining your system”, on page 41.

Recovery and restart

At any time during the operation of WebSphere MQ, there might be changes held in processor storage that have not yet been written to the page set. These changes are written out to the page set that is the least recently used by a background task within the queue manager.

If the queue manager terminates abnormally, the recovery phase of queue manager restart can recover the lost page set changes because persistent message data is

held in log records. This means that WebSphere MQ can recover persistent message data and object changes right up to the point of failure.

If a queue manager that is a member of a queue-sharing group encounters a Coupling Facility failure, the persistent messages on that queue can only be recovered if you have backed up your Coupling Facility structure.

Recovery and restart are discussed in Chapter 6, "Recovery and restart", on page 51.

Security

You can use an external security manager, such as RACF® to protect the resources that WebSphere MQ owns and manages from access by unauthorized users.

WebSphere MQ security is discussed in Chapter 7, "Security", on page 65.

Availability

There are several features of WebSphere MQ that are designed to increase system availability in the event of queue manager or communications subsystem failure. These features are discussed in Chapter 8, "Availability", on page 73.

Manipulating objects

When the queue manager is running, WebSphere MQ objects can be manipulated either through a z/OS console interface, or through an administration utility that uses ISPF services under TSO. Both mechanisms allow you to define, alter, or delete WebSphere MQ objects. You can also control and display the status of various WebSphere MQ and queue manager functions.

These facilities are discussed in Chapter 9, "Creating and managing objects", on page 77.

Monitoring and statistics

Several facilities are available to monitor your queue managers and channel initiators. You can also collect statistics for performance evaluation and accounting purposes.

These facilities are discussed in Chapter 10, "Monitoring and statistics", on page 91.

Application environments

When the queue manager has started, applications can connect to it and start using the WebSphere MQ API. These can be CICS, IMS, Batch, or WebSphere Application Server applications. WebSphere MQ applications can also access applications on CICS and IMS systems that are not aware of WebSphere MQ, using the CICS and IMS bridges. (CICS and IMS connections are not available in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

These facilities are discussed in Part 3, "WebSphere MQ and other products", on page 93.

For information about writing WebSphere MQ applications, see the following manuals:

- *WebSphere MQ Application Programming Guide*
- *WebSphere MQ Using C++*

- *WebSphere MQ Using Java*
- *WebSphere MQ Application Messaging Interface*

Internet Gateway

The Internet Gateway provides a bridge between the synchronous World Wide Web and asynchronous WebSphere MQ applications. With the gateway, Web server software and WebSphere MQ together provide an Internet-connected Web browser with access to WebSphere MQ applications. The gateway enables enterprises to take advantage of the low-cost access to global markets provided by the Internet, while benefiting from the robust infrastructure and assured message delivery of WebSphere MQ. Users interact with the gateway through HTML fill-out form POST requests; WebSphere MQ applications respond by returning HTML pages to the gateway, through a WebSphere MQ queue. The Internet Gateway supports the CGI and ICAPI Web server interfaces.

The Internet Gateway is an optional feature of WebSphere MQ, and is supplied on the WebSphere MQ product tape or cartridge. It is described in the *MQSeries Internet Gateway for MVS™ User's Guide*, which is available on the Web at:

`http://www.ibm.com/software/mqseries/library/manualsa/`

What is a channel initiator?

Note: The reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server does not support sending messages using channels between two separate queue managers. It supports only server connection channels to a single queue manager.

The *channel initiator* provides and manages resources that enable WebSphere MQ distributed queuing. WebSphere MQ uses *Message Channel Agents* (MCAs) to send messages from one queue manager to another.

To send messages from queue manager A to queue manager B, a *sending* MCA on queue manager A must set up a communications link to queue manager B. A *receiving* MCA must be started on queue manager B to receive messages from the communications link. This one-way path consisting of the sending MCA, the communications link, and the receiving MCA is known as a *channel*. The sending MCA takes messages from a transmission queue and sends them down a channel to the receiving MCA. The receiving MCA receives the messages and puts them on to the destination queues.

In WebSphere MQ for z/OS, the sending and receiving MCAs all run inside the channel initiator (the channel initiator is also known as the *mover*). The channel initiator runs as a z/OS address space under the control of the queue manager. There can only be a single channel initiator connected to a queue manager and it is run inside the same z/OS image as the queue manager. There can be thousands of MCA processes running inside the channel initiator concurrently.

Figure 2 on page 9 shows two queue managers within a sysplex. Each queue manager has a channel initiator and a local queue. Messages sent by queue managers on AIX® and Windows NT® are placed on the local queue, from where they are retrieved by an application. Reply messages are returned by a similar route.

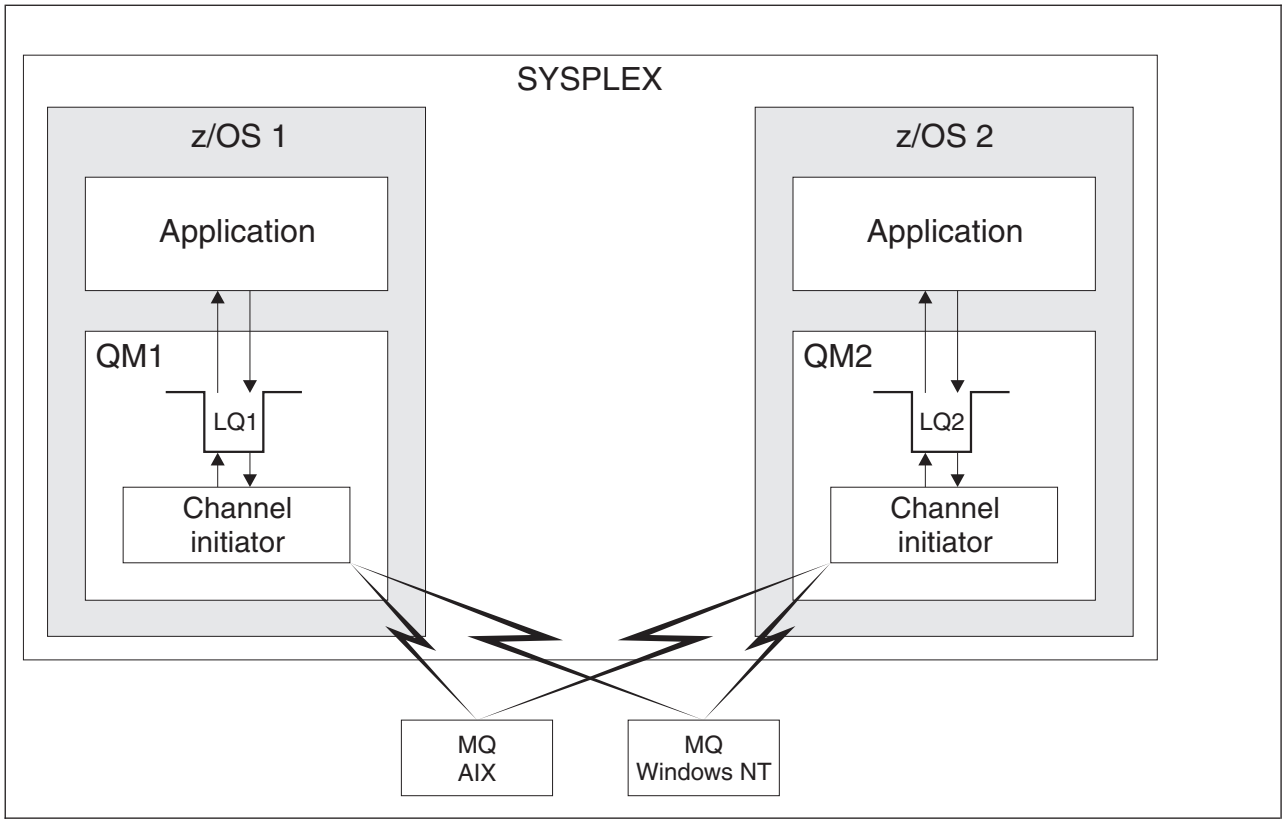

Figure 2. Communication between queue managers

The channel initiator also contains other processes concerned with the management of the channels. These include:

Listeners
These listen for inbound channel requests on a communications subsystem such as TCP, and start a named MCA when an inbound request is received.

Supervisor
This manages the channel initiator address space, for example it is responsible for restarting channels after a failure.

Name server
This is used to resolve TCP names into addresses.

SSL tasks
These are used to perform encryption and decryption and check certificate revocation lists.

Distributed queuing is described in the *WebSphere MQ Intercommunication* manual.

Queue manager clusters

You can group queue managers in a *cluster*. Queue managers in a cluster can make the queues that they host available to every other queue manager in the cluster. Any queue manager can send a message to any other queue manager in the same cluster without the need for many of the object definitions required for standard distributed queuing. Each queue manager in the cluster has a single transmission queue from which it can transmit messages to any other queue manager in the cluster.

Clusters are described in the *WebSphere MQ Queue Manager Clusters* manual.

Part 2. WebSphere MQ for z/OS concepts

Chapter 2. Shared queues and queue-sharing groups

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

This chapter describes how several queue managers can share the same queues and the messages on those queues. It discusses the following topics:

What is a shared queue?

A *shared queue* is a type of local queue. The messages on that queue can be accessed by one or more queue managers that are in a sysplex. The queue managers that can access the same set of shared queues form a group called a *queue-sharing group*.

Messages can be accessed by any queue manager

A shared queue can be accessed by any queue manager in the queue-sharing group. This means that you can put a message on to a shared queue on one queue manager, and get the same message from the queue from a different queue manager. This provides a rapid mechanism for communication within a queue-sharing group that does not require channels to be active between queue managers.

The messages on a shared queue are stored in the zSeries Coupling Facility. Figure 3 on page 14 shows three queue managers and a Coupling Facility, forming a queue-sharing group. All three queue managers can access the shared queue in the Coupling Facility.

An application can connect to any of the queue managers within the queue-sharing group. Because all the queue managers in the queue-sharing group can access all the shared queues, the application does not depend on the availability of a specific queue manager; any queue manager in the queue-sharing group can service the queue.

This gives greater availability because all the other queue managers in the queue-sharing group can continue processing the queue if one of the queue managers has a problem.

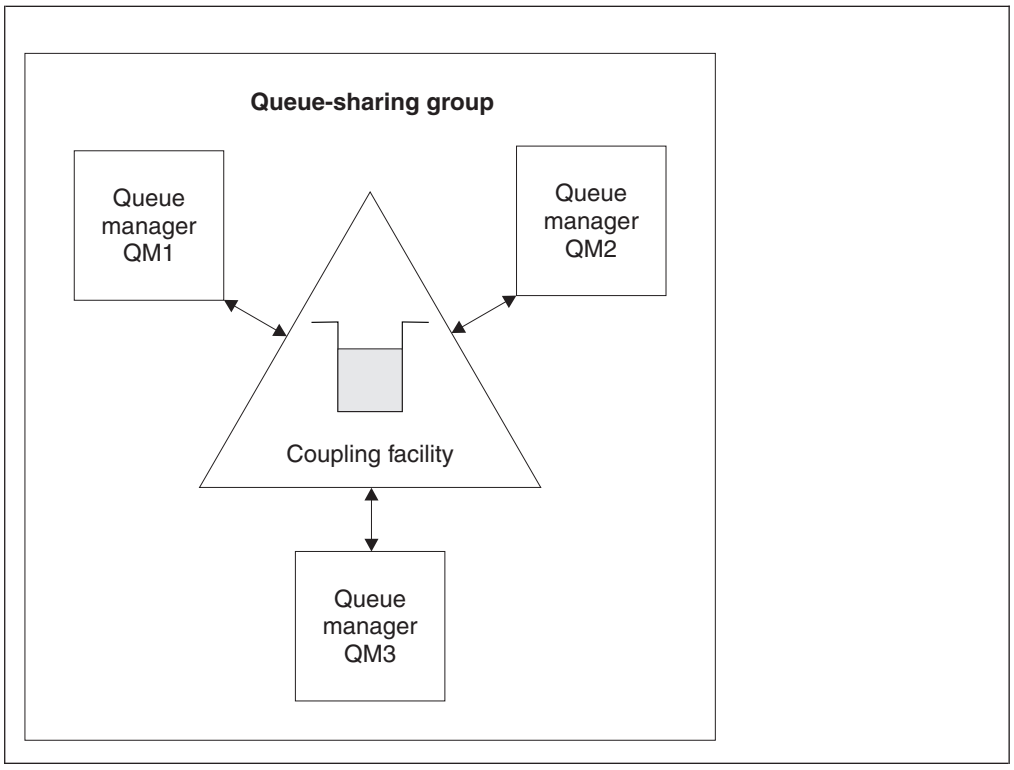

Figure 3. A queue-sharing group

Queue definition shared by all queue managers

The definition of a shared queue is stored in a DB2 shared database called the *shared repository*. Because of this, the queue need only be defined once and then it can be accessed by all the queue managers in the queue-sharing group. This means that there are fewer definitions to make.

By contrast, the definition of a non-shared queue is stored on page set zero of the queue manager that owns the queue (as described in "Page sets" on page 27).

You cannot define a shared queue if a queue with that name has already been defined on the page sets of the defining queue manager. Likewise, you cannot define a local version of a queue on the queue manager page sets if a shared queue with the same name already exists.

What is a queue-sharing group?

The group of queue managers that can access the same shared queues is called a queue-sharing group. Each member of the queue-sharing group has access to the same set of shared queues.

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.

Figure 4 illustrates a queue-sharing group that contains two queue managers. Each queue manager has a channel initiator and its own local page sets and log data sets.

Each member of the queue-sharing group must also connect to a DB2 system. The DB2 systems must all be in the same DB2 data-sharing group so that the queue managers can access the DB2 shared repository used to hold shared object definitions. These are definitions of any type of WebSphere MQ object (for example, queues and channels) that are defined once only and can then be used by any queue manager in the group. These are called *global* definitions and are described in "Private and global definitions" on page 77.

A particular data-sharing group can be referenced by more than one queue-sharing group. You specify the name of the DB2 subsystem and which data-sharing group a queue manager uses in the WebSphere MQ system parameters at startup.

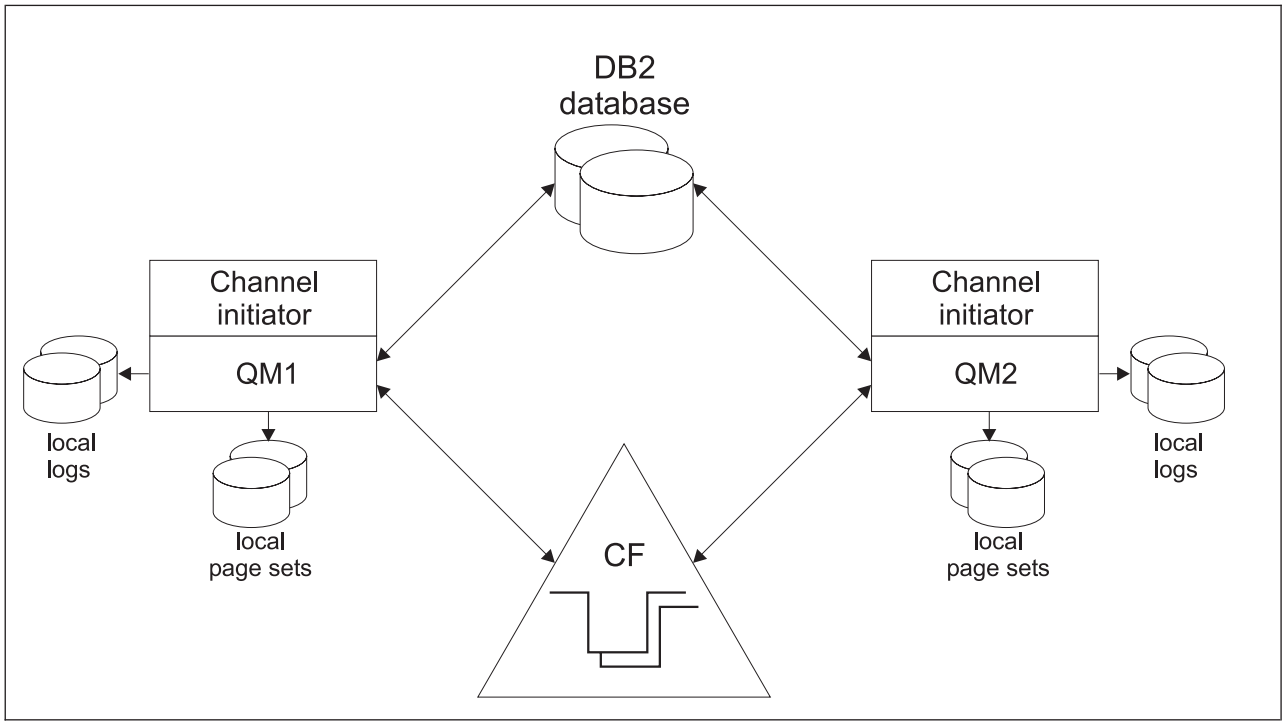

Figure 4. The components of queue managers in a queue-sharing group

When a queue manager has joined a queue-sharing group, it has access to the shared objects defined for that group, and can be used to define new shared objects within the group. If shared queues are defined within the group, this queue manager can be used to put messages to and get messages from those shared queues. Messages held on a shared queue can be retrieved by any queue manager in the group.

You can enter an MQSC command once, and have it executed on all queue managers within the queue-sharing group as if it had been entered at each queue manager individually. The *command scope* attribute is used for this. This attribute is described in "Directing commands to different queue managers" on page 79.

When a queue manager runs as a member of a queue-sharing group it must be possible to distinguish between WebSphere MQ objects defined privately to that

queue manager and WebSphere MQ objects defined globally that are available to all queue managers in the queue-sharing group. The *queue-sharing group disposition* attribute is used for this. This attribute is described in "Private and global definitions" on page 77.

You can define a single set of security profiles that control access to WebSphere MQ objects anywhere within the group. This means that the number of profiles you have to define is greatly reduced.

A queue manager can belong to one queue-sharing group only, and all queue managers in the group must be in the same sysplex. You specify which queue-sharing group the queue manager belongs to in the system parameters at startup.

Where are shared queue messages held?

The messages in shared queues are stored on list structures in the zSeries Coupling Facility. They can be accessed by many queue managers in the same sysplex. All queue managers also maintain their own logs and page sets (as shown in Figure 4 on page 15) to use non-shared local queues, and store definitions of private objects on page set zero. Messages that are put on to shared queues are not stored on page sets.

MQPUT and MQGET operations on persistent messages (on both shared and non-shared queues) are recorded on the queue manager log. This minimizes the risk of data loss in the event of a Coupling Facility or page set failure.

The Coupling Facility

The messages held on shared queues are actually stored inside a Coupling Facility. The Coupling Facility lies outside any of the z/OS images in the sysplex and is typically configured to run on a different power supply. The Coupling Facility is therefore resilient to software failures and can be configured so that it is resilient to hardware failures or power-outages. This means that messages stored in the Coupling Facility are highly available.

WebSphere MQ uses Coupling Facility list structures to store messages. This means that the maximum length for messages on a shared queue is 63 KB.

Each Coupling Facility list structure used by WebSphere MQ is dedicated to a specific queue-sharing group, but a Coupling Facility can hold structures for more than one queue-sharing group. Queue managers in different queue-sharing groups cannot share data. Up to 32 queue managers in a queue-sharing group can connect to a Coupling Facility list structure at the same time.

A single Coupling Facility list structure can contain up to 512 shared queues. The amount of message data is limited by the size of the list structure. The size of the list structure is restricted by the following factors:

- It must lie within a single Coupling Facility.
- It might share the available Coupling Facility storage with other structures for WebSphere MQ and other products.

The CF structure object

The queue manager's use of a Coupling Facility structure is specified in a CF structure (CFSTRUCT) WebSphere MQ object.

When using z/OS commands or definitions relating to a Coupling Facility structure, the first four characters of the queue-sharing group name are required. However, a WebSphere MQ CFSTRUCT object always exists within a single queue-sharing group, and so its name does not include the first four characters of the queue-sharing group name. For example, CFSTRUCT(MYDATA) defined in queue-sharing group starting with SQ03 would use Coupling Facility list structure SQ03MYDATA.

Backup and recovery

Coupling Facility list structures can be backed up using the WebSphere MQ command BACKUP CFSTRUCT. This puts a copy of the persistent messages currently within the CF structure onto the active log data set of the queue manager making the backup, and writes a record of the backup to DB2.

In the event of a Coupling Facility failure, you can use the WebSphere MQ command RECOVER CFSTRUCT. This uses the backup record from DB2 to locate and restore persistent messages from the backup of the CF structure. Any activity since the last backup is replayed using the logs of all the queue managers in the queue-sharing group, and the CF structure is then restored up to the point before the failure.

For more information about the BACKUP CFSTRUCT and RECOVER CFSTRUCT commands, see the *WebSphere MQ Script (MQSC) Command Reference*.

Advantages of using shared queues

The shared queue architecture, where cloned servers pull work from a single shared queue, has some very useful properties:

- It is scalable, by adding new instances of the server application, or even adding a new z/OS image with a queue manager (in the queue-sharing group) and a copy of the application.
- It is highly available.
- It naturally performs *pull* workload balancing, based on the available processing capacity of each queue manager in the queue-sharing group.

High availability

The following examples illustrate how a shared queue can be used to increase application availability.

Consider a WebSphere MQ scenario where client applications running in the network want to make requests of server applications running on z/OS. The client application constructs a request message and places it on a request queue. The client then waits for a reply from the server, sent to the reply-to queue named in the message descriptor of the request message.

WebSphere MQ manages the transportation of the request message from the client machine to the server's input queue on z/OS and of the server's response back to the client. By defining the server's input queue as a shared queue, any messages put to the queue can be retrieved on any queue manager in the queue-sharing group. This means that you can configure a queue manager on each z/OS image in the sysplex and, by connecting them all to the same queue-sharing group, any one of them can access messages on the server's input queue.

Messages on the server's input queue are still available, even if one of the queue managers terminates abnormally or has to be stopped for administrative reasons. In fact, an entire z/OS image can be taken off-line and the messages will still be available.

To take advantage of this availability of messages on a shared queue, run an instance of the server application on each z/OS image in the sysplex to provide higher server application capacity and availability, as shown in Figure 5.

One instance of the server application retrieves a request message from the shared queue and, based on the content, performs its processing, producing a result to be sent back to the client as a WebSphere MQ message. The response message is destined for the reply-to queue and reply-to queue manager named in the message descriptor of the request message.

There are a number of options that can be used to configure the return path; these are considered in "Distributed queuing and queue-sharing groups" on page 19.

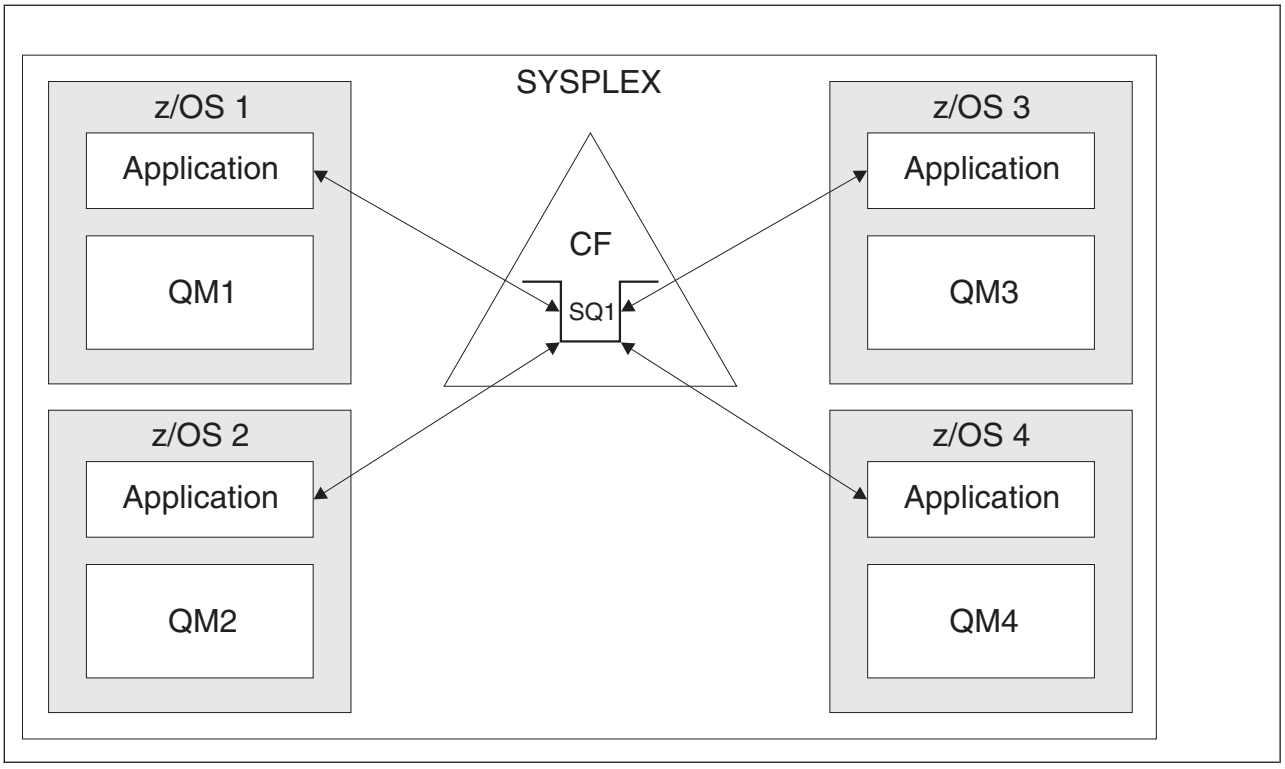

Figure 5. Multiple instances of an application servicing a shared queue

Peer recovery

To further enhance the availability of messages in a queue-sharing group, WebSphere MQ detects if another queue manager in the group disconnects from the Coupling Facility abnormally and completes units of work for that queue manager that are still pending, where possible. This is known as *peer recovery.*

Suppose a queue manager terminates abnormally at a point where an application has retrieved a request message from a queue in syncpoint, but has not yet put the response message or committed the unit of work. Another queue manager in the queue-sharing group detects the failure, and backs out the in-flight units of work being performed on the failed queue manager. This means that the request

message is put back on to the request queue and is available for one of the other server instances to process, without waiting for the failed queue manager to restart.

If WebSphere MQ cannot resolve a unit of work automatically, you can resolve the shared portion manually to enable another queue manager in the queue-sharing group to continue processing that work.

Distributed queuing and queue-sharing groups

To complement the high availability of messages on shared queues, the distributed queuing component of WebSphere MQ has additional functions to provide the following:

- Higher availability to the network.
- Increased capacity for inbound network connections to the queue-sharing group.

Figure 6 illustrates distributed queuing and queue-sharing groups. It shows two queue managers within a sysplex, both of which belong to the same queue-sharing group. They can both access shared queue SQ1. Queue managers in the network (on AIX and Windows NT for example) can put messages onto this queue through the channel initiator of either queue manager. Cloned applications on both queue managers service the queue.

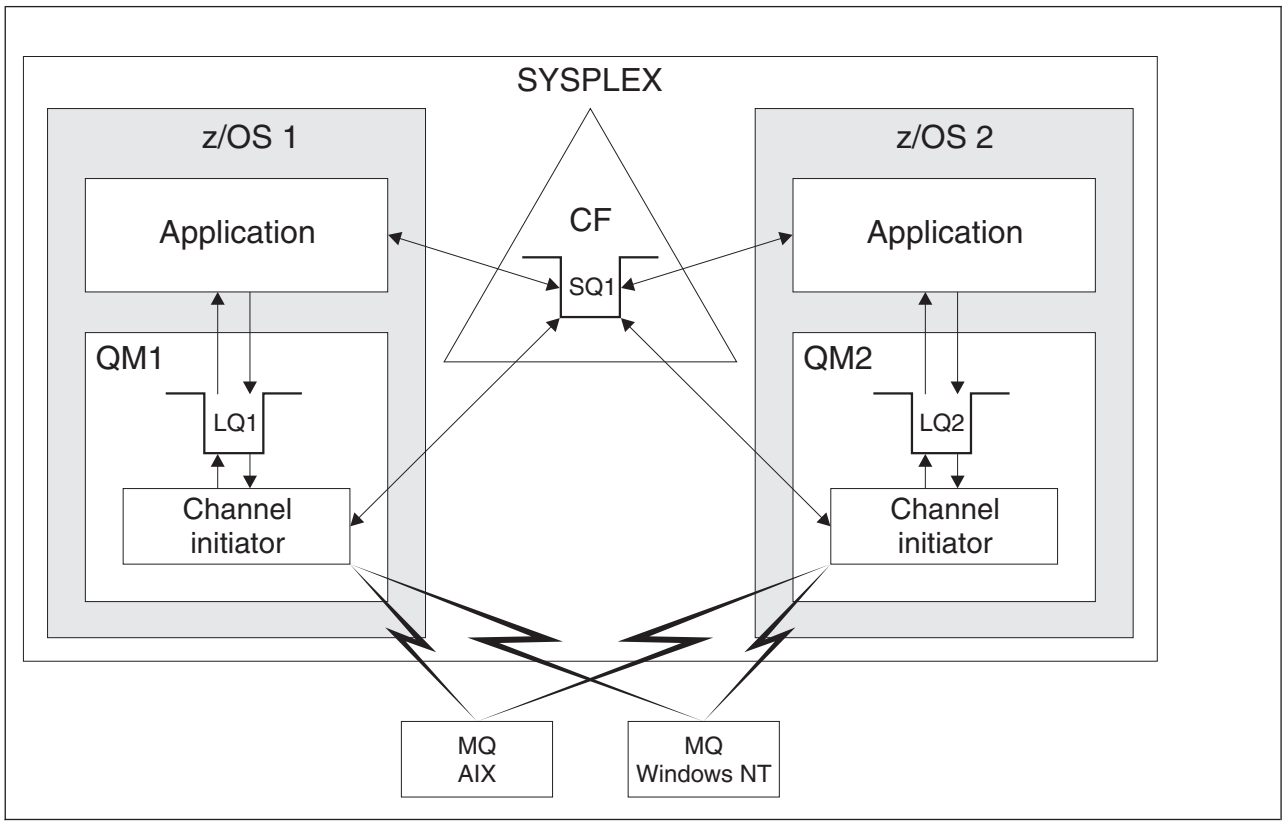

Figure 6. Distributed queuing and queue-sharing groups

Shared channels

A number of networking products provide a mechanism to hide server failures from the network, or to balance inbound network requests across a set of eligible servers. These include:

- VTAM® generic resources
- TCP/IP Domain Name System (DNS)

The channel initiator takes advantage of these products to exploit the capabilities of shared queues.

Shared inbound channels

Each channel initiator in the queue-sharing group starts an additional listener task to listen on a *generic port*. This generic port is made available to the network through one of the technologies mentioned above. This means that an inbound network attach request for the generic port can be dispatched to any one of the listeners in the queue-sharing group that are listening on the generic port.

A channel can only be started on the channel initiator to which the inbound attach is directed if the channel initiator has access to a channel definition for a channel with that name. A channel definition can be defined to be private to a queue manager or stored on the shared repository and available anywhere (a global definition). This means that a channel definition can be made available on any channel initiator in the queue-sharing group by defining it as a global definition.

There is an additional difference when starting a channel through the generic port; channel synchronization is with the queue-sharing group and not with an individual queue manager. For example, consider a client starting a channel through the generic port. When the channel first starts, it might start on queue manager QM1 and messages flow. If the channel stops and is restarted on queue manager QM2, information about the number of messages that have flowed is still correct because the synchronization is with the queue-sharing group.

An inbound channel started through the generic port can be used to put messages to any queue. The client does not know whether the target queue is shared or not. If the target queue is a shared queue, the client connects through any available channel initiator in a load-balanced fashion and the messages are put to the shared queue. If the target queue is not a shared queue, the messages might be put on any queue in the queue-sharing group with that name (the environment is one of replicated local queues), and the name of the queue determines the function, regardless of the hosting queue manager.

Shared outbound channels

An outbound channel is considered to be a shared channel if it is taking messages from a shared transmission queue. If it is shared, it holds synchronization information at queue-sharing group level. This means that the channel can be restarted on a different queue manager and channel initiator instance within the queue-sharing group if the communications subsystem, channel initiator, or queue manager fails. Restarting failed channels in this way is a feature of shared channels called *peer channel recovery*.

Workload balancing: An outbound shared channel is eligible for starting on any channel initiator within the queue-sharing group, provided that you have not specified that you want it to be started on a particular channel initiator. The channel initiator selected by WebSphere MQ is determined using the following criteria:

- Is the communications subsystem required currently available to the channel initiator?
- Is a DB2 connection available to the channel initiator?
- Which channel initiator has the lowest current workload? The workload includes channels that are active and retrying.

Shared channel summary

Shared channels differ from private channels in the following ways:

Private channel

Tied to a single channel initiator.

- Outbound channel uses a local transmission queue.
- Inbound channel started through a local port.
- Synchronization information held in SYSTEM.CHANNEL.SYNCQ queue.

Shared Channel

Workload balanced with high availability.

- Outbound channel uses a shared transmission queue.
- Inbound channel started through a generic port.
- Synchronization information held in SYSTEM.QSG.CHANNEL.SYNCQ queue.
- Shared outbound channels have a maximum message length of 63 KB.

You specify whether a channel is private or shared when you start the channel. A shared channel can be started by triggering in the same way as a private channel. However, when a shared channel is started, WebSphere MQ performs workload balancing and starts the channel on the most appropriate channel initiator within the queue-sharing group. (If required, you can specify that a shared channel is to be started on a particular channel initiator.)

Shared channel status

The channel initiators in a queue-sharing group maintain a shared channel-status table in DB2. This records which channels are active on which channel initiators. The shared channel-status table is used if there is a channel initiator or communications system failure. It indicates which channels need to be restarted on a different channel initiator in the queue-sharing group.

Intra-group queuing

You can perform fast message transfer between queue managers in a queue-sharing group without defining channels. This uses a system queue called the SYSTEM.QSG.TRANSMIT.QUEUE, which is a shared transmission queue. Each queue manager in the queue-sharing group starts a task called the intra-group queuing agent, which waits for messages to arrive on this queue that are destined for their queue manager. When such a message is detected, it is removed from the queue and placed on the correct destination queue.

Standard name resolution rules are used but, if intra-group queuing is enabled and the target queue manager is within the queue-sharing group, the SYSTEM.QSG.TRANSMIT.QUEUE is used to transfer the message to the correct destination queue manager instead of using a transmission queue and channel.

Intra-group queuing is enabled through a queue manager attribute at startup. It can be used only to move messages with a message length of up to 63 KB, including the transmission-queue header (63 KB is the maximum message length for shared queues). Intra-group queuing moves non-persistent messages outside syncpoint, and persistent messages within syncpoint. If it finds a problem delivering messages to the target queue, intra-group queuing tries to put them to the dead-letter queue. If the dead-letter queue is full or undefined, non-persistent messages are discarded, but persistent messages are backed out and returned to the SYSTEM.QSG.TRANSMIT.QUEUE, and the IGQ agent tries to deliver the messages until it is successful.

Shared queues and queue-sharing groups

An inbound shared channel that receives a message destined for a queue on a different queue manager in the queue-sharing group can use intra-group queuing to *hop* the message to the correct destination.

Clusters and queue-sharing groups

You can make your shared queues available to a cluster in a single definition. To do this you specify the name of the cluster when you define the shared queue.

Users in the network see the shared queue as being hosted by each queue manager within the queue-sharing group (the shared queue is not advertised as being hosted by the queue-sharing group). Clients can start sessions with any members of the queue-sharing group to put messages to the same shared queue.

Figure 7 shows how members of a cluster can access a shared queue through any member of the queue-sharing group.

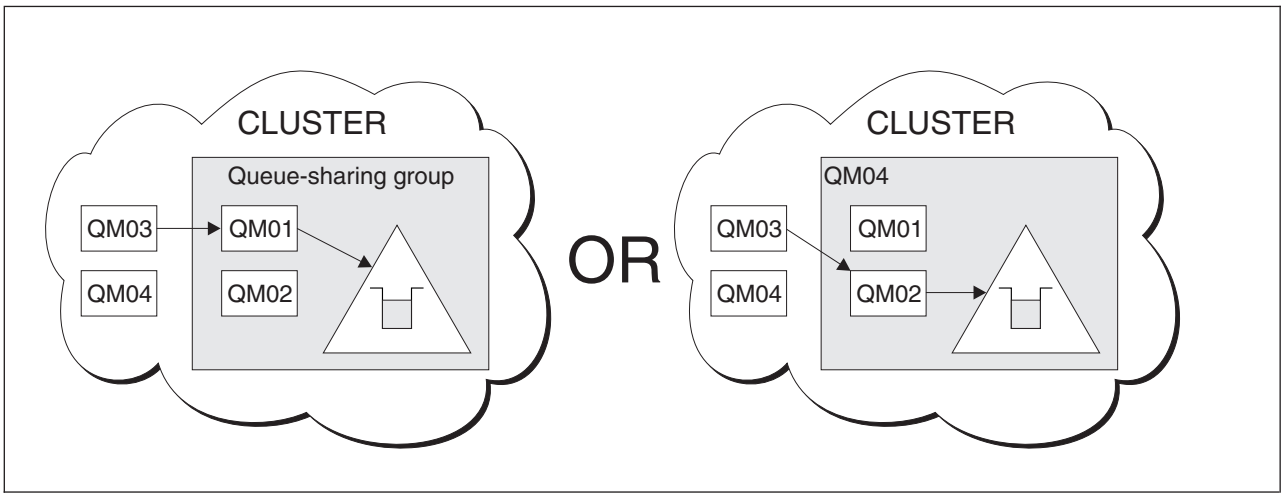

Figure 7. A queue-sharing group as part of a cluster

Application programming with shared queues

This section discusses some of the factors you need to take into account when designing new applications that will use shared queues, and when migrating existing applications to the shared queue environment.

Serializing your applications

Certain types of applications might have to ensure that messages are retrieved from a queue in exactly the same order as they arrived on the queue. For example, if WebSphere MQ is being used to shadow database updates on to a remote system, a message describing the update to a record must be processed after a message describing the insert of that record. In a local queuing environment, this is often achieved by the application that is getting the messages opening the queue with the MQOO_INPUT_EXCLUSIVE option, thus preventing any other getting application from processing the queue at the same time.

WebSphere MQ allows applications to open shared queues exclusively in the same way. However, if the application is working from a partition of a queue (for example, all database updates are on the same queue, but those for table A have a correlation identifier of A, and those for table B a correlation identifier of B), and

applications want to get messages for table A updates and table B updates concurrently, the simple mechanism of opening the queue exclusively is not possible.

If this type of application is to take advantage of the high availability of shared queues, you might decide that another instance of the application that accesses the same shared queues, running on a secondary queue manager, will take over if the primary getting application or queue manager fails.

If the primary queue manager fails, two things happen:

- Shared queue peer recovery ensures that any incomplete updates from the primary application are completed or backed out.
- The secondary application takes over processing the queue.

The secondary application might start before all the incomplete units of work have been dealt with, which could lead to the secondary application retrieving the messages out of sequence. To solve this type of problem, the application can choose to be a *serialized application*.

A serialized application uses the **MQCONNX** call to connect to the queue manager, specifying a connection tag when it connects that is unique to that application. Any units of work performed by the application are marked with the connection tag. WebSphere MQ ensures that units of work within the queue-sharing group with the same connection tag are serialized (according to the serialization options on the **MQCONNX** call).

This means that, if the primary application uses the **MQCONNX** call with a connection tag of `Database shadow retriever`, and the secondary takeover application attempts to use the **MQCONNX** call with an identical connection tag, the secondary application will not be able to connect to the second WebSphere MQ until any outstanding primary units of work have been completed, in this case by peer recovery.

You should consider using the serialized application technique for applications that depend on the exact sequence of messages on a queue. In particular:

- Applications that must not restart after an application or queue manager failure until all commit and back-out operations for the previous execution of the application are complete.

  In this case, the serialized application technique is only applicable if the application works in syncpoint.
- Applications that must not start while another instance of the same application is already running.

  In this case, the serialized application technique is only required if the application cannot open the queue for exclusive input.

Note: WebSphere MQ only guarantees to preserve the sequence of messages when certain criteria are met. These are described in the description of the **MQGET** call in the *WebSphere MQ Application Programming Reference*.

Applications that are not suitable for use with shared queues

Some features of WebSphere MQ are not supported when you are using shared queues, so applications that use these features are not suitable for the shared queue environment. You should consider the following points when designing your shared queue applications:

- Messages on shared queues cannot be greater than 63 KB in size. Because Coupling Facility storage is limited, you must also consider the number of messages to be generated by the application to ensure that the messages will not fill the queue. However, remember that you can monitor the queue and start more versions of the application on different queue managers to service it if this is a problem.
- Queue indexing is limited for shared queues. If you want to use the message identifier or correlation identifier to select the message you want to get from the queue, the queue must have the correct index defined. If you do not define a queue index, applications can only get the next available message.
- You cannot use temporary dynamic queues as shared queues. You can use permanent dynamic queues however. The models for shared dynamic queues have a DEFTYPE of SHAREDYN (shared dynamic) although they are created and destroyed in the same way as PERMDYN (permanent dynamic) queues.

Deciding whether to share non-application queues

There are queues other than application queues that you might want to consider sharing:

Initiation queues

If you define a shared initiation queue, you do not need to have a trigger monitor running on every queue manager in the queue-sharing group, as long as there is at least one trigger monitor running. (You can also use a shared initiation queue even if there is a trigger monitor running on each queue manager in the queue-sharing group.)

If you have a shared application queue and use the trigger type of EVERY (or a trigger type of FIRST with a small trigger interval, which behaves like a trigger type of EVERY) your initiation queue should always be a shared queue. For more information about when to use a shared initiation queue, see Table 1 on page 25.

Dead-letter queue

Do not define your dead-letter queue as a shared queue. This is because shared queues cannot hold messages with a size greater than 63 KB.

SYSTEM.* queues

You can define the SYSTEM.ADMIN.* queues used to hold event messages as shared queues. This can be useful to check load balancing if an exception occurs. Each event message created by WebSphere MQ contains a correlation identifier indicating which queue manager produced it.

You must define the SYSTEM.QSG.* queues used for shared channels and intra-group queuing as shared queues.

You can also change the definition of the SYSTEM.DEFAULT.LOCAL.QUEUE to be shared, or define your own default shared queue definition. This is described in "Defining system objects" on page 41.

You cannot define any of the other SYSTEM.* queues as shared queues.

Migrating your existing applications to use shared queues

Migrating your existing queues to shared queues is described in the *WebSphere MQ for z/OS System Administration Guide*.

When you are migrating your existing applications, you should consider the following things, which might work in a slightly differently way in the shared queue environment.

Reason Codes
When you are migrating your existing applications to use shared queues, remember to check for the new reason codes that can be issued.

MQINQ
When you use the **MQINQ** call to display information about a shared queue, the values of the number of **MQOPEN** calls that have the queue open for input and output relate only to the queue manager that issued the call. No information is produced about other queue managers in the queue-sharing group that have the queue open.

Triggering
If you are using a shared application queue, triggering works on committed messages only (on a non-shared application queue, triggering works on all messages).

If you use triggering to start applications, you might want to use a shared initiation queue. Table 1 describes what you need to consider when deciding which type of initiation queue to use.

Table 1. When to use a shared-initiation queue

	Non-shared application queue	**Shared application queue**
Non-shared initiation queue	OK.	If you are using trigger type of FIRST or DEPTH, you can use a non-shared initiation queue with a shared application queue. There is the possibility of extra trigger messages being generated, but this setup is good for triggering long-running applications (like the CICS bridge) and provides high availability. For trigger type FIRST or DEPTH, a trigger message triggers an instance of the application on every queue manager that is running a trigger monitor and that does not already have the application queue open for input.
Shared initiation queue	Do not use a shared initiation queue with a non-shared application queue.	If you have a shared application queue that has trigger type EVERY, use a shared initiation queue or you will lose trigger messages. For trigger type FIRST or DEPTH, one trigger message is generated by each queue manager that has the named initiation queue open for input. If the initiation queue is defined as a local queue, one trigger message is available to any trigger monitors running on that queue manager against the queue.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 2. Where to find more information about shared queues and queue-sharing groups

Topic	Where to look
Queue-sharing group recovery	Chapter 6, "Recovery and restart", on page 51
Queue-sharing group security	Chapter 7, "Security", on page 65
Private and global object definitions Directing commands to different queue managers	Chapter 9, "Creating and managing objects", on page 77
Planning your Coupling Facility environment	"Defining Coupling Facility resources" on page 131
Planning your DB2 environment	"Planning your DB2 environment" on page 134
Setting up your shared queues System parameters	*WebSphere MQ for z/OS System Setup Guide*
Utility programs Migrating queues	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*
WebSphere MQ clusters	*WebSphere MQ Queue Manager Clusters*
WebSphere MQ distributed queuing Channel name resolution	*WebSphere MQ Intercommunication*
Writing applications	*WebSphere MQ Application Programming Guide*
MQCONNX call	*WebSphere MQ Application Programming Reference*

Chapter 3. Storage management

This chapter discusses how WebSphere MQ for z/OS manages storage. It contains the following sections:

Page sets

A *page set* is a linear VSAM data set (LDS) that has been specially formatted to be used by WebSphere MQ. Page sets are used to store most messages and object definitions.

The exceptions to this are global definitions, which are stored in a shared repository on DB2, and the messages on shared queues. These are not stored on the queue manager page sets. Shared queues are discussed in Chapter 2, "Shared queues and queue-sharing groups", on page 13, and global definitions are discussed in "Private and global definitions" on page 77.

WebSphere MQ page sets can be up to 4 GB in size. Each page set is identified by a page set identifier (PSID), an integer in the range 00 through 99. Each queue manager must have its own page sets.

WebSphere MQ uses page set zero (PSID=00) to store object definitions and other important information relevant to the queue manager. For normal operation of WebSphere MQ it is essential that page set zero does not become full, so do not use it to store messages.

To improve the performance of your system, you should also separate short lived messages from long lived messages by placing them on different page sets.

Page sets must be formatted so WebSphere MQ provides a FORMAT utility for this. This is described in the *WebSphere MQ for z/OS System Administration Guide*. Page sets must also be defined to the WebSphere MQ subsystem, this is described in the *WebSphere MQ for z/OS System Setup Guide*.

If you define secondary extents for your page sets, WebSphere MQ for z/OS expands a page set dynamically if it becomes full. WebSphere MQ continues to expand the page set if required until 123 logical extents exist, provided that there is sufficient disk storage space available. The extents can span volumes if the LDS is so defined, however, WebSphere MQ cannot expand the page sets beyond 4 GB.

You cannot use page sets from one WebSphere MQ queue manager on a different WebSphere MQ queue manager, or change the queue manager name. If you want to transfer the data from one queue manager to another, you must unload all the objects and messages from the first queue manager and reload them onto another.

Storage classes

A *storage class* maps one or more queues to a page set. This means that messages for that queue are stored on that page set.

Storage classes allow you to control where non-shared message data is stored for administrative, data set space and load management, or application isolation purposes. Storage classes can also be used to define the XCF group and member name of an IMS region if you are using the IMS bridge (described in Chapter 12, "WebSphere MQ and IMS", on page 107).

Shared queues do not use storage classes to obtain a page set mapping because the messages on them are not stored on page sets.

How storage classes work

- You define a storage class, using the DEFINE STGCLASS command, specifying a page set identifier (PSID).
- When you define a queue, you specify the storage class in the STGCLASS attribute.

In the following example, the local queue QE5 is mapped to page set 21 through storage class ARC2.

```
DEFINE STGCLASS(ARC2) PSID(21)
DEFINE QLOCAL(QE5) STGCLASS(ARC2)
```

This means that messages that are put on the queue QE5 are stored on page set 21 (if they stay on the queue long enough to be written to DASD).

More than one queue can use the same storage class, and you can define as many storage classes as you like. For example, you can extend the previous example to include more storage class and queue definitions, as follows:

```
DEFINE STGCLASS(ARC1) PSID(05)
DEFINE STGCLASS(ARC2) PSID(21)
DEFINE STGCLASS(MAXI) PSID(05)
DEFINE QLOCAL(QE1) STGCLASS(ARC1) ...
DEFINE QLOCAL(QE2) STGCLASS(ARC1) ...
DEFINE QLOCAL(QE3) STGCLASS(MAXI) ...
DEFINE QLOCAL(QE4) STGCLASS(ARC2) ...
DEFINE QLOCAL(QE5) STGCLASS(ARC2) ...
```

In Figure 8 on page 29, both storage classes ARC1 and MAXI are associated with page set 05. Therefore, the queues QE1, QE2, and QE3 are mapped to page set 05. Similarly, storage class ARC2 associates queues QE4 and QE5 with page set 21.

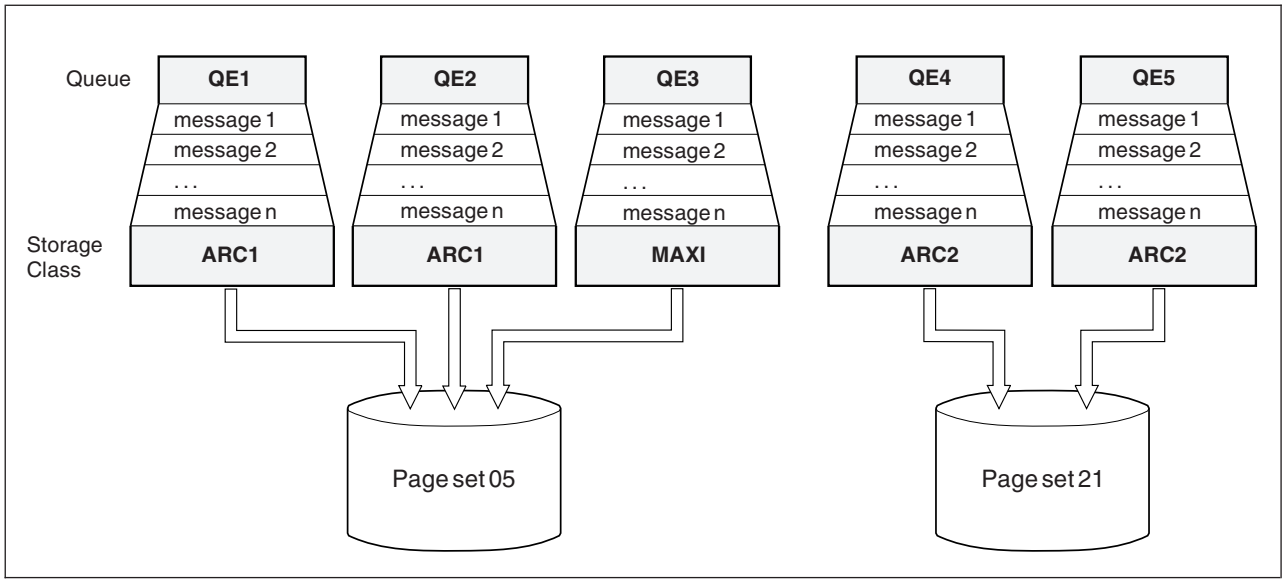

Figure 8. Mapping queues to page sets through storage classes

If you define a queue without specifying a storage class, WebSphere MQ uses a default storage class.

If a message is put on a queue that names a nonexistent storage class, the application receives an error. You must alter the queue definition to give it an existing storage class name, or create the storage class named by the queue.

A storage class can only be changed when:

- All queues that use this storage class are empty, and have no uncommitted activity.
- All queues that use this storage class are closed.

Buffers and buffer pools

For efficiency, WebSphere MQ uses a form of caching whereby messages (and object definitions) are stored temporarily in buffers before being stored in page sets on DASD. Short-lived messages, that is, messages that are retrieved from a queue shortly after they are received, might only ever be stored in the buffers. However, this is all transparent to the user because the buffers are controlled by a buffer manager, which is a component of WebSphere MQ.

The buffers are organized into *buffer pools*. You can define up to 16 buffer pools (0 through 15) for each queue manager; you are recommended to use the minimal number of buffer pools consistent with the object and message type segregation outlined in Figure 9 on page 30, and any data isolation requirements your application might have. Each buffer is 4 KB long. The maximum number of buffers is determined by the amount of storage available in the queue manager address space; do not use more than about 70% of the space for buffers. Usually, the more buffers you have, the more efficient the buffering and the better the performance of WebSphere MQ.

Figure 9 on page 30 shows the relationship between messages, buffers, buffer pools, and page sets. A buffer pool is associated with one or more page sets; each page set is associated with a single buffer pool.

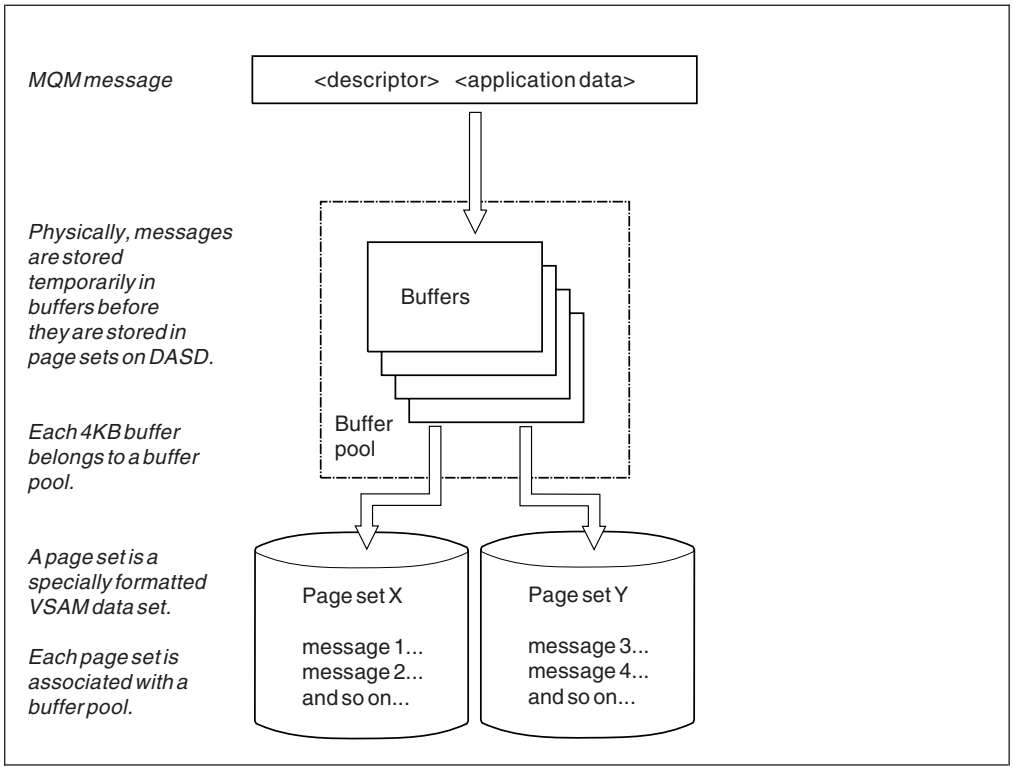

Figure 9. Buffers, buffer pools, and page sets

You specify the number of buffers in a pool with the DEFINE BUFFPOOL command. This command is described in the *WebSphere MQ Script (MQSC) Command Reference* manual.

For performance reasons, do not put messages and object definitions in the same buffer pool. Use one buffer pool (say number zero) exclusively for page set zero, where the object definitions are kept. Similarly, keep short-lived messages and long-lived messages in different buffer pools and therefore on different page sets, and in different queues.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 3. Where to find more information about storage management

Topic	Where to look
How much storage you need	Chapter 15, "Planning your storage requirements", on page 119
How large to make your page sets and buffer pools	Chapter 16, "Planning your page sets and buffer pools", on page 123
Defining page sets	*WebSphere MQ for z/OS System Setup Guide*
Formatting page sets Utility programs	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*

Chapter 4. Logging

WebSphere MQ maintains *logs* of data changes and significant events as they occur. The *bootstrap data set* (BSDS) stores information about the data sets that contain the logs.

This chapter contains the following sections:

The log does not contain information for statistics, traces, or performance evaluation. The statistical and monitoring information that WebSphere MQ collects is discussed in Chapter 10, "Monitoring and statistics", on page 91.

What logs are

WebSphere MQ records all significant events as they occur in an *active log*. The log contains the information needed to recover:
- Persistent messages
- WebSphere MQ objects, such as queues
- The WebSphere MQ queue manager

Archiving

Because the active log is finite, WebSphere MQ copies the contents of each log data set periodically to an *archive log*, which is normally a data set on a direct access storage device (DASD) or a magnetic tape. If there is a subsystem or transaction failure, WebSphere MQ uses the active log and, if necessary, the archive log for recovery.

The archive log can contain up to 1000 sequential data sets. Each data set can be cataloged using the z/OS integrated catalog facility (ICF).

Archiving is an essential component of WebSphere MQ recovery. If a unit of recovery is a long-running one, log records within that unit of recovery might be found in the archive log. In this case, recovery requires data from the archive log. However, if archiving is switched off, the active log with new log records wraps, overwriting earlier log records. This means that WebSphere MQ might not be able to back out the unit of recovery and messages might be lost. The queue manager then terminates abnormally.

Therefore, in a production environment, **never switch archiving off**. If you do, you run the risk of losing data after a system or transaction failure. Only if you are running in a test environment can you consider switching archiving off. If you need to do this, use the CSQ6LOGP macro, which is described in the *WebSphere MQ for z/OS System Setup Guide*.

To help prevent problems with unplanned long-running units of work, WebSphere MQ issues a message (CSQJ160I or CSQJ161I) if a long-running unit of work is detected during active log offload.

Dual logging

You can configure WebSphere MQ to run with either *single logging* or *dual logging*. With single logging, log records are written once to an active log data set. With dual logging, each log record is written to two different active log data sets. Dual logging minimizes the likelihood of data loss problems during restart. If possible, the two log data sets should be on separate volumes. This reduces the risk of them both being lost if one of the volumes is corrupted or destroyed. If both copies of the log are lost, the probability of data loss is high.

Note: You should always use dual logging and dual BSDSs rather than dual writing to DASD.

Single logging gives you 2 through 53 active log data sets, whereas dual logging gives you 4 through 106. Each active log data set is a single-volume, single-extent VSAM linear data set (LDS).

Although the minimum number of log data sets required is two, in practice you should have at least three, and on a busy system you might need more. This is to allow time for each log data set to be copied to archive before it is reused in the active log cycle.

Log data

The log can contain up to 140 million million ($1.4*10^{14}$) bytes. Each byte can be addressed by its offset from the beginning of the log, and that offset is known as its *relative byte address* (RBA).

The RBA is referenced by a 6-byte field giving a total addressable range of 2^{48} bytes. However, once WebSphere MQ detects that the used range is beyond x'700 000 000 000' (that is, more than $1.2*10^{14}$ bytes of the log have been used), messages CSQI045, CSQI046 and CSQI047 are issued, warning you to reset the log RBA. Do not allow the log RBA to go beyond 2^{47} bytes (x'800 000 000 000').

The log is made up of *log records*, each of which is a set of log data treated as a single unit. A log record is identified either by the RBA of the first byte of its header, or by its log record sequence number (LRSN). The RBA or LRSN uniquely identifies a record that starts at a particular point in the log.

Whether you use the RBA or LRSN to identify log points depends on whether you are using queue-sharing groups. In a queue-sharing environment, the relative byte address cannot be used to uniquely identify a log point, because multiple queue managers can update the same queue at the same time, and each has its own log. To solve this, the log record sequence number is derived from a timestamp value, and does not necessarily represent the physical displacement of the log record within the log.

Each log record has a header that gives its type, the WebSphere MQ subcomponent that made the record, and, for unit of recovery records, a unit of recovery identifier.

There are four types of log record, described under these headings:

- "Unit-of-recovery log records" on page 33
- "Checkpoint records" on page 33
- "Page set control records" on page 33
- "CF structure backup records" on page 33

Unit-of-recovery log records

Most of the log records describe changes to WebSphere MQ queues. All such changes are made within units of recovery.

WebSphere MQ uses special logging techniques involving *undo/redo* and *compensating log records* to reduce restart times and improve system availability.

One effect of this is that the restart time is bounded. If a failure occurs during a restart so that the queue manager has to be restarted a second time, all the recovery activity that completed to the point of failure in the first restart does not need to be reapplied during a second restart. This means that successive restarts do not take progressively longer times to complete.

Checkpoint records

To reduce restart time, WebSphere MQ takes periodic checkpoints during normal operation. These occur:

- When a predefined number of log records has been written. This number is defined by the checkpoint frequency operand called LOGLOAD of the system parameter macro CSQ6SYSP, described in the *WebSphere MQ for z/OS System Setup Guide*.
- At the end of a successful restart.
- At normal termination.
- Whenever WebSphere MQ switches to the next active log data set in the cycle.

At the time a checkpoint is taken, WebSphere MQ issues the DISPLAY THREAD command (described in the *WebSphere MQ Script (MQSC) Command Reference*) internally so that a list of threads currently in doubt is written to the z/OS console log.

Page set control records

These records register the page sets known to the WebSphere MQ queue manager at each checkpoint, and record information about the log ranges required to perform media recovery of the page set at the time of the checkpoint.

CF structure backup records

These records hold data read from a Coupling Facility list structure in response to a BACKUP CFSTRUCT command. In the unlikely event of a Coupling Facility structure failure, these records are used, together with unit of recovery records, by the RECOVER CFSTRUCT command to perform media recovery of the Coupling Facility structure to the point of failure.

How the log is structured

Each active log data set must be a VSAM linear data set (LDS). The physical output unit written to the active log data set is a 4 KB control interval (CI). Each CI contains one VSAM record.

Physical and logical log records

One VSAM CI is a *physical* record. The information to be logged at a particular time forms a *logical* record, whose length varies independently of the space available in the CI. So one physical record might contain:

- Several logical records

- One or more logical records and part of another logical record
- Part of one logical record only

The term “log record” refers to the *logical* record, regardless of how many *physical* records are needed to store it.

How the logs are written

WebSphere MQ writes each log record to a DASD data set called the *active log*. When the active log is full, WebSphere MQ copies its contents to a DASD or tape data set called the *archive log*. This process is called *off-loading*.

Figure 10 illustrates the process of logging. Log records typically go through the following cycle:

1. WebSphere MQ notes changes to data and significant events in recovery log records.
2. WebSphere MQ processes recovery log records and breaks them into segments, if necessary.
3. Log records are placed sequentially in *output log buffers*, which are formatted as VSAM CIs. Each log record is identified by a relative byte address in the range zero through 2^{48}–1.
4. The CIs are written to a set of predefined DASD active log data sets, which are used sequentially and recycled.
5. If archiving is active, as each active log data set becomes full, its contents are automatically off-loaded to a new archive log data set.

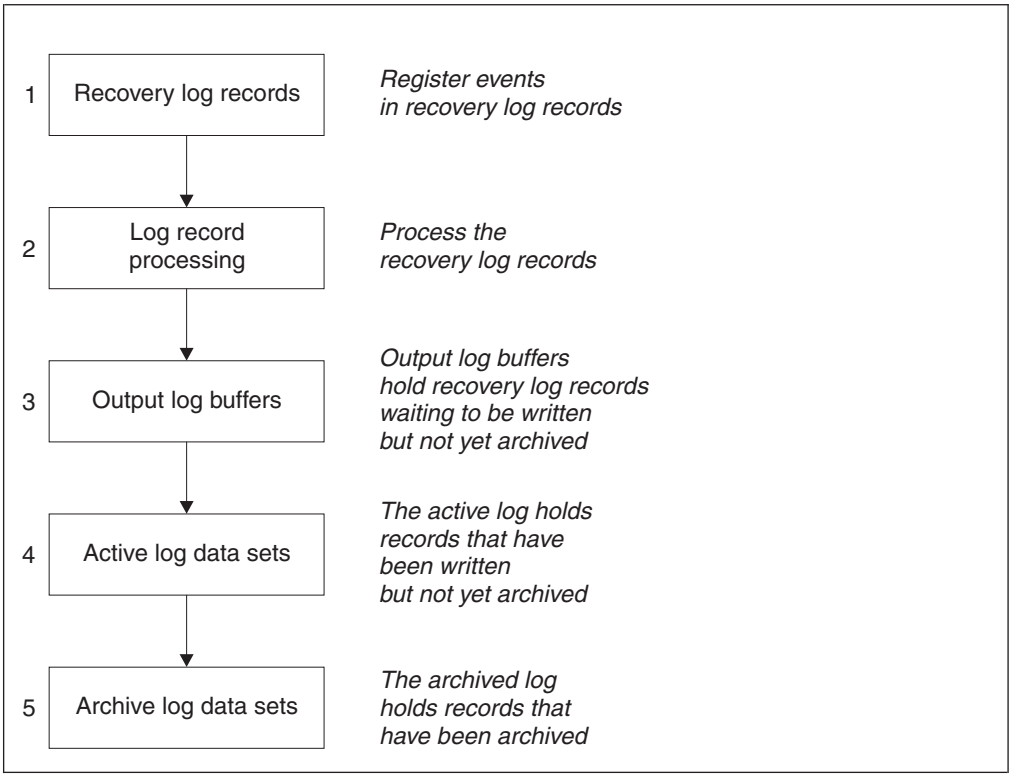

Figure 10. The logging process

When the active log is written

The in-storage log buffers are written to an active log data set whenever any of the following occur:

- The log buffers become full.
- The write threshold is reached (as specified in the CSQ6LOGP macro).
- Certain significant events occur, such as a commit point, or when a WebSphere MQ BACKUP CFSTRUCT command is issued.

When the queue manager is initialized, the active log data sets named in the BSDS are dynamically allocated for exclusive use by the queue manager and remain allocated exclusively to WebSphere MQ until the queue manager terminates. To add or replace active log data sets, you must terminate and restart the queue manager.

When the archive log is written

The process of copying active logs to archive logs is called *off-loading*. The relation of off-loading to other logging events is shown schematically in Figure 11.

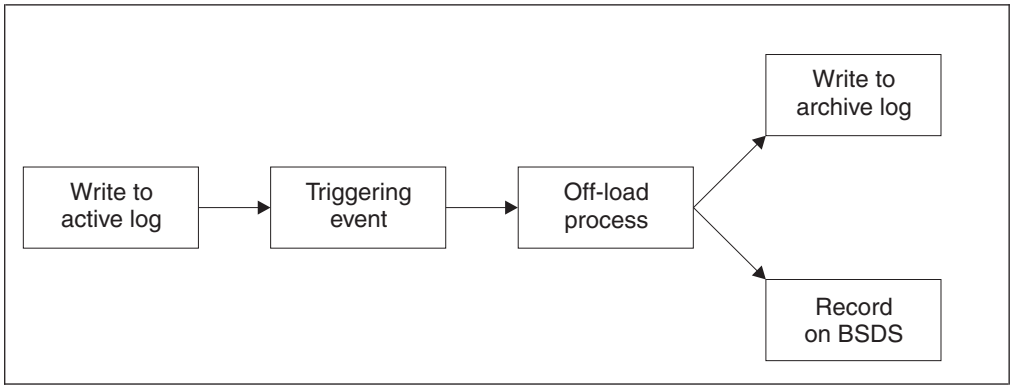

Figure 11. The off-loading process

Triggering an off-load

The off-load of an active log to an archive log can be triggered by several events. For example:

- Filling an active log data set.
- Using the MQSC ARCHIVE LOG command.
- An error occurring while writing to an active log data set.

The data set is truncated before the point of failure, and the record that failed to be written becomes the first record of the new data set. Off-load is triggered for the truncated data set as it would be for a normal full log data set. If there are dual active logs, both copies are truncated so that the two copies remain synchronized.

Message CSQJ110E is issued when the last available active log is 5% full and at 5% increments thereafter, stating the percentage of the log's capacity in use. If all the active logs become full, WebSphere MQ stops processing, until off-loading occurs, and issues this message:

```
CSQJ111A +CSQ1 OUT OF SPACE IN ACTIVE LOG DATA SETS
```

The off-load process

When all the active logs become full, WebSphere MQ runs an off-load and halts processing until the off-load has been completed. If the off-load processing fails when the active logs are full, WebSphere MQ abends.

When an active log is ready to be off-loaded, a request is sent to the z/OS console operator to mount a tape or prepare a DASD unit. The value of the ARCWTOR logging option (discussed in the *WebSphere MQ for z/OS System Setup Guide*) determines whether the request is received. If you are using tape for off-loading, specify ARCWTOR=YES. If the value is YES, the request is preceded by a WTOR (message number CSQJ008E) telling the operator to prepare an archive log data set to be allocated.

The operator need not respond to this message immediately. However, delaying the response delays the off-load process. It does not affect WebSphere MQ performance unless the operator delays the response for so long that WebSphere MQ runs out of active logs.

The operator can respond by canceling the off-load. In this case, if the allocation is for the first copy of dual archive data sets, the off-load is merely delayed until the next active log data set becomes full. If the allocation is for the second copy, the archive process switches to single copy mode, but for this data set only.

Interruptions and errors while off-loading

A request to stop the queue manager does not take effect until off-loading has finished. If WebSphere MQ fails while off-loading is in progress, off-load begins again when the queue manager is restarted. Off-load handling of I/O errors on the logs is discussed in the *WebSphere MQ for z/OS System Administration Guide*.

Messages during off-load

Off-load messages are sent to the z/OS console by WebSphere MQ and the off-load process. These messages can be used to find the RBA ranges in the various log data sets. For an explanation of the off-load messages, see the *WebSphere MQ for z/OS Messages and Codes* manual.

WebSphere MQ and SMS

WebSphere MQ parameters enable you to specify Storage Management Subsystem (MVS/DFP™ SMS) storage classes when allocating WebSphere MQ archive log data sets dynamically. WebSphere MQ initiates the archiving of log data sets, but SMS can be used to perform allocation of the archive data set.

What the bootstrap data set is for

The *bootstrap data set* (BSDS) is a VSAM key-sequenced data set (KSDS) that holds information needed by WebSphere MQ. It contains:

- An inventory of all active and archived log data sets known to WebSphere MQ. WebSphere MQ uses this inventory to:
  - Track the active and archived log data sets
  - Locate log records so that it can satisfy log read requests during normal processing
  - Locate log records so that it can handle restart processing

  WebSphere MQ stores information in the inventory each time an archive log data set is defined or an active log data set is reused. For active logs, the inventory shows which are full and which are available for reuse. The inventory holds the relative byte address (RBA) of each portion of the log held in that data set.
- A *wrap-around* inventory of all recent WebSphere MQ activity. This is needed if the queue manager has to be restarted.

The BSDS is required if the queue manager has an error and has to be restarted. WebSphere MQ must have a BSDS; it is not optional. To minimize the likelihood of problems during a restart, WebSphere MQ can be configured with dual BSDSs, each recording the same information. This is known as running in *dual mode*. If possible, the copies should be on separate volumes. This reduces the risk of them both being lost if the volume is corrupted or destroyed. You should use dual BSDSs rather than dual write to DASD.

The BSDS is set up when WebSphere MQ is customized and the inventory can be managed using the change log inventory utility (CSQJU003). This utility is discussed in the *WebSphere MQ for z/OS System Administration Guide*. It is referenced by a DD statement in the queue manager startup procedure.

Normally, WebSphere MQ keeps duplicate copies of the BSDS. If an I/O error occurs, it deallocates the failing copy and continues with a single BSDS. You can restore dual mode operation, this is described in the *WebSphere MQ for z/OS System Administration Guide*.

The active logs are first registered in the BSDS when WebSphere MQ is installed. They cannot be replaced, nor can new ones be added, without terminating and restarting the queue manager.

Archive log data sets are allocated dynamically. When one is allocated, the data set name is registered in the BSDS. The list of archive log data sets expands as archives are added, and wraps when a user-determined number of entries has been reached. The maximum number of entries is 1000 for single archive logging and 2000 for dual logging.

You can use a tape management system to delete the archive log data sets (WebSphere MQ does not have an automated method). Therefore, the information about an archive log data set can be in the BSDS long after the archive log data set has been deleted by the system administrator.

Conversely, the maximum number of archive log data sets could have been exceeded, and the data from the BSDS dropped long before the data set has reached its expiry date.

The following MQSC command can be used to determine the extent of the log, and the name of the active or archive log data set holding the earliest log RBA, required for various types of media or queue manager recovery:

```
DISPLAY USAGE TYPE(DATASET)
```

If the system parameter module specifies that archive log data sets are cataloged when allocated, the BSDS points to the integrated catalog facility (ICF) catalog for the information needed for later allocations. Otherwise, the BSDS entries for each volume register the volume serial number and unit information that is needed for later allocations.

Archive log data sets and BSDS copies

Each time a new archive log data set is created, a copy of the BSDS is also created. If the archive log is on tape, the BSDS is the first data set on the first output volume. If the archive log is on DASD, the BSDS is a separate data set.

The data set names of the archive log and the BSDS copy are the same, except that the lowest-level qualifier of the archive log name begins with A and the BSDS copy begins with B, for example:

Archive log name
CSQ.ARCHLOG1.E00186.T2336229.*A*0000001

BSDS copy name
CSQ.ARCHLOG1.E00186.T2336229.*B*0000001

If there is a read error while copying the BSDS, the copy is not created, message CSQJ125E is issued, and the off-load to the new archive log data set continues without the BSDS copy.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 4. Where to find more information about logging

Topic	Where to look
How many logs are required How large to make the logs	Chapter 18, "Planning your logging environment", on page 137
Setting up your logs System parameter macros	*WebSphere MQ for z/OS System Setup Guide*
Day to day logging tasks Resolving problems with logs Utility programs	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*

Logging

Chapter 5. Defining your system

This chapter discusses the following topics:

Setting system parameters

In WebSphere MQ for z/OS, a system parameter module controls the logging, archiving, tracing, and connection environments that WebSphere MQ uses in its operation. The system parameters are specified by three assembler macros, as follows:

CSQ6SYSP
: System parameters, including setting the connection and tracing environment.

CSQ6LOGP
: Logging parameters.

CSQ6ARVP
: Log archive parameters.

There is also a channel initiator parameter module that controls how the channel initiator operates; its macro is called CSQ6CHIP.

Default parameter modules are supplied with WebSphere MQ for z/OS. If these do not contain the values that you want to use, you can create your own parameter modules using the samples supplied with WebSphere MQ. The samples are thlqual.SCSQPROC(CSQ4ZPRM) and thlqual.SCSQPROC(CSQ4XPRM).

You can alter some system parameters while a queue manager is still running. See the SET SYSTEM, SET LOG and SET ARCHIVE commands in the *WebSphere MQ Script (MQSC) Command Reference*. Put the SET commands in your initialization input data sets so that they take effect every time you start the queue manager.

Defining system objects

There are several objects that you need to define before you can use WebSphere MQ for z/OS. These are called the system objects and are described here. Sample definitions are supplied with WebSphere MQ to help you define these objects. These samples are described in "Sample definitions supplied with WebSphere MQ" on page 46.

System default objects

The system default objects are used to provide default attributes when you define an object and do not specify the name of another object to base the definition on.

The names of the default system object definitions begin with the characters "SYSTEM.DEFAULT" or "SYSTEM.DEF". For example, the system default local queue is named SYSTEM.DEFAULT.LOCAL.QUEUE.

These objects define the system defaults for the attributes of these WebSphere MQ objects:

- Local queues
- Model queues
- Alias queues
- Remote queues
- Processes
- Namelists
- Channels
- Storage classes
- Authentication information

Shared queues are a special type of local queue, so when you define a shared queue, the definition is based on the SYSTEM.DEFAULT.LOCAL.QUEUE. You need to remember to supply a value for the Coupling Facility structure name because one is not specified in the default definition. Alternatively, you could define your own default shared queue definition to use as a basis for shared queues so that they all inherit the required attributes. Remember that you need to define a shared queue on one queue manager in the queue-sharing group only.

System command objects

The names of the system command objects begin with the characters SYSTEM.COMMAND. You must define these objects before the WebSphere MQ operations and control panels can be used to issue commands to a WebSphere MQ subsystem.

There are two system-command objects:

1. The system-command input queue is a local queue on which commands are put before they are processed by the WebSphere MQ command processor. It must be called SYSTEM.COMMAND.INPUT.
2. SYSTEM.COMMAND.REPLY.MODEL is a model queue that defines the system-command reply-to queue.

Commands are normally sent using nonpersistent messages so both the system-command objects should have the DEFPSIST(NO) attribute so that applications using them (including the supplied applications like the utility program and the operations and control panels) get nonpersistent messages by default. If you have an application that uses persistent messages for commands, set the DEFTYPE(PERMDYN) attribute for the reply-to queue, because the reply messages to such commands are persistent.

System administration objects

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

These queues are used for event messages. The names of the system administration objects begin with the characters SYSTEM.ADMIN.

There are four system-administration objects:

- The SYSTEM.ADMIN.QMGR.EVENT queue

- The SYSTEM.ADMIN.PERFM.EVENT queue
- The SYSTEM.ADMIN.CHANNEL.EVENT queue
- The SYSTEM.ADMIN.CONFIG.EVENT queue

Channel queues

To use distributed queuing, you need to define the following objects:

- A local queue with the name SYSTEM.CHANNEL.SYNCQ, which is used to maintain sequence numbers and logical units of work identifiers (LUWID) of channels. To improve channel performance, you should define this queue with an index type of MSGID (as shown in the supplied sample queue definition).
- A local queue with the name SYSTEM.CHANNEL.INITQ, which is used for channel commands.

You cannot define these queues as shared queues.

Cluster queues

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

To use WebSphere MQ clusters, you need to define the following objects:

- A local queue called the SYSTEM.CLUSTER.COMMAND.QUEUE, which is used to communicate repository changes between queue managers. Messages written to this queue contain updates to the repository data to be applied to the local copy of the repository, or requests for repository data.
- A local queue called SYSTEM.CLUSTER.REPOSITORY.QUEUE, which is used to hold a persistent copy of the repository.
- A local queue called SYSTEM.CLUSTER.TRANSMIT.QUEUE, which is the transmission queue for all destinations in the cluster. For performance reasons you should define this queue with an index type of CORRELID (as shown in the sample queue definition).

These queues typically contain large numbers of messages.

You cannot define these queues as shared queues.

Queue-sharing group queues

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

To use shared channels and intra-group queuing, you need to define the following objects:

- A shared queue with the name SYSTEM.QSG.CHANNEL.SYNCQ, which is used to hold synchronization information for shared channels.
- A shared queue with the name SYSTEM.QSG.TRANSMIT.QUEUE, which is used as the transmission queue for intra-group queuing. If you are running in a queue-sharing group, you must define this queue, even if you are not using intra-group queuing.

Storage classes

You are recommended to define the following six storage classes. You must define four of them because they are required by WebSphere MQ. The other storage class definitions are recommended because they are used in the sample queue definitions.

DEFAULT (required)
This storage class is used for all message queues that are not performance critical and that don't fit in to any of the other storage classes. It is also the supplied default storage class if you do not specify one when defining a queue.

NODEFINE (required)
This storage class is used if the storage class specified when you define a queue is not defined.

REMOTE (required)
This storage class is used primarily for transmission queues, that is, system related queues with short-lived performance-critical messages.

SYSLNGLV
This storage class is used for long-lived, performance-critical messages.

SYSTEM (required)
This storage class is used for performance critical, system related message queues, for example the SYSTEM.CHANNEL.SYNQ and the SYSTEM.CLUSTER.* queues.

SYSVOLAT
This storage class is used for short-lived, performance-critical messages.

You can modify their attributes and add other storage class definitions as required.

Dead-letter queue

The dead-letter queue is used if the message destination is not valid. WebSphere MQ puts such messages on a local queue called the dead-letter queue. Although having a dead-letter queue is not mandatory, it should be regarded as essential, especially if you are using either distributed queuing or one of the WebSphere MQ bridges.

Do **not** define the dead-letter queue as a shared queue.

If you decide to define a dead-letter queue, you must also tell the queue manager its name. To do this use the ALTER QMGR command, as shown in the sample.

Default transmission queue

The default transmission queue is used when no other suitable transmission queue is available for sending messages to another queue manager. If you define a default transmission queue, you must also define a channel to serve the queue. If you do not do this, messages that are put on to the default transmission queue are not transmitted to the remote queue manager and remain on the queue.

If you decide to define a default transmission queue, you must also tell the queue manager its name. To do this use the ALTER QMGR command.

Tuning your queue manager

There are a number of ways in which you can improve the performance of your queue manager, which are controlled by queue manager attributes set by the ALTER QMGR command. This section discusses how you can do this by setting the maximum number of messages allowed on the queue manager, or by performing 'housekeeping' on the queue manager.

Syncpoints

One of the roles of the queue manager is syncpoint control within an application. An application constructs a unit of work containing any number of MQPUT or MQGET calls terminated with an MQCMIT call. However, as the number of MQPUT or MQGET calls within the scope of one MQCMIT increases, the performance cost of the commit increases significantly.

You can limit the number of messages within any single syncpoint by using the MAXUMSGS queue manager attribute. You should not normally exceed 100 MQPUTs within a single MQCMIT.

You are recommended to set a fairly low MAXUMSGS value, such as 100, to avoid performance problems, and to protect against looping applications. If you have a need for a larger value, it would be better to change MAXUMSGS temporarily while the application that requires the larger value is run, rather than using a permanently high value.

Expired messages

Messages that have expired are discarded by the the next appropriate MQGET call. However, if no such call occurs, the expired message is not discarded, and, for some queues, a large number of expired messages can accumulate. To remedy this, you can set the queue manager to scan queues periodically and discard expired messages on one or more queues.

There are two ways of doing this:

Periodic scan
: You can specify a period using the EXPRYINT (expiry interval) queue manager attribute. Each time the expiry interval is reached, the queue manager looks for candidate queues that are worth scanning to discard expired messages.

  The queue manager maintains information about the expired messages on each queue, and knows whether a scan for expired messages is worthwhile. So, only a selection of queues is scanned at any time.

  Shared queues are scanned only by one queue manager in a queue-sharing group. Generally, it is the first queue manager to restart, or the first to have EXPRYINT set. Set the expiry interval value for all queue managers within a queue-sharing group to the same value.

Explicit request
: Issue the REFRESH QMGR TYPE(EXPIRY) command, specifying the queue or queues that you want scanned.

Sample definitions supplied with WebSphere MQ

The following sample definitions are supplied with WebSphere MQ in the thlqual.SCSQPROC library. You can use them to define the system objects and to customize your own objects. You can include some of them in the initialization input data sets (described in "Initialization commands" on page 85).

Note: In Table 5, the second column (full function) gives the sample names for the complete WebSphere MQ for z/OS Version 5 Release 3.1 product; the third column (reduced function) gives the sample names for the reduced function form of WebSphere MQ for z/OS Version 5 Release 3.1 available with WebSphere Application Server. *None* indicates that there is no equivalent sample in the reduced function form, because the full function sample features functions that do not apply to the reduced function form.

Table 5. WebSphere MQ sample definitions for system objects

Initialization input data set	Sample name (full function)	Sample name (reduced function)
CSQINP1	CSQ4INP1	CSQ4INPR
CSQINP2	CSQ4INSG CSQ4INSS CSQ4INSX CSQ4INYC CSQ4INYD CSQ4INYG	CSQ4INSR *None* *None* *None* CSQ4INYR CSQ4INYR
Other	CSQ4DISP CSQ4DISQ CSQ4INPX CSQ4IVPQ CSQ4IVPG	CSQ4DISP *None* CSQ4INPX CSQ4IVPQ *None*

The CSQINP1 sample

The sample CSQINP1 data set thlqual.SCSQPROC(CSQ4INP1) (CSQ4INPR when using reduced function WebSphere MQ) contains definitions of buffer pools, page set to buffer pool associations, and an ALTER SECURITY command. The sample should be included in the CSQINP1 concatenation of your queue manager started task procedure.

CSQ4INSG system object sample

The sample CSQINP2 data set thlqual.SCSQPROC(CSQ4INSG) (CSQ4INSR when using reduced function WebSphere MQ) contains definitions for the following system objects for general use:

- System default objects
- System command objects
- System administration objects

You must define the objects in this sample, but you need to do it only once when the subsystem is first started. Including the definitions in the CSQINP2 data set is the best way to do this. They are maintained across queue manager shutdown and restart. You must not change the object names, but you can change their attributes if required.

CSQ4INSS system object sample

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

You can define additional system objects if you are using queue-sharing groups.

Sample data set thlqual.SCSQPROC(CSQ4INSS) contains sample commands for use with CF structures and a set of definitions for the system objects required for shared channels and intra-group queuing.

You cannot use this sample as is; you must customize it before use. Then you can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to the COMMAND function of the CSQUTIL utility to issue the required commands.

When you are defining group or shared objects, you only need to include them in the CSQINP2 DD concatenation for one queue manager in the queue-sharing group.

CSQ4INSX system object sample

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

You must define additional system objects if you are using distributed queuing and clustering.

Sample data set thlqual.SCSQPROC(CSQ4INSX) contains the queue definitions required. You can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to the COMMAND function in CSQUTIL utility to issue the required DEFINE commands.

There are two types of object definitions:
- SYSTEM.CHANNEL.xx, needed for any distributed queuing
- SYSTEM.CLUSTER.xx, needed for clustering

CSQ4INYC object sample

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

If you are using clustering, definitions equivalent to the channel definitions and remote queue definitions of distributed queuing are created automatically, when needed. However, some manual channel definitions are needed – a cluster-receiver channel for the cluster and a cluster-sender definition to at least one cluster repository queue manager.

Sample data set thlqual.SCSQPROC(CSQ4INYC) contains the following sample definitions that you can use for customizing your clustering objects:
- Definitions for the queue manager
- Definitions for the receiving channel
- Definitions for the sending channel
- Definitions for cluster queues
- Definitions for lists of clusters

You cannot use this sample as is; you must customize it before use. Then you can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to the COMMAND function of the CSQUTIL utility to issue the required DEFINE commands. (This is preferable because it means that you don't have to redefine these objects each time that you restart WebSphere MQ).

CSQ4INYD object sample

If you are using distributed queuing and not clustering, you need to set up your own queues, processes, and channels.

Sample data set thlqual.SCSQPROC(CSQ4INYD) (CSQ4INYR when using reduced function WebSphere MQ) contains sample definitions that you can use for customizing your distributed queuing objects. It comprises:

- A set of definitions for the sending end
- A set of definitions for the receiving end
- A set of definitions for using clients

You cannot use this sample as is; you must customize it before use. Then you can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to the COMMAND function of the CSQUTIL utility to issue the required DEFINE commands. (This is preferable because it means that you don't have to redefine these objects each time you restart the queue manager).

CSQ4INYG object sample

Sample data set thlqual.SCSQPROC(CSQ4INYG) (CSQ4INYR when using reduced function WebSphere MQ) contains the following sample definitions that you can use for customizing your own objects for general use:

- Storage classes
- Dead-letter queue
- Default transmission queue
- CICS adapter objects

You cannot use this sample as is, you must customize it before use. Then you can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to the COMMAND function of the CSQUTIL utility to issue the required DEFINE commands. (This is preferable because it means that you don't have to redefine these objects each time that you restart WebSphere MQ).

In addition to the sample definitions here, you can use the system object definitions as the basis for your own resource definitions. For example, you could make a working copy of SYSTEM.DEFAULT.LOCAL.QUEUE and name it MY.DEFAULT.LOCAL.QUEUE. You can then change any of the parameters in this copy as required. You could then issue a DEFINE command by whichever method you choose, provided you have the authority to create resources of that type.

Default transmission queue

Read "Default transmission queue" on page 44 before you decide whether you want to define a default transmission queue.

- If you decide that you do want to define a default transmission queue, remember that you must also define a channel to serve it.
- If you decide that you do not want to define one, remember to remove the DEFXMITQ statement from the ALTER QMGR command in the sample.

CICS adapter objects

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

The sample defines an initiation queue named CICS01.INITQ. This queue is used by the WebSphere MQ-supplied CKTI transaction. You can change the name of this queue; however it must match the name specified in the CICS system initialization table (SIT) or SYSIN override in the INITPARM statement.

CSQ4DISP display sample

Sample data set thlqual.SCSQPROC(CSQ4DISP) contains a set of generic DISPLAY commands that display all the defined resources on your queue manager. This includes the definitions for all WebSphere MQ objects and definitions such as storage classes and trace. These commands can generate a large amount of output. This sample can be used in the CSQINP2 data set or as input to the COMMAND function of the CSQUTIL utility.

CSQ4DISQ distributed queuing using CICS sample

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

Sample data set thlqual.SCSQPROC(CSQ4DISQ) contains a set of commands that are required to implement distributed queuing using CICS ISC. (This is described in the *WebSphere MQ for z/OS System Setup Guide*).

You can include this member in the CSQINP2 DD concatenation of the queue manager startup procedure, or you can use it as input to COMMAND function of the CSQUTIL utility to issue the required DEFINE commands.

CSQ4INPX sample

Sample data set thlqual.SCSQPROC(CSQ4INPX) contains a set of commands that you might want to execute each time the channel initiator starts. You must customize this sample before use; you can then include it in the CSQINPX data set for the channel initiator.

CSQ4IVPQ and CSQ4IVPG samples

Sample data sets thlqual.SCSQPROC(CSQ4IVPQ) and thlqual.SCSQPROC(CSQ4IVPG) contain sets of DEFINE commands that are required to run the installation verification programs (IVPs). (CSQ4IVPG is not available when using reduced function WebSphere MQ.)

You can include these samples in the CSQINP2 data set. When you have run the IVPs successfully, you do not need to run them again each time the queue manager is restarted. Therefore, you do not need to keep these samples in the CSQINP2 concatenation permanently.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 6. Where to find more information about system parameters and system objects

Topic	Where to look
Using initialization input data sets System parameter macros Installation verification program	*WebSphere MQ for z/OS System Setup Guide*
Utility programs	*WebSphere MQ for z/OS System Administration Guide*
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*
WebSphere MQ clusters	*WebSphere MQ Queue Manager Clusters*
WebSphere MQ events	*WebSphere MQ Event Monitoring*

Chapter 6. Recovery and restart

This chapter describes how a queue manager recovers after it has stopped, and what happens when it is restarted. It contains the following sections:

How changes are made to data

WebSphere MQ must interact with other subsystems to keep all the data consistent. This section discusses *units of recovery*; what they are and how they are used in *back outs*.

Units of recovery

A *unit of recovery* is the processing done by a single queue manager for an application program, that changes WebSphere MQ data from one point of consistency to another. A *point of consistency* – also called a *syncpoint* or *commit point* – is a point in time when all the recoverable data that an application program accesses is consistent.

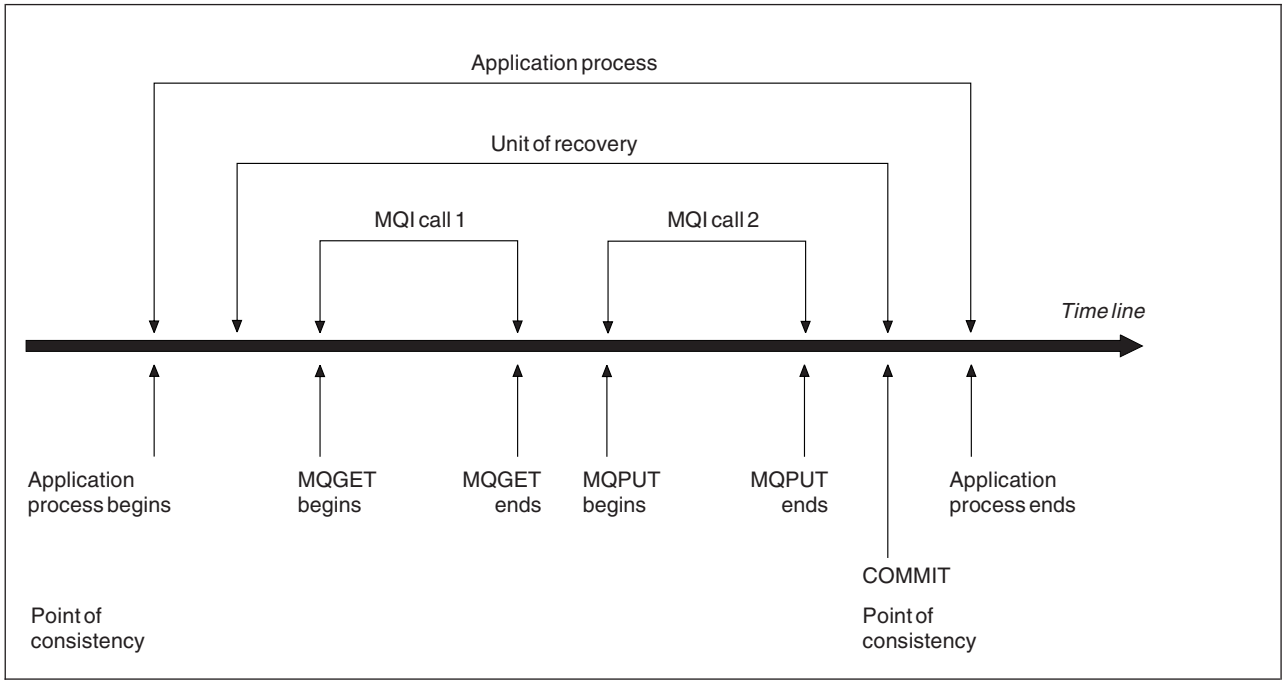

Figure 12. A unit of recovery within an application program. Typically, the unit of recovery consists of more than one MQI call. More than one unit of recovery can occur within an application program.

A unit of recovery begins with the first change to the data after the beginning of the program or following the previous point of consistency; it ends with a later point of consistency. Figure 12 shows the relationship between units of recovery,

the point of consistency, and an application program. In this example, the application program makes changes to queues through MQI calls 1 and 2. The application program can include more than one unit of recovery or just one. However, any complete unit of recovery ends in a commit point.

For example, a bank transaction transfers funds from one account to another. First, the program subtracts the amount from the first account, account A. Then, it adds the amount to the second account, B. After subtracting the amount from A, the two accounts are inconsistent and WebSphere MQ cannot commit. They become consistent when the amount is added to account B. When both steps are complete, the program can announce a point of consistency through a commit, making the changes visible to other application programs.

Normal termination of an application program automatically causes a point of consistency. Some program requests in CICS and IMS programs also cause a point of consistency, for example, EXEC CICS SYNCPOINT.

Backing out work

If an error occurs within a unit of recovery, WebSphere MQ removes any changes to data, returning the data to its state at the start of the unit of recovery; that is, WebSphere MQ backs out the work. The events are shown in Figure 13.

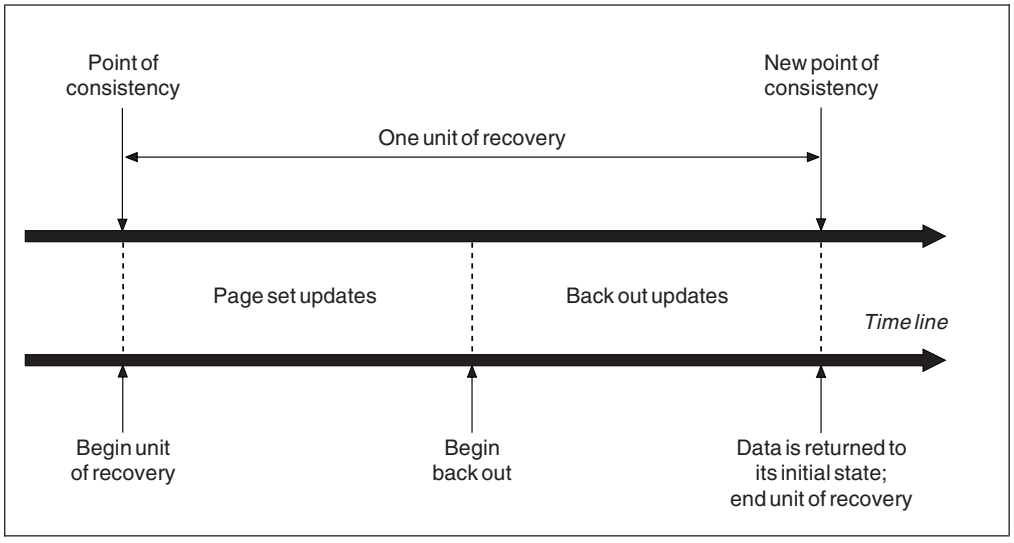

Figure 13. A unit of recovery showing back out

How consistency is maintained

If data in WebSphere MQ is to be consistent with batch, CICS, IMS, or TSO, any data changed in one must be matched by a change in the other. Before one system commits the changed data, it must know that the other system can make the corresponding change. So, the systems must communicate.

During a *two-phase commit* (for example under CICS), one subsystem coordinates the process. That subsystem is called the *coordinator*; the other is the *participant*. CICS or IMS is always the coordinator in interactions with WebSphere MQ, and WebSphere MQ is always the participant. In the batch or TSO environment, WebSphere MQ can participate in two-phase commit protocols coordinated by z/OS RRS.

During a *single-phase commit* (for example under TSO or batch), WebSphere MQ is always the coordinator in the interactions and completely controls the commit process.

In a WebSphere Application Server environment, the semantics of the JMS session object determine whether single-phase or two-phase commit coordination is used.

Consistency with CICS or IMS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

The connection between WebSphere MQ and CICS or IMS supports the following syncpoint protocols:

- Two-phase commit – for transactions that update resources owned by more than one resource manager.

  This is the standard distributed syncpoint protocol. It involves more logging and message flows than a single-phase commit.
- Single-phase commit – for transactions that update resources owned by a single resource manager (WebSphere MQ).

  This protocol is optimized for logging and message flows.
- Bypass of syncpoint – for transactions that involve WebSphere MQ but which do nothing in the queue manager that requires a syncpoint (for example, browsing a queue).

In each case, CICS or IMS acts as the syncpoint manager.

The stages of the two-phase commit that WebSphere MQ uses to communicate with CICS or IMS are:

1. In phase 1, each system determines independently whether it has recorded enough recovery information in its log, and can commit its work.

At the end of the phase, the systems communicate. If they agree, each begins the next phase.

2. In phase 2, the changes are made permanent. If one of the systems abends during phase 2, the operation is completed by the recovery process during restart.

Illustration of the two-phase commit process

Figure 14 illustrates the two-phase commit process. Events in the CICS or IMS coordinator are shown on the upper line, events in WebSphere MQ on the lower line.

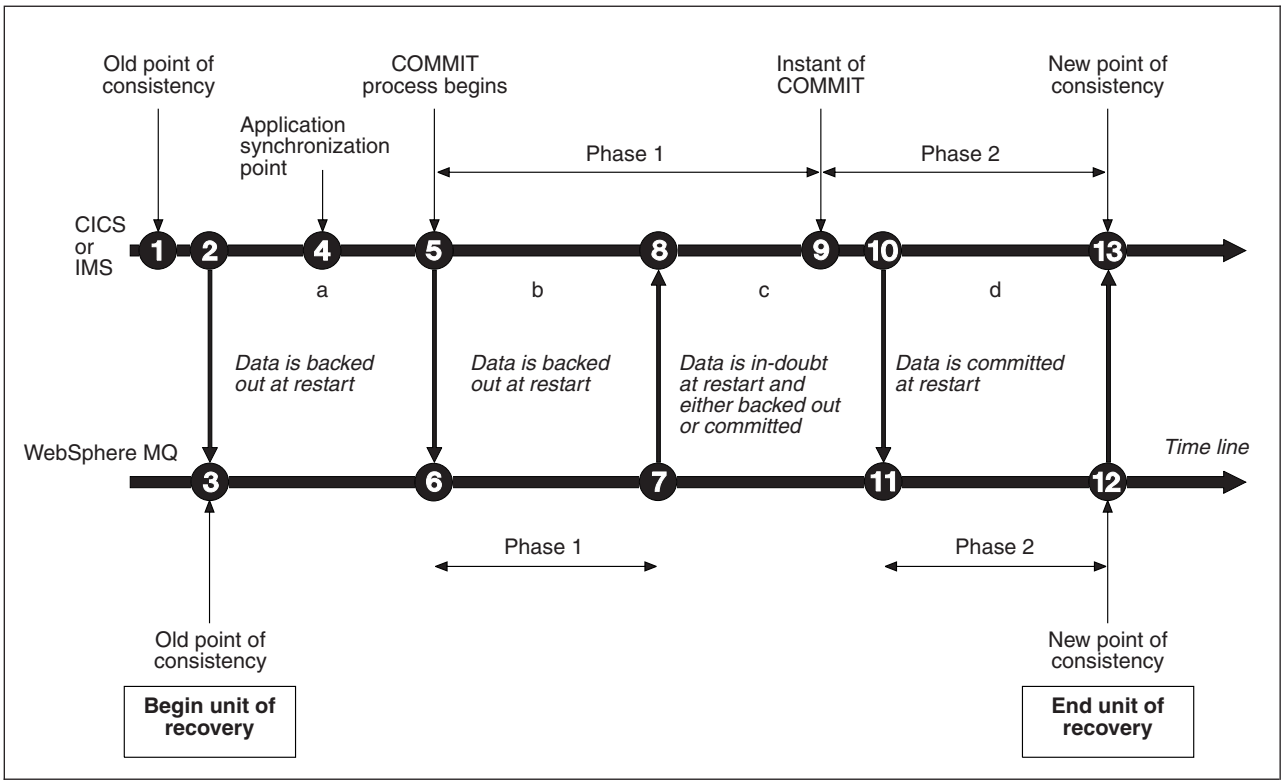

Figure 14. The two-phase commit process

The numbers in the following discussion are linked to those in the figure.

1. The data in the coordinator is at a point of consistency.
2. An application program in the coordinator calls WebSphere MQ to update a queue by adding a message.
3. This starts a unit of recovery in WebSphere MQ.
4. Processing continues in the coordinator until an application synchronization point is reached.
5. The coordinator then starts commit processing. CICS programs use a SYNCPOINT command or a normal application termination to start the commit. IMS programs can start the commit by using a CHKP call, a SYNC call, a GET UNIQUE call to the IOPCB, or a normal application termination. Phase 1 of commit processing begins.
6. As the coordinator begins phase 1 processing, so does WebSphere MQ.
7. WebSphere MQ successfully completes phase 1, writes this fact in its log, and notifies the coordinator.
8. The coordinator receives the notification.
9. The coordinator successfully completes its phase 1 processing. Now both subsystems agree to commit the data changes, because both have completed phase 1 and could recover from any errors. The coordinator records in its log the instant of commit – the irrevocable decision of the two subsystems to make the changes.

The coordinator now begins phase 2 of the processing – the actual commitment.

10. The coordinator notifies WebSphere MQ to begin its phase 2.
11. WebSphere MQ logs the start of phase 2.
12. Phase 2 is successfully completed, and this is now a new point of consistency for WebSphere MQ. WebSphere MQ then notifies the coordinator that it has finished its phase 2 processing.
13. The coordinator finishes its phase 2 processing. The data controlled by both subsystems is now consistent and available to other applications.

How consistency is maintained after an abnormal termination

When a queue manager is restarted after an abnormal termination, it must determine whether to commit or to back out units of recovery that were active at the time of termination. For certain units of recovery, WebSphere MQ has enough information to make the decision. For others, it does not, and must get information from the coordinator when the connection is reestablished.

Figure 14 shows four periods within the two phases: a, b, c, and d. The status of a unit of recovery depends on the period in which termination happened. The status can be:

In flight
The queue manager ended before finishing phase 1 (period a or b); during restart, WebSphere MQ backs out the updates.

In doubt
The queue manager ended after finishing phase 1 and before starting phase 2 (period c); only the coordinator knows whether the error happened before or after the commit (point 9). If it happened before, WebSphere MQ must back out its changes; if it happened after, WebSphere MQ must make its changes and commit them. At restart, WebSphere MQ waits for information from the coordinator before processing this unit of recovery.

In commit
The queue manager ended after it began its own phase 2 processing (period d); it makes committed changes.

In backout
The queue manager ended after a unit of recovery began to be backed out but before the process was complete (not shown in the figure); during restart, WebSphere MQ continues to back out the changes.

What happens during termination

A queue manager terminates normally in response to the STOP QMGR command. If a queue manager stops for any other reason, the termination is considered to be abnormal.

Normal termination

In a normal termination, WebSphere MQ stops all activity in an orderly way. You can stop WebSphere MQ using either quiesce, force, or restart mode. The effects are given in Table 7 on page 56.

Table 7. Termination using QUIESCE, FORCE, and RESTART

Thread type	QUIESCE	FORCE	RESTART
Active threads	Run to completion	Back out	Back out
New threads	Can start	Not permitted	Not permitted
New connections	Not permitted	Not permitted	Not permitted

Batch applications are notified if a termination occurs while the application is still connected.

With CICS, a current thread runs only to the end of the unit of recovery. With CICS, stopping a queue manager in quiesce mode stops the CICS adapter, and so if an active task contains more than one unit of recovery, the task does not necessarily run to completion.

If you stop a queue manager in force or restart mode, no new threads are allocated, and work on connected threads is rolled back. Using these modes can create in-doubt units of recovery for threads that are between commit processing phases. They are resolved when WebSphere MQ is reconnected with the controlling CICS, IMS, or RRS subsystem.

When you stop a queue manager, in any mode, the steps are:

1. Connections are ended.
2. WebSphere MQ ceases to accept commands.
3. WebSphere MQ ensures that any outstanding updates to the page sets are completed.
4. The DISPLAY USAGE command is issued internally by WebSphere MQ so that the restart RBA is recorded on the z/OS console log.
5. The shutdown checkpoint is taken and the BSDS is updated.

Terminations that specify quiesce mode do not affect in-doubt units of recovery. Any unit that is in doubt remains in doubt.

Abnormal termination

An abnormal termination can leave data in an inconsistent state, for example:

- A unit of recovery has been interrupted before reaching a point of consistency.
- Committed data has not been written to page sets.
- Uncommitted data has been written to page sets.
- An application program has been interrupted between phase 1 and phase 2 of the commit process, leaving the unit of recovery in doubt.

WebSphere MQ resolves any data inconsistencies arising from abnormal termination during restart and recovery.

What happens during restart and recovery

WebSphere MQ uses its recovery log and the bootstrap data set (BSDS) to determine what to recover when it restarts. The BSDS identifies the active and archive log data sets, and the location of the most recent WebSphere MQ checkpoint in the log.

After WebSphere MQ has been initialized, the queue manager restart process takes place as follows:

- Log initialization
- Current status rebuild
- Forward log recovery
- Backward log recovery
- Queue index rebuilding

When recovery has been completed:

- Committed changes are reflected in the data.
- In-doubt activity is reflected in the data. However, the data is locked and cannot be used until WebSphere MQ recognizes and acts on the in-doubt decision.
- Interrupted in-flight and in-abort changes have been removed from the queues. The messages are consistent and can be used.
- A new checkpoint has been taken.
- New indexes have been built for indexed queues containing persistent messages (described in "Rebuilding queue indexes").

Batch applications are not notified when restart occurs *after* the application has requested a connection.

If dual BSDSs are in use, WebSphere MQ checks the consistency of the time stamps in the BSDS:

- If both copies of the BSDS are current, WebSphere MQ tests whether the two time stamps are equal. If they are not, WebSphere MQ issues message CSQJ120E and terminates. This can happen when the two copies of the BSDS are maintained on separate DASD volumes and one of the volumes was restored while the queue manager was stopped. WebSphere MQ detects the situation at restart.
- If one copy of the BSDS was deallocated, and logging continued with a single BSDS, a problem could arise. If both copies of the BSDS are maintained on a single volume, and the volume was restored, or if both BSDS copies were restored separately, WebSphere MQ might not detect the restoration. In that case, log records not noted in the BSDS would be unknown to the system.

Rebuilding queue indexes

To increase the speed of **MQGET** operations on a queue where messages are not retrieved sequentially, you can specify that you want WebSphere MQ to maintain an index of the message or correlation identifiers or groupid for all the messages on that queue (as described in the *WebSphere MQ Application Programming Guide*).

When a queue manager is restarted, these indexes are rebuilt for each queue. This only applies to persistent messages; nonpersistent messages are deleted at restart. If your indexed queues contain large numbers of persistent messages, this increases the time taken to restart the queue manager.

Recovery and restart

You can choose to have indexes rebuilt asynchronously to queue manager startup by using the QINDXBLD parameter of the CSQ6SYSP macro. If you set QINDXBLD=NOWAIT, WebSphere MQ restarts without waiting for indexes to rebuild.

How in-doubt units of recovery are resolved

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

If WebSphere MQ loses its connection to CICS, IMS, or RRS, it normally attempts to recover all inconsistent objects at restart. The information needed to resolve in-doubt units of recovery must come from the coordinating system. The next section describes the process of resolution.

How in-doubt units of recovery are resolved from CICS

The resolution of in-doubt units has no effect on CICS resources. CICS is in control of recovery coordination and, when it restarts, automatically commits or backs out each unit, depending on whether there was a log record marking the beginning of the commit. The existence of in-doubt objects does not lock CICS resources while WebSphere MQ is being reconnected.

One of the functions of the CICS adapter is to keep data synchronized between CICS and WebSphere MQ. If a queue manager abends while connected to CICS, it is possible for CICS to commit or back out work without WebSphere MQ being aware of it. When the queue manager restarts, that work is termed *in doubt*.

WebSphere MQ cannot resolve these in-doubt units of recovery (that is, commit or back out the changes made to WebSphere MQ resources) until the connection to CICS is restarted or reconnected.

A process to resolve in-doubt units of recovery is initiated during startup of the CICS adapter. The process starts when the adapter requests a list of in-doubt units of recovery. Then:

- The adapter receives a list of in-doubt units of recovery for this connection ID from WebSphere MQ, and passes them to CICS for resolution.
- CICS compares entries from this list with entries in its own log. CICS determines from its own list what action it took for each in-doubt unit of recovery.

Under some circumstances, CICS cannot run the WebSphere MQ process to resolve in-doubt units of recovery. When this happens, WebSphere MQ sends one of these messages:

- CSQC404E
- CSQC405E
- CSQC406E
- CSQC407E

followed by the message CSQC408I.

For details of what these messages mean, see the *WebSphere MQ for z/OS Messages and Codes* manual.

For all resolved units, WebSphere MQ updates the queues as necessary and releases the corresponding locks. Unresolved units can remain after restart. Resolve them by the methods described in the *WebSphere MQ for z/OS System Administration Guide*.

How in-doubt units of recovery are resolved from IMS

Resolving in-doubt units of recovery in IMS has no effect on DL/I resources. IMS is in control of recovery coordination and, when it restarts, automatically commits or backs out incomplete DL/I work. The decision to commit or back out for online regions (non-fast-path) is on the presence or absence of IMS log record types X'3730' and X'3801' respectively. The existence of in-doubt units of recovery does not imply that DL/I records are locked until WebSphere MQ connects.

During queue manager restart, WebSphere MQ makes a list of in-doubt units of recovery. IMS builds its own list of residual recovery entries (RREs). The RREs are logged at IMS checkpoints until all entries are resolved.

During reconnection of an IMS region to WebSphere MQ, IMS indicates to WebSphere MQ whether to commit or back out units of work marked in WebSphere MQ as in doubt.

When in-doubt units are resolved:

1. If WebSphere MQ recognizes that it has marked an entry for commit and IMS has marked it to be backed out, WebSphere MQ issues message CSQQ010E. WebSphere MQ issues this message for all inconsistencies of this type between WebSphere MQ and IMS.
2. If WebSphere MQ has any remaining in-doubt units, the adapter issues message CSQQ008I.

For all resolved units, WebSphere MQ updates queues as necessary and releases the corresponding locks.

WebSphere MQ maintains locks on in-doubt work that was not resolved. This can cause a backlog in the system if important locks are being held. The connection remains active so you can resolve the IMS RREs. Recover the in-doubt threads by the methods described in the *WebSphere MQ for z/OS System Administration Guide*.

All in-doubt work should be resolved unless there are software or operating problems, such as with an IMS cold start. In-doubt resolution by the IMS control region takes place in two circumstances:

1. At the start of the connection to WebSphere MQ, during which resolution is done synchronously.
2. When a program abends, during which the resolution is done asynchronously.

How in-doubt units of recovery are resolved from RRS

One of the functions of the RRS adapter is to keep data synchronized between WebSphere MQ and other RRS-participating resource managers. If a failure occurs when WebSphere MQ has completed phase one of the commit and is waiting for a decision from RRS (the commit coordinator), the unit of recovery enters the in-doubt state.

When communication is reestablished between RRS and WebSphere MQ, RRS automatically commits or backs out each unit of recovery, depending on whether there was a log record marking the beginning of the commit. WebSphere MQ

cannot resolve these in-doubt units of recovery (that is, commit or back out the changes made to WebSphere MQ resources) until the connection to RRS is reestablished.

Under some circumstances, RRS cannot resolve in-doubt units of recovery. When this happens, WebSphere MQ sends one of the following messages to the z/OS console:

- CSQ3011I
- CSQ3013I
- CSQ3014I
- CSQ3016I

For details of what these messages mean, see the *WebSphere MQ for z/OS Messages and Codes* manual.

For all resolved units of recovery, WebSphere MQ updates the queues as necessary and releases the corresponding locks. Unresolved units of recovery can remain after restart. Resolve them by the method described in the *WebSphere MQ for z/OS System Administration Guide*.

Shared queue recovery

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

This section describes WebSphere MQ recovery in the queue-sharing group environment.

Transactional recovery

When an application issues an **MQBACK** call or terminates abnormally (for example, because of an EXEC CICS ROLLBACK or an IMS abend) thread level information stored in the queue manager ensures that the in-flight unit of work is rolled back. **MQPUT** and **MQGET** operations within syncpoint on shared queues are rolled back in the same way as updates to non-shared queues.

Peer recovery

If a queue manager fails, it disconnects abnormally from the Coupling Facility structures that it is currently connected to. If the connection between the z/OS instance and the Coupling Facility fails (for example, physical link failure or power off of a Coupling Facility or partition) this is also detected as an abnormal termination of the connection between the queue manager and the Coupling Facility structures involved. Other queue managers in the same queue-sharing group that remain connected to that structure detect the abnormal disconnection and all attempt to initiate *peer recovery* for the failed queue manager on that structure. Only one of these queue managers initiates peer recovery successfully, but all the other queue managers cooperate in the recovery of units of work that were owned by the queue manager that failed.

If a queue manager fails when there are no peers connected to a structure, recovery is performed when another queue manager connects to that structure, or when the queue manager that failed restarts.

Peer recovery is performed on a structure by structure basis and it is possible for a single queue manager to be participating in the recovery of more than one

structure at the same time. However, the set of peers cooperating in the recovery of different structures might vary depending on which queue managers were connected to the different structures at the time of failure.

When the failed queue manager restarts, it reconnects to the structures that it was connected to at the time of failure, and recovers any remaining unresolved units of work that were not recovered by peer recovery.

Peer recovery is a multi-phase process. During the first phase, units of work that had progressed beyond the in-flight phase are recovered; this might involve committing messages for units of work that are in-commit and locking messages for units of work that are in-doubt. During the second phase, queues that had threads active against them in the failing queue manager are checked, uncommitted messages related to in-flight units of work are rolled back, and information about active handles on shared queues in the failed queue manager are reset. This means that WebSphere MQ resets any indicators that the failing queue manager had a shared queue open for input-exclusive, allowing other active queue managers to open the queue for input.

Shared queue definitions

The queue objects that represent the attributes of a shared queue are held in the shared DB2 repository used by the queue-sharing group. You should ensure that adequate procedures are in place for the backup and recovery of the DB2 tables used to hold WebSphere MQ objects. You can also use the WebSphere MQ CSQUTIL utility to create MQSC commands for replay into a queue manager to redefine WebSphere MQ objects, including shared queue and group definitions stored in DB2.

Logging

Queue sharing groups can now support persistent messages, as the messages on shared queues can be logged in the queue manager logs.

Coupling Facility failure

In the unlikely event of a Coupling Facility failure, any nonpersistent messages stored in the affected CF structures will be lost. Persistent messages can be recovered, using the RECOVER CFSTRUCT command.

To ensure that you can recover a CF structure in a reasonable time, you must take frequent backups, using the BACKUP CFSTRUCT command. You can choose to 'round-robin' the backups across all the queue managers in the queue-sharing group, or dedicate one queue manager to do all the backups.

Each backup is output to the active log data set of the queue manager taking the backup. The shared queue DB2 repository records the name of the CF structure being backed up, the name of the queue manager doing the backup, the RBA range for this backup on that queue manager's log, and the backup time.

You recover a CF structure by issuing a RECOVER CFSTRUCT command to the queue manager that will do the recovery; you can specify a single CF structure to be recovered, or you can recover several CF structures simultaneously. The command uses the backup, located through the DB2 repository information, and forward recovers this to the point of failure. It does this by applying log records from any queue manager in the queue-sharing group that has performed an MQPUT or MQGET between the start of the backup and the time of failure, on

any shared queue that maps to the CF structure. The resulting merging of the logs might require reading a considerable amount of log data, and so you are strongly advised to make frequent (say, hourly) backups.

If a recoverable application structure has failed, any further application activity is prevented until the structure has been recovered. If the administration structure has also failed, all the queue managers in the queue-sharing group must be started before the RECOVER CFSTRUCT command can be issued.

If a CF structure fails, the action taken by connected queue managers depends on:

- the structure type (application or administration)
- the queue manager level (Version 5.2 or Version 5.3)
- The CFLEVEL of the CFSTRUCT object (1, 2, or 3)
- the type of failure reported by the XES component of z/OS to WebSphere MQ.

If an administration structure fails, all queue managers in the queue-sharing group terminate.

If an application structure fails and the queue manager is running at Version 5.2, the queue manager terminates.

If an application structure fails and the queue manager is running at Version 5.3, and the CFLEVEL of the CFSTUCT is less than 3, the queue manager terminates.

If an application structure fails and the queue manager is running at Version 5.3, and the CFLEVEL is 3, and the error reported by XES is a connection failure, the queue manager terminates to allow other queue managers in the queue-sharing group that might still have connectivity to perform peer level recovery, and so make messages more available.

If a CF structure fails, Version 5.3 queue managers connected to a CFLEVEL(3) CFSTRUCT continue to run, and applications that do not use the queues in the failed structure can continue normal processing. However, applications that attempt operations on queues in the failed structure will receive errors, until the RECOVER CFSTRUCT command has successfully rebuilt the failed structure, at which point new requests to open queues in the structure are allowed.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 8. Where to find more information about recovery and restart

Topic	Where to look
Planning your backup strategy	Chapter 19, “Planning for backup and recovery”, on page 143
System parameters	*WebSphere MQ for z/OS System Setup Guide*
Routine backup and recovery procedures Resolving in-doubt threads Resolving problems during restart Utility programs	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*

Table 8. Where to find more information about recovery and restart (continued)

Topic	Where to look
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*

Recovery and restart

Chapter 7. Security

This chapter discusses WebSphere MQ security. It contains the following sections:

Why you need to protect WebSphere MQ resources

Because WebSphere MQ handles the transfer of information that is potentially valuable, it needs the safeguard of a security system. This is to ensure that the resources WebSphere MQ owns and manages are protected from unauthorized access that might lead to the loss or disclosure of the information. It is essential that none of the following are accessed or changed by any unauthorized user or process:

- Connections to WebSphere MQ
- WebSphere MQ objects such as queues, processes, and namelists
- WebSphere MQ transmission links
- WebSphere MQ system control commands
- WebSphere MQ messages
- Context information associated with messages

To provide the necessary security, WebSphere MQ uses the z/OS system authorization facility (SAF) to route authorization requests to an External Security Manager (ESM), for example Resource Access Control Facility (RACF). WebSphere MQ does no security verification of its own. Where distributed queuing or clients are being used, additional security measures might be required, for which WebSphere MQ provides channel exits and the MCAUSER channel attribute.

The decision to allow access to an object is made by the ESM and WebSphere MQ follows that decision. If the ESM cannot make a decision, WebSphere MQ prevents access to the object.

If you do nothing

If you do nothing about security, the most likely effect is that ***all*** users can access and change ***every*** resource. This includes not only local users, but also those on remote systems using distributed queuing or clients, where the logon security controls might be less strict than is normally the case for z/OS.

To enable security checking you must do the following:

- Install and activate an ESM (for example, RACF)
- Define the MQADMIN class if you are using an ESM other than RACF
- Activate the MQADMIN class

Security controls and options

You can specify whether security is turned on for the whole WebSphere MQ subsystem, and whether you want to perform security checks at queue manager or queue-sharing group level. You can also control the number of user IDs checked for API-resource security.

Subsystem security

Subsystem security is a control that specifies whether any security checking is done for the whole queue manager. If you do not require security checking (for example, on a test system), or if you are satisfied with the level of security on all the resources that can connect to WebSphere MQ (including clients and channels), you can turn security checking off for the queue manager or queue-sharing group so that no further security checking takes place.

This is the only check that can turn security off completely and determine whether any other security checks are performed or not. That is, if you turn off checking for the queue manager or queue-sharing group, no other WebSphere MQ checking is done; if you leave it turned on, WebSphere MQ checks your security requirements for other WebSphere MQ resources.

Security can also be switched on or off for particular sets of resources, such as commands.

Queue manager or queue-sharing group level checking

Security can be implemented at queue manager level or at queue-sharing group level. If you implement security at queue-sharing group level, all the queue managers in the group share the same profiles. This means that there are fewer profiles to define and maintain, making security management easier. It also makes it easy to add a new queue manager to the queue-sharing group as it inherits the existing security profiles.

It is also possible to implement a combination of both if your installation requires it, for example, during migration or if you have one queue manager in the queue-sharing group that requires different levels of security to the other queue managers in the group.

Queue-sharing group level security

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

Queue-sharing group level security checking is performed for the entire queue-sharing group. It enables you to simplify security administration because it requires you to define fewer security profiles. The authorization of a user ID to use a particular resource is handled at the queue-sharing group level, and is independent of which queue manager that user ID is using to access the resource.

For example, say a server application runs under user ID SERVER and wants access to a queue called SERVER.REQUEST, and you want to run an instance of SERVER on each z/OS image in the sysplex. Rather than

permitting SERVER to open SERVER.REQUEST on each queue manager individually (queue manager level security), you can permit access once at the queue-sharing group level.

Queue-sharing group level security profiles can be used to protect all types of resource, whether local or shared.

Queue manager level security
Queue manager level security checking is performed at subsystem level using security profiles specific to that queue manager. This is how security is managed on Version 2.1 of MQSeries for OS/390, or earlier.

Queue manager level security profiles can be used to protect all types of resource, whether local or shared.

Combination of both levels
You can use a combination of both queue manager and queue-sharing group level security.

You can override queue-sharing group level security settings for a particular queue manager that is a member of that group. This means that you can perform a different level of security checks on an individual queue manager to those performed on the other queue managers in the group.

Controlling the number of user IDs checked

RESLEVEL is a RACF profile that controls the number of user IDs checked for WebSphere MQ resource security. Normally, when a user attempts to access a WebSphere MQ resource, RACF checks the relevant user ID or IDs to see if access is allowed to that resource. By defining a RESLEVEL profile you can control whether zero, one or, where applicable, two user IDs are checked.

These controls are done on a connection by connection basis, and last for the life of the connection.

There is only one RESLEVEL profile for each queue manager. Control is implemented by the access that a user ID has to this profile.

Resources you can protect

When a queue manager starts, or when instructed by an operator command, WebSphere MQ determines which resources you want to protect. You can control which security checks are performed for each individual queue manager. For example, you could implement a number of security checks on a production queue manager, but none on a test queue manager.

Connection security

Connection security checking is carried out either when an application program tries to connect to a queue manager by issuing an **MQCONN** or **MQCONNX** request, or when the channel initiator, or CICS or IMS adapter issues a connection request.

If you are using queue manager level security, you can turn connection security checking off for a particular queue manager, but if you do any user can connect to that queue manager.

For the CICS adapter, only the CICS address space user ID is used for the connection security check—not the individual CICS terminal user ID. For the IMS

adapter, when the IMS control or dependent regions connect to WebSphere MQ, the IMS address space user ID is checked. For the channel initiator, the user ID used by the channel initiator address space is checked.

Connection security checking can be turned on or off at either queue manager or queue-sharing group level.

API-resource security

Resources are checked when an application opens an object with an **MQOPEN** or an **MQPUT1** call. The access needed to open an object depends on what open options are specified when the queue is opened.

API-resource security is subdivided into these checks:
- Queue
- Process
- Namelist
- Alternate user
- Context

No security checks are performed when opening the queue manager object or when accessing storage class objects.

Queue security

Queue security checking controls who is allowed to open which queue, and what options they are allowed to open it with. For example, a user might be allowed to open a queue called PAYROLL.INCREASE.SALARY to browse the messages on the queue (via the MQOO_BROWSE option), but not to remove messages from the queue (via one of the MQOO_INPUT_* options). If you turn checking for queues off, any user can open any queue with any valid open option (that is, any valid MQOO_* option on an **MQOPEN** or **MQPUT1** call).

Queue security checking can be turned on or off at either queue manager or queue-sharing group level.

Process security

Process security checking is carried out when a user opens a process definition object. If you turn checking for processes off, any user can open any process.

Process security checking can be turned on or off at either queue manager or queue-sharing group level.

Namelist security

Namelist security checking is carried out when a user opens a namelist. If you turn checking for namelists off, any user can open any namelist.

Namelist security checking can be turned on or off at either queue manager or queue-sharing group level.

Alternate user security

Alternate user security controls whether one user ID can use the authority of another user ID to open a WebSphere MQ object.

For example:

- A server program running under user ID PAYSERV retrieves a request message from a queue that was put on the queue by user ID USER1.

- When the server program gets the request message, it processes the request and puts the reply back into the reply-to queue specified with the request message.
- Instead of using its own user ID (PAYSERV) to authorize opening the reply-to queue, the server can specify some other user ID, in this case, USER1. In this example, alternate user security would control whether user ID PAYSERV is allowed to specify user ID USER1 as an alternate user ID when opening the reply-to queue.

The alternate user ID is specified in the *AlternateUserId* field of the object descriptor (MQOD).

You can use alternate user IDs on any WebSphere MQ object, for example, processes or namelists. It does not affect the user ID used by any other resource managers, for example, for CICS security or for z/OS data set security.

If alternate user security is not active, any user can use any other user ID as an alternate user ID.

Alternate user security checking can be turned on or off at either queue manager or queue-sharing group level.

Context security

Context is information that is applicable to a particular message and is contained in the message descriptor (MQMD) that is part of the message. The context information comes in two sections:

Identity section
The user of the application that first put the message to a queue. It consists of the following fields:
- *UserIdentifier*
- *AccountingToken*
- *ApplIdentityData*

Origin section
The application that put the message on the queue where it is currently stored. It consists of the following fields:
- *PutApplType*
- *PutApplName*
- *PutDate*
- *PutTime*
- *ApplOriginData*

Applications can specify the context data when either an **MQPUT** or an **MQPUT1** call is made. This data might be generated by the application, it might be passed on from another message, or it might be generated by the queue manager by default. For example, context data can be used by server programs to check the identity of the requester, that is, did this message come from the correct application? Typically, the *UserIdentifier* field is used to determine the user ID of an alternate user.

You use context security to control whether the user can specify any of the context options on any **MQOPEN** or **MQPUT1** call. For information about the context options, see the *WebSphere MQ Application Programming Guide*; for descriptions of the message descriptor fields relating to context, see the *WebSphere MQ Application Programming Reference* manual.

If you turn context security checking off, any user can use any of the context options that the queue security allows.

Context security checking can be turned on or off at either queue, queue manager or queue-sharing group level.

Command security

Command security checking is carried out when a user issues an MQSC command from any of the sources described in "Issuing commands" on page 77. A separate check can be made on the resource specified by the command as described in "Command resource security".

If you turn off command checking, issuers of commands are not checked to see whether they have the authority to issue the command.

If MQSC commands are entered from a console, the console must have the z/OS SYS console authority attribute. Commands that are issued from the CSQINP1 or CSQINP2 data sets, or internally by the queue manager, are exempt from all security checking while those for CSQINPX use the user ID of the channel initiator address space. You should control who is allowed to update these data sets through normal data set protection.

Command security checking can be turned on or off at either queue manager or queue-sharing group level.

Command resource security

Some MQSC commands, for example defining a local queue, involve the manipulation of WebSphere MQ resources. When command resource security is active, each time a command involving a resource is issued, WebSphere MQ checks to see if the user is allowed to change the definition of that resource.

You can use command resource security to help enforce naming standards. For example, a payroll administrator might be allowed to delete and define only queues with names beginning "PAYROLL". If command resource security is inactive, no security checks are made on the resource that is being manipulated by the command. Do not confuse command resource security with command security; the two are independent.

Turning off command resource security checking does not affect the resource checking that is done specifically for other types of processing that do not involve commands.

Command resource security checking can be turned on or off at either queue manager or queue-sharing group level.

Channel security

When you are using channels, the security features available depend on which communications protocol you are going to use. If you use TCP, there are no security features provided with the communications protocol, although you can use SSL. If you are using APPC, you can flow user ID information from the sending MCA through the network to the destination MCA for verification.

For both protocols, you can specify which user IDs you want to check for security purposes, and how many. Again, the choices available to you depend on which

protocol you are using, what you specify when you define the channel, and the RESLEVEL settings for the channel initiator.

You can also write your own security exit programs to be called by the MCA.

Secure Sockets Layer (SSL)

You can also use the Secure Sockets Layer (SSL) to provide channel security. SSL support is provided with WebSphere MQ for z/OS Version 5 Release 3. SSL is an industry-standard protocol that provides a data security layer between application protocols and the communications layer, usually TCP/IP. SSL uses encryption techniques, digital signatures and digital certificates to provide message privacy, message integrity and mutual authentication between clients and servers.

For a detailed explanation of SSL, see the *WebSphere MQ Security* book.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 9. Where to find more information about security

Topic	Where to look
Setting up your security Channel security	*WebSphere MQ for z/OS System Setup Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
MQSC commands Defining channels	*WebSphere MQ Script (MQSC) Command Reference*
Channel security exits	*WebSphere MQ Intercommunication*
Secure Sockets Layer (SSL)	*WebSphere MQ Security*

Security

Chapter 8. Availability

There are several features of WebSphere MQ that are designed to increase system availability if the queue manager or channel initiator fails. These are discussed in the following sections:

Sysplex considerations

In a *sysplex* a number of z/OS operating system images collaborate in a single system image and communicate using a Coupling Facility. WebSphere MQ can use the facilities of the sysplex environment for enhanced availability.

Removing the affinities between a queue manager and a particular z/OS image allows a queue manager to be restarted on a different z/OS image in the event of an image failure. The restart mechanism can be manual, use ARM, or use system automation if you ensure the following:

- All page sets, logs, bootstrap data sets, code libraries, and queue manager configuration data sets must be defined on shared volumes.
- The subsystem definition must have sysplex scope and a unique name within the sysplex.
- The level of *early code* installed on every z/OS image at IPL time must be at the same level.
- TCP virtual IP addresses (VIPA) must be available on each TCP stack in the sysplex, and you must configure WebSphere MQ TCP listeners and inbound connections to use VIPAs rather than default host names.

For more information about VIPA, see the *z/OS Communications server IP User's Guide*, SC31-8780. For more information about using TCP in a sysplex, see *Redbook: TCP/IP in a sysplex*, SG24–5235.

You can additionally configure multiple WebSphere MQ queue managers running on different operating system images in a sysplex to operate as a queue-sharing-group, which can take advantage of shared queues and shared channels for higher availability and workload balancing.

Shared queues

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

In the queue-sharing group environment, an application can connect to any of the queue managers within the queue-sharing group. Because all the queue managers in the queue-sharing group can access the same set of shared queues, the application does not depend on the availability of a particular queue manager; any

queue manager in the queue-sharing group can service any queue. This gives greater availability if a queue manager stops because all the other queue managers in the queue-sharing group can continue processing the queue. For information about high availability of shared queues, see "Advantages of using shared queues" on page 17.

To further enhance the availability of messages in a queue-sharing group, WebSphere MQ detects if another queue manager in the group disconnects from the Coupling Facility abnormally, and completes units of work for that queue manager that are still pending, where possible. This is known as *peer recovery* and is described in "Peer recovery" on page 60.

Peer recovery cannot recover units of work that were in doubt at the time of the failure. You can use the Automatic Restart Manager (ARM) to restart all the systems involved in the failure (CICS, DB2, and WebSphere MQ for example), and to ensure that they are all restarted on the same new processor. This means that they can resynchronize, and gives rapid recovery of in-doubt units of work. This is described in "Using the z/OS Automatic Restart Manager (ARM)" on page 75.

Shared channels

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

In the queue-sharing group environment, WebSphere MQ provides functions that give high availability to the network. The channel initiator enables you to use networking products that balance network requests across a set of eligible servers and hide server failures from the network (for example, VTAM generic resources). WebSphere MQ uses a generic port for inbound requests so that attach requests can be routed to any available channel initiator in the queue-sharing group. This is described in "Shared channels" on page 19.

Shared outbound channels take the messages they send from a shared transmission queue. Information about the status of a shared channel is held in one place for the whole queue-sharing group level. This means that a channel can be restarted automatically on a different channel initiator in the queue-sharing group if the channel initiator, queue manager, or communications subsystem fails. This is called *peer channel recovery* and is described in "Shared outbound channels" on page 20.

WebSphere MQ network availability

WebSphere MQ messages are carried from queue manager to queue manager in a WebSphere MQ network using channels. You can change the configuration at a number of levels to improve the network availability of a queue manager, and the ability of a WebSphere MQ channel to detect a network problem and to reconnect.

TCP *Keepalive* is available for TCP/IP channels. It causes TCP to send packets periodically between sessions to detect network failures. The KAINT channel attribute determines the frequency of these packets for a channel.

AdoptMCA allows a channel, blocked in receive processing as a result of a network outage, to be terminated and replaced by a new connection request. You control AdoptMCA using the ADOPTMCA channel initiator parameter, set by the CSQ6CHIP macro.

ReceiveTimeout prevents a channel from being permanently blocked in a network receive call. The RCVTIME and RCVTMIN channel initiator parameters, set by the CSQ6CHIP macro, determine the receive timeout characteristics for channels, as a function of their heartbeat interval.

Using the z/OS Automatic Restart Manager (ARM)

WebSphere MQ for z/OS can be used in conjunction with the z/OS automatic restart manager (ARM). If a queue manager or a channel initiator has failed, ARM restarts it on the same z/OS image. If z/OS fails, a whole group of related subsystems and applications also fail. ARM can restart all the failed systems automatically, in a predefined order, on another z/OS image within the sysplex. This is called a cross-system restart.

ARM enables rapid recovery of in-doubt transactions in the shared queue environment. It also gives higher availability if you are not using queue-sharing groups.

You can use ARM to restart a queue manager on a different z/OS image within the sysplex in the event of z/OS failure.

To enable automatic restart:

- You must set up an ARM coupling data set.
- You must define the automatic restart actions that you want z/OS to perform in an *ARM policy*.
- The ARM policy must be started.

If you want to restart queue managers in different z/OS images automatically, every queue manager must be defined in each z/OS image on which that queue manager might be restarted, with a sysplex-wide unique 4-character subsystem name.

Using ARM with WebSphere MQ is described in the *WebSphere MQ for z/OS System Administration Guide*.

Using the z/OS Extended Recovery Facility (XRF)

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

WebSphere MQ can be used in an extended recovery facility (XRF) environment. All WebSphere MQ-owned data sets (executable code, BSDSs, logs, and page sets) must be on DASD shared between the active and alternate XRF processors.

If you use XRF for recovery, you must stop the queue manager on the active processor and start it on the alternate. For CICS, this can be done using the command list table (CLT) provided by CICS, or manually by the system operator. For IMS, this is a manual operation and must be done after the coordinating IMS system has completed the processor switch.

WebSphere MQ utilities must be completed or terminated before the queue manager can be switched to the alternate processor. Consider the effect of this potential interruption carefully when planning your XRF recovery plans.

Availability

Take care to prevent the queue manager starting on the alternate processor before the queue manager on the active processor terminates. A premature start can cause severe integrity problems in data, the catalog, and the log. Using global resource serialization (GRS) helps avoid the integrity problems by preventing simultaneous use of WebSphere MQ on the two systems. The BSDS must be included as a protected resource, and the active and alternate XRF processors must be included in the GRS ring.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 10. Where to find more information about availability

Topic	Where to look
Queue-sharing groups	Chapter 2, "Shared queues and queue-sharing groups", on page 13
System parameters	*WebSphere MQ for z/OS System Setup Guide*
Using the Automatic Restart Manager Utility programs	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*

Chapter 9. Creating and managing objects

This chapter discusses how to use the MQSC commands and utilities to create objects and manage your queue managers. It includes the following sections:

Issuing commands

WebSphere MQ for z/OS supports MQSC commands, which can be issued from the following sources:

- The z/OS console or equivalent (such as SDSF/TSO).
- The initialization input data sets.
- The supplied batch utility, CSQUTIL, processing a list of commands in a sequential data set.
- A suitably authorized application, by sending a command as a message to the command queue. This can be either:
  - A batch region program
  - A CICS application
  - An IMS application
  - A TSO application
  - An application program or utility on another WebSphere MQ system

Table 13 on page 82 summarizes the MQSC commands and the sources from which they can be issued.

Much of the functionality of these commands is available in a user-friendly way from the WebSphere MQ for z/OS operations and controls panels.

Changes made to the resource definitions of a queue manager using the commands (directly or indirectly) are preserved across restarts of the WebSphere MQ subsystem.

Private and global definitions

When you define an object on WebSphere MQ for z/OS, you can choose whether you want to share that definition with other queue managers (a *global* definition), or whether the object definition is to be used by one queue manager only (a *private* definition). This is called the object *disposition*.

Global definition

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

If your queue manager belongs to a queue-sharing group, you can choose to share any object definitions you make with the other members of the group. This means that an object has to be defined once only, reducing the total number of definitions required for the whole system.

Global object definitions are held in a *shared repository* (a DB2 shared database), and are available to all the queue managers in the queue-sharing group. These objects have a disposition of GROUP.

Private definition
If you want to create an object definition that is required by one queue manager only, or if your queue manager is not a member of a queue-sharing group, you can create object definitions that are not shared with other members of a queue-sharing group.

Private object definitions are held on page set zero of the defining queue manager. These objects have a disposition of QMGR.

You can create private definitions for all types of WebSphere MQ objects except CF structures (channels, namelists, process definitions, queues, queue managers, storage class definitions, and authentication information), and global definitions for all types of objects except queue managers.

WebSphere MQ automatically copies the definition of a group object to page set zero of each queue manager that uses it. You can alter the copy of the definition temporarily if you want, and WebSphere MQ allows you to refresh the page set copies from the repository copy if required. WebSphere MQ always tries to refresh the page set copies from the repository copy on start up (for channel commands, this is done when the channel initiator restarts). This ensures that the page set copies reflect the version on the repository, including any changes that were made when the queue manager was inactive. There are circumstances under which the refresh is not performed, for example:

- If a copy of the queue is open, a refresh that changes the usage of the queue will fail.
- If a copy of a queue has messages on it, a refresh that deletes that queue will fail.

In these circumstances, the refresh is not performed on that copy, but is performed on the copies on all other queue managers.

If the queue manager is shut down and then restarted stand-alone, any local copies of objects are deleted, unless for example, the queue has associated messages.

There is a third object disposition that applies to local queues only. This allows you to create shared queues. The definition for a shared queue is held on the shared repository and is available to all the queue managers in the queue-sharing group. In addition, the messages on a shared queue are also available to all the queue managers in the queue sharing group. This is described in Chapter 2, “Shared queues and queue-sharing groups”, on page 13. Shared queues have an object disposition of SHARED.

The following table summarizes the effect of the object disposition options for queue managers started stand-alone, and as a member of a queue-sharing group.

Disposition	Stand-alone queue manager	Member of a queue-sharing group
QMGR	Object definition held on page set zero.	Object definition held on page set zero.
GROUP	Not allowed.	Object definition held in the shared repository. Local copy held on page set zero of each queue manager in the group.

Disposition	Stand-alone queue manager	Member of a queue-sharing group
SHARED	Not allowed.	Queue definition held in the shared repository. Messages available to any queue manager in the group.

Manipulating global definitions

If you want to change the definition of an object that is held in the shared repository, you need to specify whether you want to change the version on the repository, or the local copy on page set zero. Use the object disposition as part of the command to do this.

Directing commands to different queue managers

You can choose to execute a command on the queue manager where it is entered, or on a different queue manager in the queue-sharing group. You can also choose to issue a particular command in parallel on all the queue managers in a queue-sharing group.

This is determined by the *command scope*. The command scope is used in conjunction with the object disposition to determine which version of an object you want to work with.

For example, you might want to alter some of the attributes of an object, the definition of which is held in the shared repository.

- You might want to change the version on one queue manager only, and not make any changes to the version on the repository or those in use by other queue managers.
- You might want to change the version in the shared repository for future users, but leave existing copies unchanged.
- You might want to change the version in the shared repository, but also want your changes to be reflected immediately on all the queue managers in the queue-sharing group that hold a copy of the object on their page set zero.

Use the command scope to specify whether the command is executed on this queue manager, another queue manager, or all queue managers. Use the object disposition to specify whether the object you are manipulating is in the shared repository (a group object), or is a local copy on page set zero (a queue manager object).

You do not have to specify the command scope and object disposition to work with a shared queue because every queue manager in the queue-sharing group sees the shared queue as a single queue.

MQSC command summary

Table 11 summarizes the MQSC commands that are available on WebSphere MQ for z/OS to alter, define, delete and display WebSphere MQ objects.

Table 11. Summary of the main MQSC commands

Object	ALTER	DEFINE	DISPLAY	DELETE
AUTHINFO	✔	✔	✔	✔
CFSTRUCT	✔	✔	✔	✔
CHANNEL	✔	✔	✔	✔

Table 11. Summary of the main MQSC commands (continued)

Object	ALTER	DEFINE	DISPLAY	DELETE
NAMELIST	✓	✓	✓	✓
PROCESS	✓	✓	✓	✓
QALIAS	✓	✓	✓	✓
QCLUSTER			✓	
QLOCAL	✓	✓	✓	✓
QMGR	✓		✓	
QMODEL	✓	✓	✓	✓
QREMOTE	✓	✓	✓	✓
QUEUE			✓	
STGCLASS	✓	✓	✓	✓

You can use the MQSC commands to manage other WebSphere MQ resources, and carry out other actions in addition to those summarized in Table 11 on page 79. Table 12 summarizes all the MQSC commands. An asterisk (*) indicates a complete resource, or individual commands within a resource, that are not available within WebSphere Application Server embedded messaging.

Table 12. Summary of all the MQSC commands

Resource	Task	Command
Authentication information objects	Create a new authentication information object Change an authentication information object's settings Display information about an authentication information object Delete an authentication information object	DEFINE AUTHINFO ALTER AUTHINFO DISPLAY AUTHINFO DELETE AUTHINFO
BSDS	Reestablish a dual bootstrap data set that had a data set error	RECOVER BSDS
Buffer pools	Define a buffer pool and the number of 4 KB buffers that it contains	DEFINE BUFFPOOL
*CF structures	Create a new CF structure Change a CF structure's settings Display information about a CF structure Display CF structure status information Backup a CF structure Recover a CF structure Delete a CF structure	DEFINE CFSTRUCT ALTER CFSTRUCT DISPLAY CFSTRUCT DISPLAY CFSTATUS BACKUP CFSTRUCT RECOVER CFSTRUCT DELETE CFSTRUCT
Channels	Create a new channel Change a channel's settings Start a channel Stop a channel Test a channel Reset channel sequence numbers Resolve in-doubt messages on a channel Display information about a channel Display channel status information Delete a channel	DEFINE CHANNEL ALTER CHANNEL START CHANNEL STOP CHANNEL PING CHANNEL RESET CHANNEL RESOLVE CHANNEL DISPLAY CHANNEL DISPLAY CHSTATUS DELETE CHANNEL
Channel initiators	Start a channel initiator Stop a channel initiator Display information about channel initiators	START CHINIT STOP CHINIT DISPLAY DQM

Table 12. Summary of all the MQSC commands (continued)

Resource	Task	Command
Channel listeners	Start a channel listener Stop a channel listener	START LISTENER STOP LISTENER
*Clusters	Refresh locally-held cluster information Perform special cluster operations Join a cluster Leave a cluster Display cluster information about the queue managers in a cluster	REFRESH CLUSTER RESET CLUSTER RESUME QMGR SUSPEND QMGR DISPLAY CLUSQMGR
Command servers	Start the command server Stop the command server Display command server attributes	START CMDSERV STOP CMDSERV DISPLAY CMDSERV
*IMS Tpipes	Reset sequence numbers for an IMS transaction pipe	RESET TPIPE
Logs	Copy the current active log to an archive log Control logging Display logging status	ARCHIVE LOG SET LOG SET ARCHIVE SUSPEND QMGR RESUME QMGR DISPLAY LOG DISPLAY ARCHIVE
Namelists	Create a new namelist Change a namelist's settings Display informaiton about a namelist Delete a namelist	DEFINE NAMELIST ALTER NAMELIST DISPLAY NAMELIST DELETE NAMELIST
Page sets	Define a page set and an associated buffer pool Display the current state of a page set	DEFINE PSID DISPLAY USAGE
Processes	Create a new process Change a process's settings Display information about a process Delete a process	DEFINE PROCESS ALTER PROCESS DISPLAY PROCESS DELETE PROCESS
Queues	Create a new queue Change a queue's settings Display information about a queue Display who is using a queue Move all the messages from one local queue to another Display and reset queue statistics Clear the messages from a local queue Delete a queue	DEFINE queues ALTER Queues DISPLAY QUEUE DISPLAY QSTATUS MOVE QLOCAL RESET QSTATS CLEAR QLOCAL DELETE Queues
Queue managers	Change a queue manager's settings Start a queue manager Stop a queue manager Display information about a queue manager Test a queue manager	ALTER QMGR START QMGR STOP QMGR DISPLAY QMGR PING QMGR
*Queue-sharing groups	Display information about a queue-sharing group	DISPLAY GROUP
Security	Refresh in-storage ESM tables Request security reverification for a user Change security settings Display security settings	REFRESH SECURITY RVERIFY SECURITY ALTER SECURITY DISPLAY SECURITY

Table 12. Summary of all the MQSC commands (continued)

Resource	Task	Command
Storage classes	Create a new storage class Change a storage class's settings Display information about a storage class Delete a storage class	DEFINE STGCLASS ALTER STGCLASS DISPLAY STGCLASS DELETE STGCLASS
System information	Control system Display system status	SET SYSTEM REFRESH QMGR DISPLAY SYSTEM DISPLAY USAGE
Threads	Display information about threads Resolve in-doubt units of recovery manually	DISPLAY THREAD RESOLVE INDOUBT
Traces	Start a WebSphere MQ trace Stop a WebSphere MQ trace Alter WebSphere MQ trace settings Display WebSphere MQ trace settings	START TRACE STOP TRACE ALTER TRACE DISPLAY TRACE

Table 13 shows, for each MQSC command, from where the command can be issued:

- CSQINP1 initialization input data set
- CSQINP2 initialization input data set
- z/OS console (or equivalent)
- SYSTEM.COMMAND.INPUT queue and command server (from applications, CSQUTIL, or the CSQINPX initialization input data set)

Table 13. Sources from which to run MQSC commands

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
ALTER AUTHINFO	No	Yes	Yes	Yes
ALTER CFSTRUCT	No	Yes	Yes	Yes
ALTER CHANNEL	No	Yes	Yes	Yes
ALTER NAMELIST	No	Yes	Yes	Yes
ALTER PROCESS	No	Yes	Yes	Yes
ALTER QALIAS	No	Yes	Yes	Yes
ALTER QLOCAL	No	Yes	Yes	Yes
ALTER QMGR	No	Yes	Yes	Yes
ALTER QMODEL	No	Yes	Yes	Yes
ALTER SECURITY	Yes	Yes	Yes	Yes
ALTER STGCLASS	No	Yes	Yes	Yes
ALTER TRACE	Yes	Yes	Yes	Yes
ARCHIVE LOG	Yes	Yes	Yes	Yes
BACKUP CFSTRUCT	No	No	Yes	Yes
CLEAR QLOCAL	No	Yes	Yes	Yes
DEFINE PROCESS	No	Yes	Yes	Yes
DEFINE QALIAS	No	Yes	Yes	Yes
DEFINE QLOCAL	No	Yes	Yes	Yes

Table 13. Sources from which to run MQSC commands (continued)

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
DEFINE QMODEL	No	Yes	Yes	Yes
DEFINE QREMOTE	No	Yes	Yes	Yes
DELETE AUTHINFO	No	Yes	Yes	Yes
DELETE CFSTRUCT	No	Yes	Yes	Yes
DELETE CHANNEL	No	Yes	Yes	Yes
DELETE NAMELIST	No	Yes	Yes	Yes
DELETE PROCESS	No	Yes	Yes	Yes
DELETE QALIAS	No	Yes	Yes	Yes
DELETE QLOCAL	No	Yes	Yes	Yes
DELETE QMODEL	No	Yes	Yes	Yes
DELETE QREMOTE	No	Yes	Yes	Yes
DELETE STGCLASS	No	Yes	Yes	Yes
DISPLAY ARCHIVE	Yes	Yes	Yes	Yes
DISPLAY AUTHINFO	No	Yes	Yes	Yes
DISPLAY CFSTATUS	No	No	Yes	Yes
DISPLAY CFSTRUCT	No	Yes	Yes	Yes
DISPLAY CHANNEL	No	Yes	Yes	Yes
DISPLAY CHSTATUS	No	Yes	Yes	Yes
DISPLAY CLUSQMGR	No	Yes	Yes	Yes
DISPLAY CMDSERV	Yes	Yes	Yes	Yes
DISPLAY DQM	No	Yes	Yes	Yes
DISPLAY GROUP	No	Yes	Yes	Yes
DISPLAY LOG	Yes	Yes	Yes	Yes
DISPLAY MAXSMSGS	No	Yes	Yes	Yes
DISPLAY NAMELIST	No	Yes	Yes	Yes
DISPLAY PROCESS	No	Yes	Yes	Yes
DISPLAY QALIAS	No	Yes	Yes	Yes
DISPLAY QCLUSTER	No	Yes	Yes	Yes
DISPLAY QLOCAL	No	Yes	Yes	Yes
DISPLAY QMGR	No	Yes	Yes	Yes
DISPLAY QMODEL	No	Yes	Yes	Yes
DISPLAY QREMOTE	No	Yes	Yes	Yes
DISPLAY QSTATUS	No	Yes	Yes	Yes
DISPLAY QUEUE	No	Yes	Yes	Yes
DISPLAY SECURITY	Yes	Yes	Yes	Yes
DISPLAY STGCLASS	No	Yes	Yes	Yes
DISPLAY SYSTEM	Yes	Yes	Yes	Yes
DISPLAY THREAD	No	Yes	Yes	Yes

Table 13. Sources from which to run MQSC commands (continued)

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
DISPLAY TRACE	Yes	Yes	Yes	Yes
DISPLAY USAGE	No	Yes	Yes	Yes
MOVE QLOCAL	No	Yes	Yes	Yes
PING CHANNEL	No	Yes	Yes	Yes
RECOVER BSDS	Yes	Yes	Yes	Yes
RECOVER CFSTRUCT	No	No	Yes	Yes
REFRESH CLUSTER	No	Yes	Yes	Yes
REFRESH QMGR	No	Yes	Yes	Yes
REFRESH SECURITY	No	Yes	Yes	Yes
RESET AUTHINFO	No	Yes	Yes	Yes
RESET BUFFPOOL	Yes	No	No	No
RESET CFSTRUCT	No	Yes	Yes	Yes
RESET CHANNEL	No	Yes	Yes	Yes
RESET CLUSTER	No	Yes	Yes	Yes
RESET MAXSMSGS	No	Yes	Yes	Yes
RESET NAMELIST	No	Yes	Yes	Yes
RESET PSID	Yes	No	No	No
RESET QSTATS	No	Yes	Yes	Yes
RESET STGCLASS	No	Yes	Yes	Yes
RESET TPIPE	No	No	Yes	Yes
RESOLVE CHANNEL	No	Yes	Yes	Yes
RESOLVE INDOUBT	No	Yes	Yes	Yes
RESUME QMGR	No	Yes	Yes	Yes
RVERIFY SECURITY	No	Yes	Yes	Yes
SET ARCHIVE	Yes	Yes	Yes	Yes
SET LOG	Yes	Yes	Yes	Yes
SET SYSTEM	Yes	Yes	Yes	Yes
START CHANNEL	No	Yes	Yes	Yes
START CHINIT	No	Yes	Yes	Yes
START CMDSERV	Yes	Yes	Yes	No
START LISTENER	No	Yes	Yes	Yes
START QMGR	No	No	Yes	No
START TRACE	Yes	Yes	Yes	Yes
STOP CHANNEL	No	Yes	Yes	Yes
STOP CHINIT	No	Yes	Yes	Yes
STOP CMDSERV	Yes	Yes	Yes	No
STOP LISTENER	No	Yes	Yes	Yes
STOP QMGR	No	No	Yes	Yes

Table 13. Sources from which to run MQSC commands (continued)

Command	CSQINP1	CSQINP2	z/OS console	Command input queue and server
STOP TRACE	Yes	Yes	Yes	Yes
SUSPEND QMGR	No	Yes	Yes	Yes

In the *WebSphere MQ Script (MQSC) Command Reference*, each command description identifies the sources from which that command can be run.

Initialization commands

Commands in the initialization input data sets are processed when WebSphere MQ is initialized on queue manager startup. Three types of command can be issued from the initialization input data sets:

- Commands to define WebSphere MQ entities that cannot be recovered (DEFINE BUFFPOOL and DEFINE PSID for example)

  These commands must reside in the data set identified by the DDname CSQINP1. They are processed before the restart phase of initialization. They cannot be issued through the console, operations and control panels, or an application program. The responses to these commands are written to the sequential data set that you refer to in the CSQOUT1 statement of the started task procedure.
- Commands to define WebSphere MQ objects that are recoverable after restart. These definitions must be specified in the data set identified by the DDname CSQINP2. They are stored in page set zero. CSQINP2 is processed after the restart phase of initialization. The responses to these commands are written to the sequential data set that you refer to in the CSQOUT2 statement of the started task procedure.
- Commands to manipulate WebSphere MQ objects. These commands must also be specified in the data set identified by the DDname CSQINP2. For example, the WebSphere MQ-supplied sample contains an ALTER QMGR command to specify a dead-letter queue for the subsystem. The response to these commands is written to the CSQOUT2 output data set.

Note: If WebSphere MQ objects are defined in CSQINP2, WebSphere MQ attempts to redefine them each time the queue manager is started. If the queues already exist, the attempt to define them fails. If you need to define your objects in CSQINP2, you can avoid this problem by using the REPLACE parameter of the DEFINE commands, however, this will override any changes that were made during the previous run of the queue manager.

Sample initialization data set members are supplied with WebSphere MQ for z/OS. They are described in "Sample definitions supplied with WebSphere MQ" on page 46.

Initialization commands for distributed queuing

You can also use the CSQINP2 initialization data set for the START CHINIT command, and follow it with a series of other commands to define your distributed queuing environment (for example, defining your channels). If you stop and restart the channel initiator however, CSQINP2 is not reprocessed, so WebSphere MQ provides a third initialization input data set, called CSQINPX, that you can choose to process as part of the channel initiator started task procedure.

The MQSC commands contained in the data set are executed at the end of channel initiator initialization, and output is written to the data set specified by the CSQOUTX DD statement. You might use the CSQINPX initialization data set to start listeners for example.

A sample channel initiator initialization data set member is supplied with WebSphere MQ for z/OS. It is described in “Sample definitions supplied with WebSphere MQ” on page 46.

The WebSphere MQ for z/OS utilities

WebSphere MQ for z/OS supplies a set of utility programs to help you perform various administrative tasks. These utilities:

- Perform™ backup, restoration, and reorganization tasks
- Issue commands and process object definitions
- Generate data-conversion exits
- Modify the bootstrap data set
- List information about the logs
- Print the logs
- Set up DB2 tables and other DB2 utilities
- Process messages on the dead-letter queue

The CSQUTIL utility

The CSQUTIL utility program is provided with WebSphere MQ for z/OS to help you perform backup, restoration, and reorganization tasks, and to issue commands and process object definitions. Through this utility program, you can invoke the following functions:

COMMAND
To issue MQSC commands, to record object definitions, and to make client-channel definition files.

COPY To read the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set, and put them into a sequential file and retain the original queue.

COPYPAGE
To copy whole page sets to larger page sets.

EMPTY
To delete the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set, retaining the definitions of the queues.

FORMAT
To format WebSphere MQ for z/OS page sets.

LOAD
To restore the contents of a named WebSphere MQ for z/OS message queue or the contents of all the queues of a named page set from a sequential file created by the COPY function.

PAGEINFO
To extract page set information from one or more page sets.

RESETPAGE
To copy whole page sets to other page set data sets and reset the log information in the copy.

SCOPY
To copy the contents of a queue to a data set while the queue manager is offline.

SDEFS
To produce a set of define commands for objects while the queue manager is offline.

The data conversion exit utility

The WebSphere MQ for z/OS data conversion exit utility (CSQUCVX) runs as a stand-alone utility to create data conversion exit routines.

The change log inventory utility

The WebSphere MQ for z/OS change log inventory utility program (CSQJU003) runs as a stand-alone utility to change the bootstrap data set (BSDS). The utility can be used to:

- Add or delete active or archive log data sets
- Supply passwords for archive logs

The print log map utility

The WebSphere MQ for z/OS print log map utility program (CSQJU004) runs as a stand-alone utility to list the following information:

- Log data set name and log RBA association for both copies of all active and archive log data sets. If dual logging is not active, there is only one copy of the data sets.
- Active log data sets available for new log data.
- Contents of the queue of checkpoint records in the bootstrap data set (BSDS).
- Contents of the archive log command history record.
- System and utility time stamps.

The log print utility

The log print utility program (CSQ1LOGP) is run as a stand-alone utility. You can run the utility specifying:

- A bootstrap data set (BSDS)
- Active logs (with no BSDS)
- Archive logs (with no BSDS)

The queue-sharing group utility

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

The queue-sharing group utility program (CSQ5PQSG) runs as a stand-alone utility to set up DB2 tables and perform other DB2 tasks required for queue-sharing groups.

The active log preformat utility

The active log preformat utility (CSQJUFMT) formats active log data sets before they are used by a queue manager. If the active log data sets are preformatted by the utility, log write performance is improved on the queue manager's first pass through the active logs.

The dead-letter queue handler utility

The dead-letter queue handler utility program (CSQUDLQH) runs as a stand-alone utility. It checks messages that are on the dead-letter queue and processes them according to a set of rules that you supply to the utility.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 14. Where to find more information about creating and managing objects

Topic	Where to look
Using initialization input data sets System parameter macros Installation verification program	*WebSphere MQ for z/OS System Setup Guide*
Writing administration programs Operations and control panels Utility programs	*WebSphere MQ for z/OS System Administration Guide*
Queue indexes MQSC commands	*WebSphere MQ Script (MQSC) Command Reference*
WebSphere MQ clusters	*WebSphere MQ Queue Manager Clusters*
WebSphere MQ events	*WebSphere MQ Event Monitoring*

Chapter 10. Monitoring and statistics

WebSphere MQ supplies facilities for monitoring the system and collecting statistics. These are discussed in the following sections:

WebSphere MQ trace

WebSphere MQ supplies a trace facility that can be used to gather the following information while the queue manager is running:

Performance statistics

The statistics trace gathers the following information to help you monitor performance and tune your system:

- Counts of different MQI requests (message manager statistics)
- Counts of different object requests (data manager statistics)
- Information about DB2 usage (DB2 manager statistics)
- Information about Coupling Facility usage (Coupling Facility manager statistics)
- Information about buffer pool usage (buffer manager statistics)
- Information about logging (log manager statistics)
- Information about storage usage (storage manager statistics)
- Information about lock requests (lock manager statistics)

Accounting data

- The accounting trace gathers information about the CPU time spent processing MQI calls and about the number of **MQPUT** and **MQGET** requests made by a particular user.
- WebSphere MQ can also gather information about each task using WebSphere MQ. This data is gathered as a thread-level accounting record. For each thread, WebSphere MQ also gathers information about each queue used by that thread.

The data generated by the trace is sent to the System Management Facility (SMF) or the generalized trace facility (GTF).

Events

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

WebSphere MQ events provide information about errors, warnings, and other significant occurrences in a queue manager. By incorporating these events into your own system management application, you can monitor the activities across many queue managers, for multiple WebSphere MQ applications. In particular, you can monitor all the queue managers in your system from a single queue manager.

Events can be reported through a user-written reporting mechanism to an administration application that supports the presentation of the events to an

operator. Events also enable applications acting as agents for other administration networks, for example NetView®, to monitor reports and create the appropriate alerts.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 15. Where to find more information about monitoring and statistics

Topic	Where to look
WebSphere MQ trace	*WebSphere MQ for z/OS System Setup Guide*
Trace commands	*WebSphere MQ Script (MQSC) Command Reference*
WebSphere MQ events	*WebSphere MQ Event Monitoring*

Part 3. WebSphere MQ and other products

Chapter 11. WebSphere MQ and CICS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

This chapter discusses how WebSphere MQ works with CICS. The CICS adapter and the CICS bridge allow you to connect your queue manager to CICS.

- The CICS adapter enables CICS applications to use the MQI.
- The CICS bridge enables applications to run a CICS program or transaction that does not use the MQI. This means that you can use your legacy applications with WebSphere MQ, without the need to rewrite them.

These topics are described in the following sections:

The CICS adapter

The CICS adapter connects a CICS subsystem to a queue manager, enabling CICS application programs to use the MQI.

You can start and stop CICS and WebSphere MQ independently, and you can establish or terminate a connection between them at any time. You can also allow CICS to connect to WebSphere MQ automatically.

The CICS adapter provides two main facilities:

- A set of control functions for use by system programmers and administrators to manage the adapter.
- MQI support for CICS applications.

In a CICS multiregion operation or intersystem communication (ISC) environment, each CICS address space can have its own attachment to the queue manager subsystem. A single CICS address space can be connected to only one queue manager at a time. However, multiple CICS address spaces can connect to the same queue manager.

You can use WebSphere MQ with the CICS Extended Recovery Facility (XRF) to aid recovery from a CICS error.

The CICS adapter is supplied with WebSphere MQ and runs as a CICS External Resource Manager. WebSphere MQ also provides CICS transactions to manage the interface.

The CICS adapter uses standard CICS command-level services where required, for example, EXEC CICS WAIT and EXEC CICS ABEND. A portion of the CICS adapter runs under the control of the transaction issuing the messaging requests. Therefore, these calls for CICS services appear to be issued by the transaction.

Control functions

The CICS adapter's control functions (the CKQC transaction) let you manage the connections between CICS and WebSphere MQ dynamically. These functions can be invoked using the CICS adapter panels, from the command line, or from a CICS application. You can use the adapter's control function to:

- Start or stop a connection to a queue manager.
- Modify the current connection. For example, you can reset the connection statistics, change the adapter's trace ID number, and enable or disable the API-crossing exit.
- Display the current status of a connection and the statistics associated with that connection.
- Start or stop an instance of the task initiator transaction, CKTI. ("Task initiator transaction" is CICS terminology; in WebSphere MQ terminology, this is a trigger monitor.)
- Display details of the current instances of CKTI.
- Display details of the CICS tasks currently using the adapter.

These functions and the different methods of invoking them are described in the *WebSphere MQ for z/OS System Administration Guide*.

MQI support

The CICS adapter implements the MQI for use by CICS application programs. The MQI calls, and how they are used, are described in the *WebSphere MQ Application Programming Guide*. The adapter also supports an *API-crossing exit*, (see "The API-crossing exit" on page 99), and a trace facility.

All application programs that run under CICS must have the supplied *API stub program* called CSQCSTUB link-edited with them if they are to access WebSphere MQ, unless the program is using dynamic calls. This stub provides the application with access to all MQI calls. (For information about calling the CICS stub dynamically, see the *WebSphere MQ Application Programming Guide*.)

For performance, the CICS adapter can handle up to eight MQI calls concurrently. For transaction integrity, the adapter supports syncpointing under the control of the CICS syncpoint manager, so that units of work can be committed or backed out as required. The adapter also supports security checking of WebSphere MQ resources when used with an appropriate security management product, such as RACF. The adapter provides high availability with automatic reconnection after a queue manager termination, and automatic resource resynchronization after a restart. It also features an alert monitor that responds to unscheduled events such as a shut down of the queue manager.

Adapter components

Figure 15 on page 97 shows the relationship between CICS, the CICS adapter, and WebSphere MQ. CICS and the adapter share the same address space; the queue manager executes in its own address space.

Part of the adapter is a CICS task-related user exit that communicates with the WebSphere MQ message manager. CICS management modules call the exit directly; application programs call it through the supplied API stub program (CSQCSTUB). Task-related user exits and stub programs are described in the *CICS Customization Guide*.

Each CKTI transaction is normally in an **MQGET WAIT** state, ready to respond to any trigger messages that are placed on its initiation queue.

The adapter management interface provides the operation and control functions described in the *WebSphere MQ for z/OS System Administration Guide*.

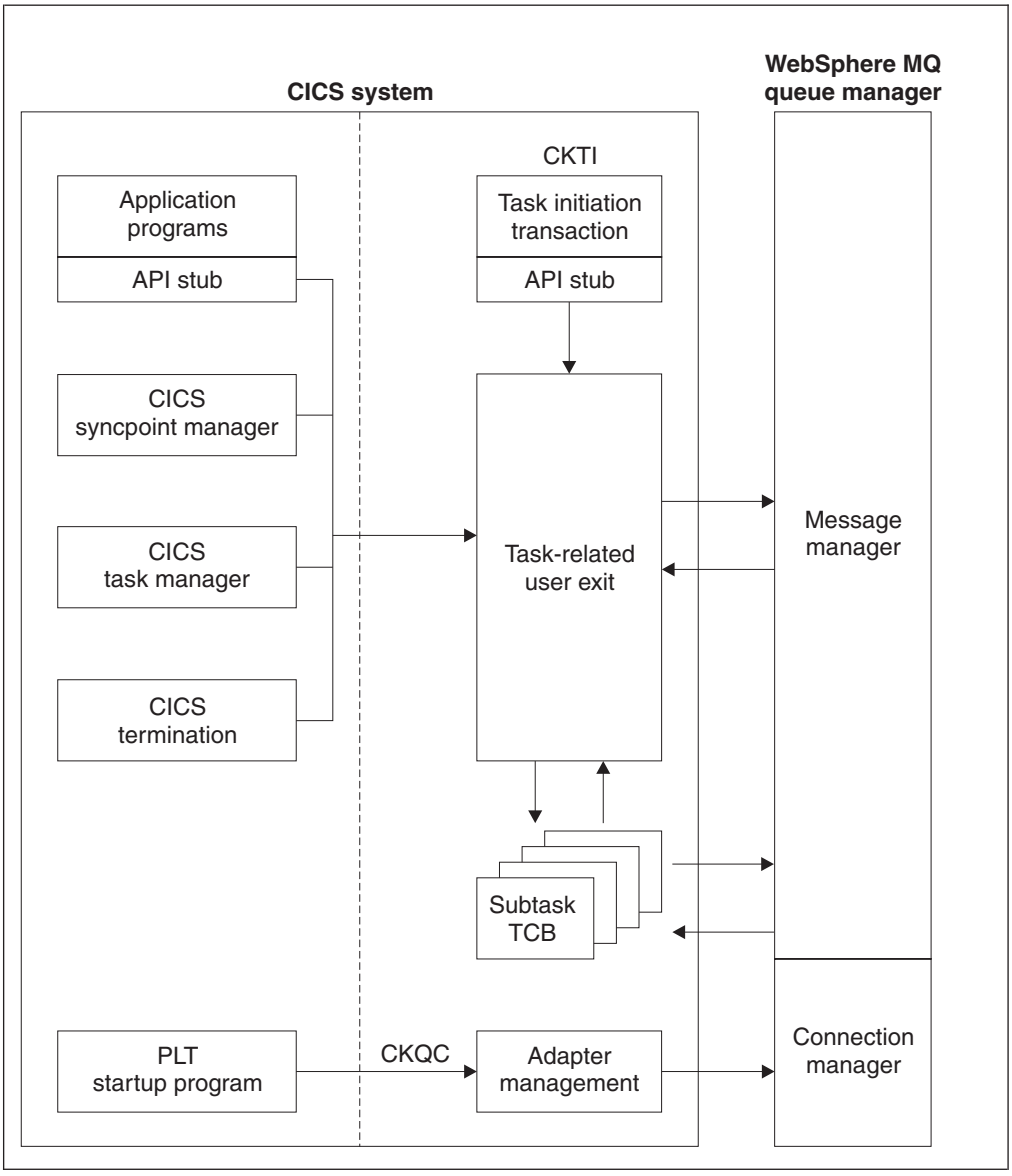

Figure 15. How CICS, the CICS adapter, and a WebSphere MQ queue manager are related

Alert monitor

The *alert monitor* transaction, CKAM, handles unscheduled events—known as *pending events*—that occur as a result of connect requests to instances of WebSphere MQ. The alert monitor generates messages that are sent to the system console.

There are two kinds of pending events:

1. **Deferred connection**
   If CICS tries to connect to WebSphere MQ before the queue manager is started, a *pending event* called a *deferred connection* is activated. When the queue manager is started, a connection request is issued by the CICS adapter, a connection is made, and the pending event is canceled.

   There can be multiple deferred connections, one of which will be connected when the queue manager is started.

2. **Termination notification**
   When a connection is successfully made to WebSphere MQ, a pending event called *termination notification* is created. This pending event expires when:
   - The queue manager shuts down normally with MODE(QUIESCE). The alert monitor issues a quiesce request on the connection.
   - The queue manager shuts down with MODE(FORCE) or terminates abnormally. After an abnormal termination, the CICS adapter waits for ten seconds and then tries a connect call. This enables the CICS system to be automatically reconnected to the queue manager when the latter is restarted.
   - The connection is shut down from the CKQC transaction.

The maximum number of pending events that can be handled is 99. If this limit is reached, no more events can be created until at least one current event expires.

The alert monitor terminates itself when all pending events have expired. It is subsequently restarted automatically by any new connect request.

Auto-reconnect

When CICS is connected to WebSphere MQ and the queue manager terminates, the CICS adapter tries to issue a connect request ten seconds after the stoppage has been detected. This request uses the same connect parameters that were used in the previous connect request. If the queue manager has not been restarted within the ten seconds, the connect request is deferred until the queue manager is restarted later.

Task initiator

CKTI is a WebSphere MQ-supplied CICS transaction that starts a CICS transaction when a WebSphere MQ trigger message is read, for example when a message is put onto a specific queue.

When a message is put onto an application message queue, a trigger is generated if the trigger conditions are met. The queue manager then writes a message, containing user-defined data, known as a *trigger message*, to the initiation queue that has been specified for that message queue. In a CICS environment, an instance of CKTI can be set up to monitor an initiation queue and to retrieve the trigger messages from it as they arrive. CKTI starts another CICS transaction, (specified using the DEFINE PROCESS command), which typically reads the message from

the application message queue and then processes it. The process must be named on the application queue definition, not the initiation queue.

Each copy of CKTI services a single initiation queue. To start or stop a copy of CKTI, you must supply the name of the queue that this CKTI is to serve, or is serving. You cannot start more than one instance of CKTI against the same initiation queue from a single CICS subsystem.

At CICS system initialization or at connect time, you can define a default initiation queue. If you issue a STARTCKTI or a STOPCKTI command without specifying an initiation queue, these commands are automatically interpreted as referring to the default initiation queue.

Notes:

1. If you are using version 4.1 of CICS, any transaction entries processed by CKTI, for example EXEC CICS START, are locked by the CKTI task until it terminates. Any attempt to CEDA INSTALL such entries after altering them will fail: CEDA rejects the install request because the transaction entry is being used by another task.

   In this situation, you must stop the CKTI task using the CICS adapter, and restart it after the CEDA install.
2. This restriction also applies to intersystem connection (ISC) and multi-region operation (MRO) links. For example, if CKTI has started a remote transaction, a connection cannot be reinstalled until CKTI has been stopped.

Multitasking

The CICS adapter optimizes the performance of a CICS to WebSphere MQ connection by exploiting multi-processors and by removing work from the main CICS task control block (TCB), allowing multiple MQI calls to be handled concurrently.

The adapter enables some MQI calls to be executed under subtasks, rather than under the main CICS TCB that runs the application code. All the CICS adapter administration code, including connection and disconnection from WebSphere MQ, runs under the main CICS TCB.

The adapter attaches up to eight z/OS subtasks (TCBs) to be used by this CICS system. *You cannot modify this number*. Each subtask makes a connect call to WebSphere MQ. Each CICS system connected takes up nine of the connections specified on the CTHREAD system parameter. This means that you must increase the value specified for CTHREAD by nine for each CICS system connected. MQI calls can flow over those connections. When the main connection is terminated, the subtasks are disconnected and terminated automatically.

The API-crossing exit

WebSphere MQ provides an API-crossing exit for use with the CICS adapter; it runs in the CICS address space. You can use this exit to intercept MQI calls as they are being run, for monitoring, testing, maintenance, or security purposes.

Using the API-crossing exit degrades WebSphere MQ performance. You should plan your use of it carefully.

CICS adapter conventions

There are a number of conventions that must be observed in applications using the adapter.

Temporary storage queue names

The CICS adapter display function uses two temporary storage queues (MAIN) for each invoking task to store the output data for browsing. The names of the queues are *tttt*CKRT and *tttt*CKDP, where *tttt* is the terminal identifier of the terminal from which the display function is requested.

Do not try to access these queues.

MQGET

When the CICS adapter puts a task on a CICS wait because the WAIT option was used with the **MQGET** call and there was no message available, the RESOURCE NAME used is GETWAIT and the RESOURCE_TYPE is MQSeries.

When the CICS adapter puts a task on a CICS wait because of a need to perform task switching the RESOURCE NAME used is TASKSWCH and the RESOURCE_TYPE is MQSeries.

ENQUEUE names

The CICS adapter uses the name:

```
CSQ.genericapplid(8).QMGR
```

to issue CICS ENQ and CICS DEQ calls during processing, for example, starting and stopping the connection.

You should not use similar names for CICS ENQ or DEQ purposes.

The CICS bridge

The WebSphere MQ-CICS bridge is the component of WebSphere MQ for z/OS that allows direct access from WebSphere MQ applications to applications on your CICS system. In bridge applications there are no WebSphere MQ calls within the CICS application (the bridge enables *implicit MQI support*). This means that you can re-engineer legacy applications that were controlled by 3270-connected terminals to be controlled by WebSphere MQ messages, without having to rewrite, recompile, or re-link them.

WebSphere MQ applications use the CICS header (the MQCIH structure) in the message data to ensure that the applications can execute as they did when driven by nonprogrammable terminals.

The bridge enables an application that is not running in a CICS environment to run a *program* or *transaction* on CICS and get a response back. This non-CICS application can be run from any environment that has access to a WebSphere MQ network that encompasses WebSphere MQ for z/OS.

A *program* is a CICS program that can be invoked using the EXEC CICS LINK command. It must conform to the DPL subset of the CICS API, that is, it must not use CICS terminal or syncpoint facilities.

A *transaction* is a CICS transaction designed to run on a 3270 terminal. This transaction can use BMS or TC commands. It can be conversational or part of a pseudoconversation. It is permitted to issue syncpoints. For information about the transactions that can be run, see the *CICS Internet and External Interfaces Guide*.

When to use the CICS bridge

The CICS bridge allows an application to run a single CICS program or a 'set' of CICS programs (often referred to as a unit of work). It caters for the application that waits for a response to come back before it runs the next CICS program (synchronous processing) and for the application that requests one or more CICS programs to run, but doesn't wait for a response (asynchronous processing).

The CICS bridge also allows an application to run a 3270-based CICS transaction, without knowledge of the 3270 data stream.

The CICS bridge uses standard CICS and WebSphere MQ security features and can be configured to authenticate, trust, or ignore the requestor's user ID.

Given this flexibility, there any many instances where the CICS bridge can be used. For example, when you want:

- To write a new WebSphere MQ application that needs access to logic or data (or both) that reside on your CICS server.
- To run CICS programs from a Lotus® Notes™ application.
- To access your CICS applications from
  - Your WebSphere MQ Classes for Java™ client application
  - A web browser using the WebSphere MQ Internet gateway

System configuration for the CICS bridge

When you are setting your system up, you should ensure that:

- Both WebSphere MQ and CICS are running in the same z/OS image.

- The WebSphere MQ request queue is local to the CICS bridge, however the response queue can be local or remote.
- The CICS bridge tasks run in the same CICS as the bridge monitor. The user programs can be in the same or a different CICS system.
- The WebSphere MQ-CICS adapter is enabled.

Running CICS DPL programs

Data necessary to run the program is provided in the WebSphere MQ message. The bridge builds a COMMAREA from this data, and runs the program using EXEC CICS LINK. Figure 16 shows the step sequence taken to process a single message to run a CICS DPL program:

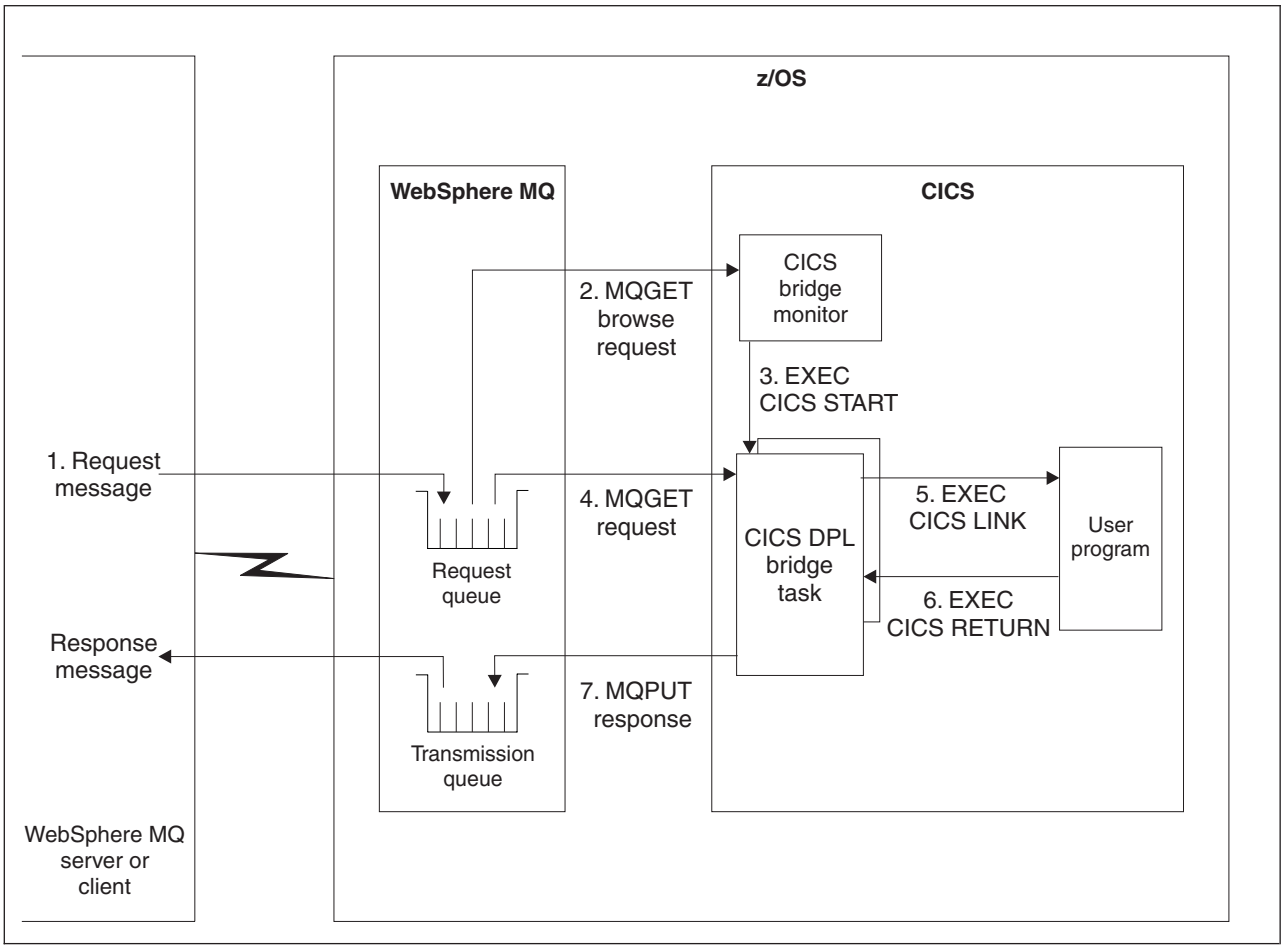

Figure 16. Components and data flow to run a CICS DPL program

The following takes each step in turn, and explains what takes place:

1. A message, with a request to run a CICS program, is put on the request queue.
2. The CICS bridge monitor task, which is constantly browsing the queue, recognizes that a 'start unit of work' message is waiting (*CorrelId*=MQCI_NEW_SESSION).
3. Relevant authentication checks are made, and a CICS DPL bridge task is started with the appropriate authority.
4. The CICS DPL bridge task removes the message from the request queue.
5. The CICS DPL bridge task builds a COMMAREA from the data in the message and issues an EXEC CICS LINK for the program requested in the message.

6. The program returns the response in the COMMAREA used by the request.
7. The CICS DPL bridge task reads the COMMAREA, creates a message, and puts it on the reply-to queue specified in the request message. All response messages (normal and error, requests and replies) are put to the reply-to queue with default context.
8. The CICS DPL bridge task ends.

A unit of work can be just a single user program, or it can be multiple user programs. There is no limit to the number of messages you can send to make up a unit of work.

In this scenario, a unit of work made up of many messages works in the same way, with the exception that the CICS bridge task waits for the next request message in the final step unless it is the last message in the unit of work.

Running CICS 3270 transactions

Data necessary to run the transaction is provided in the WebSphere MQ message. The CICS transaction runs as if it has a real 3270 terminal, but instead uses one or more MQ messages to communicate between the CICS transaction and the WebSphere MQ application

Unlike traditional 3270 emulators, the bridge does not work by replacing the VTAM flows with WebSphere MQ messages. Instead, the message consists of a number of parts called vectors, each of which corresponds to an EXEC CICS request. Therefore the application is talking directly to the CICS transaction, rather than through an emulator, using the actual data used by the transaction (known as application data structures or ADSs).

Figure 17 on page 104 shows the step sequence taken to process a single message to run a CICS 3270 transaction.

The following takes each step in turn, and explains what takes place:

1. A message, with a request to run a CICS transaction, is put on the request queue.
2. The CICS bridge monitor task, which is constantly browsing the queue, recognizes that a 'start unit of work' message is waiting (*CorrelId*=MQCI_NEW_SESSION).
3. Relevant authentication checks are made, and a CICS 3270 bridge task is started with the appropriate authority.
4. The MQ-CICS bridge exit removes the message from the queue and changes task to run a user transaction
5. Vectors in the message provide data to answer all terminal related input EXEC CICS requests in the transaction.
6. Terminal related output EXEC CICS requests result in output vectors being built.
7. The MQ-CICS bridge exit builds all the output vectors into a single message and puts this on the reply-to queue.
8. The CICS 3270 bridge task ends.

Note: The CICS bridge exit is a WebSphere MQ supplied CICS exit associated with the bridge transaction.

WebSphere MQ and CICS

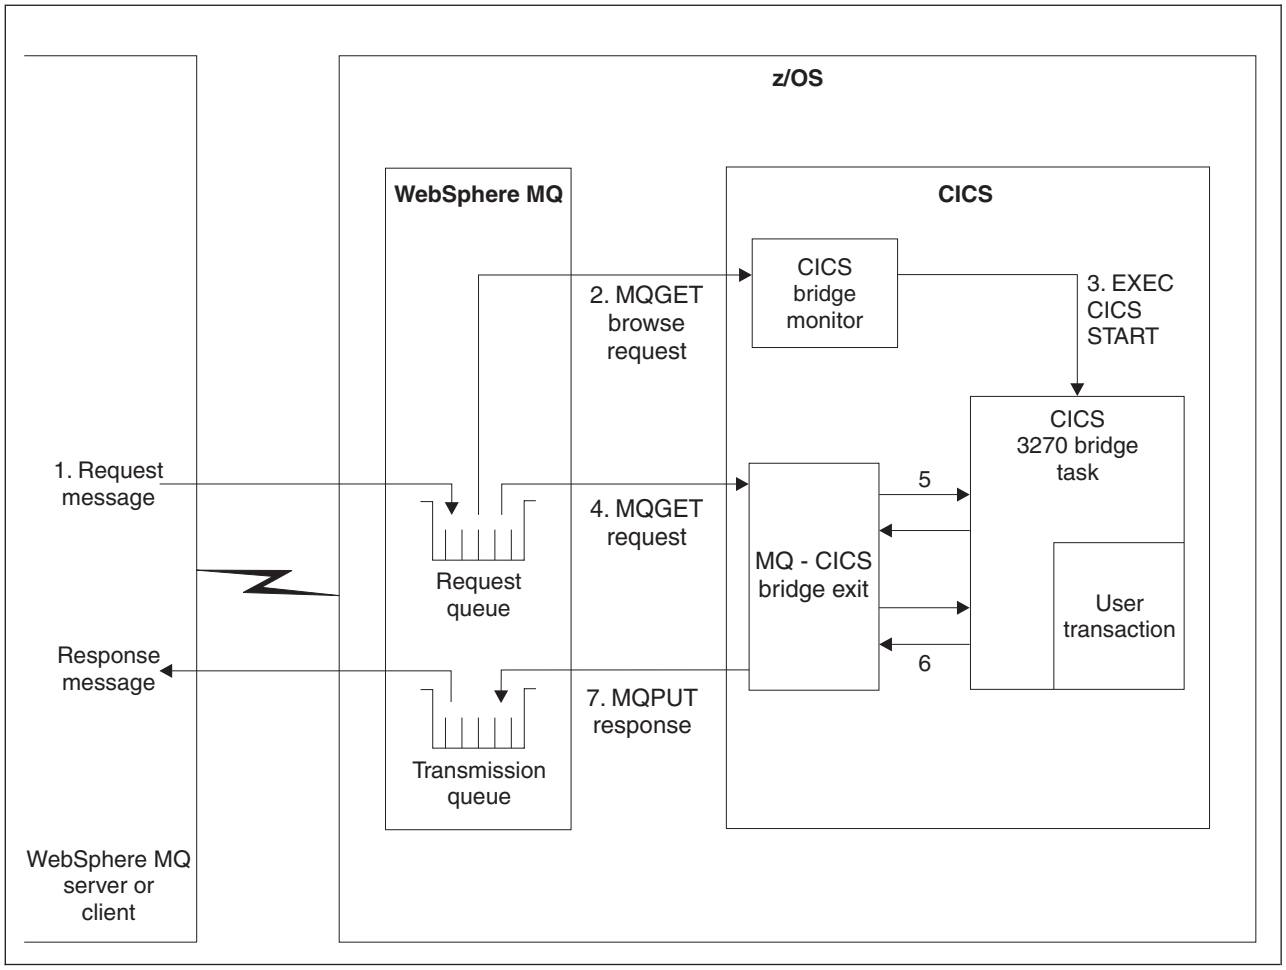

Figure 17. Components and data flow to run a CICS 3270 transaction

A traditional CICS application usually consists of one or more transactions linked together as a pseudoconversation. In general, each transaction is started by the 3270 terminal user entering data onto the screen and pressing an AID key. This model of application can be emulated by a WebSphere MQ application. A message is built for the first transaction, containing information about the transaction, and input vectors. This is put on the queue. The reply message consists of the output vectors, the name of the next transaction to be run, and a token that is used to represent the pseudoconversation. The WebSphere MQ application builds a new input message, with the transaction name set to the next transaction and the facility token set to the value returned on the previous message. Vectors for this second transaction are added to the message, and the message put on the queue. This process is continued until the application ends.

An alternative approach to writing CICS applications is the conversational model. In this model, the original message might not contain all the data to run the transaction. If the transaction issues a request that cannot be answered by any of the vectors in the message, a message is put onto the reply-to queue requesting more data. The WebSphere MQ application gets this message and puts a new message back to the queue with a vector to satisfy the request.

For more information about this, see the *CICS Internet and External Interfaces Guide*.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 16. Where to find more information about using CICS with WebSphere MQ

Topic	Where to look
System parameters Setting up the CICS adapter Setting up the CICS bridge	*WebSphere MQ for z/OS System Setup Guide*
Operating the CICS adapter Operating the CICS bridge	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
Writing CICS applications API-crossing exit	*WebSphere MQ Application Programming Guide*
Writing CICS bridge applications	*CICS Internet and External Interfaces Guide*

Chapter 12. WebSphere MQ and IMS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

This chapter discusses how WebSphere MQ works with IMS. The IMS adapter and the IMS bridge allow you to connect your queue manager to IMS.

- The IMS adapter enables IMS applications to use the MQI.
- The IMS bridge enables applications to run an IMS application that does not use the MQI. This means that you can use your legacy applications with WebSphere MQ, without the need to rewrite them.

These topics are described in the following sections:

The IMS adapter

The WebSphere MQ adapters enable different application environments to send and receive messages through a message queuing network. The IMS adapter is the interface between IMS application programs and a WebSphere MQ subsystem. It makes it possible for IMS application programs to use the MQI.

The IMS adapter receives and interprets requests for access to WebSphere MQ using the External Subsystem Attach Facility (ESAF) provided by IMS. This facility is described in the *IMS Customization Guide*. Usually, IMS connects to WebSphere MQ automatically without operator intervention.

The IMS adapter provides access to WebSphere MQ resources for programs running in:

- Task (TCB) mode
- Problem state
- Non-cross-memory mode
- Non-access register mode

The adapter provides a connection thread from an application task control block (TCB) to WebSphere MQ.

The adapter supports a two-phase commit protocol for changes made to resources owned by WebSphere MQ with IMS acting as the syncpoint coordinator.

The adapter also provides a trigger monitor transaction (CSQQTRMN). This is described in "The IMS trigger monitor" on page 108.

You can use WebSphere MQ with the IMS Extended Recovery Facility (XRF) to aid recovery from a IMS error. For more information about XRF, see the *IMS Administration Guide: System* manual.

Using the adapter

The application programs and the IMS adapter run in the same address space. The queue manager is separate, in its own address space.

Each program that issues one or more MQI calls must be link-edited to a suitable IMS language interface module, and, unless it uses dynamic MQI calls, the WebSphere MQ-supplied API stub program, CSQQSTUB. When the application issues an MQI call, the stub transfers control to the adapter through the IMS external subsystem interface, which manages the processing of the request by the message queue manager.

System administration and operation with IMS

An authorized IMS terminal operator can issue IMS commands to control and monitor the connection to WebSphere MQ. However, the IMS terminal operator has no control over the WebSphere MQ address space. For example, the operator cannot shut down WebSphere MQ from an IMS address space.

The IMS trigger monitor

The IMS trigger monitor (CSQQTRMN) is a WebSphere MQ-supplied IMS application that starts an IMS transaction when a WebSphere MQ event occurs, for example, when a message is put onto a specific queue.

How it works

When a message is put onto an application message queue, a trigger is generated if the trigger conditions are met. The queue manager then writes a message (containing some user-defined data), known as a *trigger message*, to the initiation queue that has been specified for that message queue. In an IMS environment, an instance of CSQQTRMN can be started to monitor an initiation queue and to retrieve the trigger messages from it as they arrive. Typically, CSQQTRMN schedules another IMS transaction by an INSERT (ISRT) to the IMS message queue. The started IMS application reads the message from the application message queue and then processes it. CSQQTRMN must run as a non-message BMP.

Each copy of CSQQTRMN services a single initiation queue. Once started, the trigger monitor runs until WebSphere MQ or IMS ends.

The APPLCTN macro for CSQQTRMN must specify SCHDTYP=PARALLEL.

Because the trigger monitor is a batch-oriented BMP, IMS transactions started by the trigger monitor will contain:

- Blanks in the LTERM field of the IOPCB
- The PSB name of the trigger monitor BMP in the Userid field of the IOPCB

If the target IMS transaction is RACF protected, you might need to define CSQQTRMN as a user ID to RACF.

The IMS bridge

The WebSphere MQ-IMS bridge is the component of WebSphere MQ for z/OS that allows direct access from WebSphere MQ applications to applications on your IMS system (the bridge enables *implicit MQI support*). This means that you can re-engineer legacy applications that were controlled by 3270-connected terminals to be controlled by WebSphere MQ messages, without having to rewrite, recompile, or re-link them. The bridge is an IMS *Open Transaction Manager Access* (OTMA) client.

In bridge applications there are no WebSphere MQ calls within the IMS application. The application gets its input using a GET UNIQUE (GU) to the IOPCB and sends its output using an ISRT to the IOPCB. WebSphere MQ applications use the IMS header (the MQIIH structure) in the message data to ensure that the applications can execute as they did when driven by nonprogrammable terminals. If you are using an IMS application that processes multi-segment messages, note that all segments should be contained within one WebSphere MQ message.

The IMS bridge is illustrated in Figure 18.

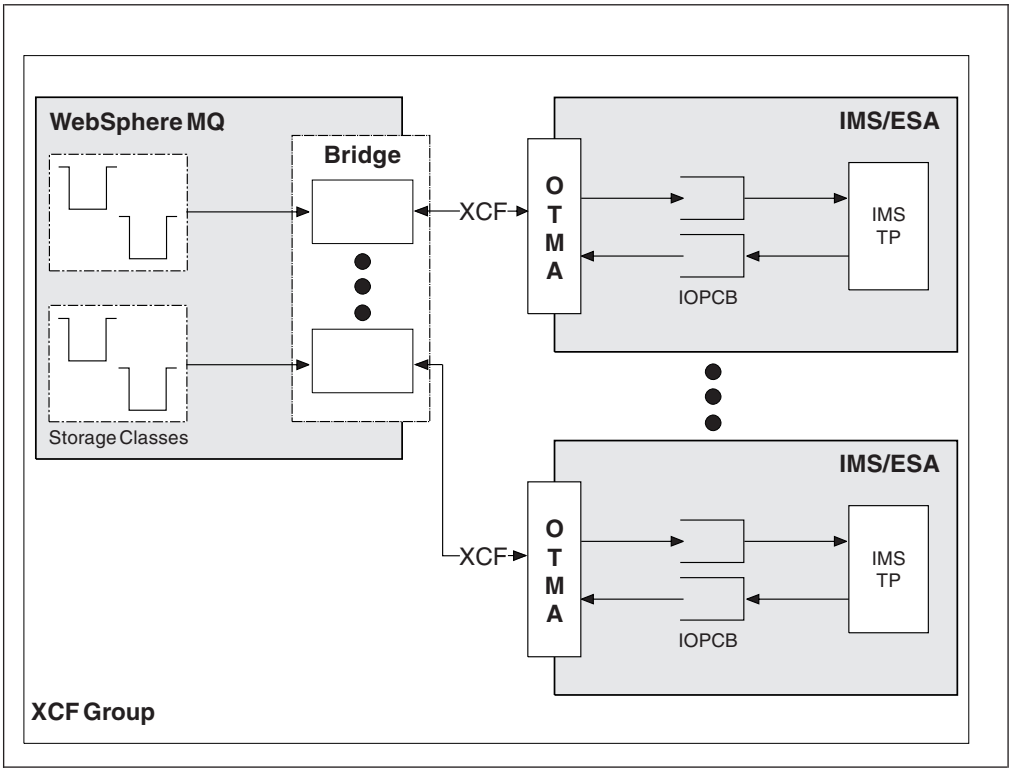

Figure 18. The WebSphere MQ-IMS bridge

A queue manager can connect to one or more IMS systems, and more than one queue manager can connect to one IMS system. The only restriction is that they must all belong to the same XCF group and must all be in the same sysplex.

What is OTMA?

The IMS OTMA facility is a transaction-based connectionless client/server protocol that runs on IMS Version 5.1 or later. It functions as an interface for host-based communications servers accessing IMS TM applications through the z/OS *Cross Systems Coupling Facility* (XCF).

OTMA enables clients to connect to IMS in a high performance manner enabling the client to support interactions with IMS for a large network or large number of sessions. OTMA is implemented in a z/OS sysplex environment. Therefore, the domain of OTMA is restricted to the domain of XCF.

Submitting IMS transactions from WebSphere MQ

To submit an IMS transaction that uses the bridge, applications put messages on a WebSphere MQ queue as usual. The messages contain IMS transaction data; they can have an IMS header (the MQIIH structure) or allow the WebSphere MQ-IMS bridge to make assumptions about the data in the message.

WebSphere MQ then puts the message to an IMS queue (it is queued in WebSphere MQ first to enable the use of syncpoints to assure data integrity). The storage class of the WebSphere MQ queue determines whether the queue is an *OTMA queue* (that is, a queue used to transmit messages to the WebSphere MQ-IMS bridge) and the particular IMS partner to which the message data is sent.

Remote queue managers can also start IMS transactions by writing to these OTMA queues on WebSphere MQ for z/OS.

Data returned from the IMS system is written directly to the WebSphere MQ reply-to queue specified in the message descriptor structure (MQMD). (This might be a transmission queue to the queue manager specified in the *ReplyToQMgr* field of the MQMD.)

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 17. Where to find more information about using IMS with WebSphere MQ

Topic	Where to look
Setting up the IMS adapter Setting up the IMS bridge	*WebSphere MQ for z/OS System Setup Guide*
Operating the IMS adapter Operating the IMS bridge	*WebSphere MQ for z/OS System Administration Guide*
Console messages	*WebSphere MQ for z/OS Messages and Codes*
Writing IMS applications	*WebSphere MQ Application Programming Guide*
IMS Open Transaction Manager Access (OTMA)	*IMS/ESA Open Transaction Manager Access Guide*
MQIIH structure	*WebSphere MQ Application Programming Reference*

Chapter 13. WebSphere MQ and z/OS Batch and TSO

This chapter discusses how WebSphere MQ works with z/OS Batch and TSO. It contains the following sections:

Introduction to the Batch adapters

The Batch/TSO adapters are the interface between z/OS application programs running under JES, TSO, or UNIX® System Services and WebSphere MQ. They enable z/OS application programs to use the MQI.

The adapters provide access to WebSphere MQ resources for programs running in:

- Task (TCB) mode
- Problem or supervisor state
- Non-cross-memory mode
- Non-access register mode

Connections between application programs and WebSphere MQ are at the task level. The adapters provide a connection thread from an application task control block (TCB) to WebSphere MQ.

The Batch/TSO adapter supports a single-phase commit protocol for changes made to resources owned by WebSphere MQ. It does not support multi-phase commit protocols. The RRS adapter enables WebSphere MQ applications to participate in two-phase commit protocols with other RRS-enabled products, coordinated by z/OS Resource Recovery Services (RRS).

The adapters use the z/OS STIMERM service to schedule an asynchronous event every second. This event runs an interrupt request block (IRB) that does not involve any waiting by the batch application's task. This IRB checks to see if the WebSphere MQ termination ECB has been posted. If the termination ECB has been posted, the IRB posts any application ECBs that are waiting on an event in WebSphere MQ (for example, a signal or a wait).

The Batch/TSO adapter

The WebSphere MQ Batch/TSO adapter provides WebSphere MQ support for z/OS Batch and TSO applications. All application programs that run under z/OS Batch or TSO must have the API stub program CSQBSTUB link-edited with them. The stub provides the application with access to all MQI calls. You use single-phase commit and backout for applications by issuing the MQI calls **MQCMIT** and **MQBACK**.

The RRS adapter

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

Resource Recovery Services (RRS) is a subcomponent of z/OS that provides a system-wide service for coordinating two-phase commit across z/OS products. The WebSphere MQ Batch/TSO RRS adapter (the RRS adapter) provides WebSphere MQ support for z/OS Batch and TSO applications that want to use these services. The RRS adapter enables WebSphere MQ to become a full participant in RRS coordination. Applications can participate in two-phase commit processing with other products that support RRS (for example, DB2).

The RRS adapter provides two stubs; application programs that want to use RRS must be link-edited with one of these stubs.

CSQBRSTB
This stub allows you to use two-phase commit and backout for applications by using the RRS callable resource recovery services instead of the MQI calls **MQCMIT** and **MQBACK**.

You must also link-edit module ATRSCSS from library SYS1.CSSLIB with your application. If you use the MQI calls **MQCMIT** and **MQBACK**, you will receive return code `MQRC_ENVIRONMENT_ERROR`.

CSQBRRSI
This stub allows you to use MQI calls **MQCMIT** and **MQBACK**; WebSphere MQ actually implements these calls as the **SRRCMIT** and **SRRBACK** RRS calls.

For information about building application programs that use the RRS adapter, see the *WebSphere MQ Application Programming Guide*.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 18. Where to find more information about using z/OS Batch with WebSphere MQ

Topic	Where to look
Setting up the Batch adapters	*WebSphere MQ for z/OS System Setup Guide*
Writing applications to use the Batch adapter	*WebSphere MQ Application Programming Guide*
RRS callable resource recovery services	*MVS Programming: Callable Services for High Level Languages*

Chapter 14. WebSphere MQ and WebSphere Application Server

This chapter discusses the use of WebSphere MQ for z/OS by the WebSphere Application Server.

Applications written in the Java programming language running under WebSphere Application Server can use the Java Messaging Service (JMS) specification to perform messaging. Point-to-point messaging in this environment is provided by a WebSphere MQ queue manager. Two classes of configuration are available:

- A *reduced function* form of WebSphere MQ.

  Here, the queue manager providing the point-to-point messaging runs the WebSphere MQ for z/OS base function code provided with WebSphere Application Server. This environment is suited to simple messaging between JMS applications running in the WebSphere Application Server environment.
- *Full function* WebSphere MQ.

  Here, the queue manager providing the point-to-point messaging runs WebSphere MQ for z/OS full function code installed as a separate product from WebSphere Application Server. In this environment, JMS applications running in the WebSphere Application Server can participate fully in the functionality of a WebSphere MQ network. For example, they can make use of the IMS Bridge, or exchange messages with WebSphere MQ queue managers running on other platforms.

"Using the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server" on page 114 tells you what *reduced function* means in more detail by identifying those functions not available in the reduced function form of WebSphere MQ supplied with WebSphere Application Server

Connection between WebSphere Application Server and a queue manager

If the queue manager providing point-to-point messaging is a reduced function queue manager, connection must be through a TCP client/server channel. In JMS terms, this means choosing *client transport* for the queue connection factory object.

If the queue manager providing point-to-point messaging is a full function queue manager, you can choose either *client transport* or *bindings transport* for the queue connection factory object. If you choose bindings transport, the WebSphere Application Server and the queue manager must both exist on the same z/OS image. If you choose client transport, you must install the *Client Attachment* feature of WebSphere MQ for z/OS.

Both types of connection support transactional applications: the client transport by using XA protocols; the bindings transport by using a WebSphere Application Server stub, CSQBWSTB, which uses RRS services.

For more information about configuring queue connection factories, see *WebSphere MQ Using Java*.

Using WebSphere MQ functions from JMS applications

By default, JMS messages held on WebSphere MQ queues use an MQRFH2 header to hold some of the JMS message header information. Many legacy 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 WebSphere MQ Workflow applications. The section ″Mapping JMS to a native WebSphere MQ application″ in the chapter ″JMS Messages″ of *WebSphere MQ Using Java* covers these special considerations.

Using the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server

The WebSphere MQ for z/OS information tells you what reduced function means by identifying those functions not available in the reduced function form of WebSphere MQ supplied with WebSphere Application Server, both in general terms, and at the command level. The list below summarizes the restrictions in the reduced function form of WebSphere MQ:

Queue sharing groups

Queue sharing groups are not available:

- If you set the QSGDATA system parameter to anything other than blank, the queue manager does not start.
- CSQ5PQSG does not run.

Clusters

Clusters are not available. You cannot:

- Define SYSTEM.CLUSTER queues
- Define, alter, or delete CLUSSDR or CLUSRCVR channels
- Set queue CLUSTER or CLUSNL attributes to anything other than blank
- Set queue manager REPOS or REPOSNL attributes to anything other than blank

The following commands have no effect:

- RESET or REFRESH CLUSTER
- SUSPEND or RESUME QMGR CLUSTER or CLUSNL
- DISPLAY CLUSQMGR

The following functions do not run:

- Mover repository manager
- CLUSSDR or CLUSRCVR channels

Distributed queuing

The reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server does not support channels other than server-connection channels. This means that:

- MCA channels are not available:
  - You cannot define, alter, or delete SDR, SVR, RCVR, or RQSTR channels
  - SDR, SVR, RCVR, or RQSTR channels do not run
- Only SVRCONN channels are available; you cannot define or alter CLNTCONN channels

TCP types other than OESOCKET

If you set the TCPTYPE system parameter to anything other than OESOCKET, the mover does not start.

LU62 communications

You cannot use LU6.2 channels and listeners. Specifically, you must not set the LU62CHL system parameter to anything other than 0 or the mover does not start.

Events

Events are not available; you cannot:

- Define SYSTEM.ADMIN queues
- Enable queue manager event attributes

CICS connection

There is no support for the WebSphere MQ for z/OS CICS connection. As a result, an MQCONN or MQCONNX from a CICS transaction fails with an MQRC_CONNECTION ERROR.

IMS connection

There is no support for the WebSphere MQ for z/OS IMS connection. As a result, an MQCONN or MQCONNX from an IMS application fails with an MQRC_CONNECTION ERROR.

RRS connection

There is no support for the WebSphere MQ for z/OS RRS connection. As a result, an MQCONN or MQCONNX from an RRS application fails with an MQRC_CONNECTION ERROR.

IMS Bridge

There is no support for the WebSphere MQ for z/OS IMS Bridge; the resource manager does not start.

Measured usage license charge (MULC)

This function is bypassed.

Part 4. Planning your WebSphere MQ environment

Chapter 15. Planning your storage requirements

This chapter discusses the storage requirements for WebSphere MQ for z/OS. It contains the following sections:

Address space storage

Each WebSphere MQ for z/OS subsystem has the following approximate storage requirements:

CSA 4 KB

ECSA 800 KB, plus the size of the trace table specified in the TRACTBL parameter of the CSQ6SYSP system parameter macro. This macro is described in the *WebSphere MQ for z/OS System Setup Guide*.

In addition, each concurrent WebSphere MQ task requires about 5 KB of ECSA. When a task ends, this storage can be reused by other WebSphere MQ tasks. WebSphere MQ does not release the storage until the queue manager is shut down, so the maximum amount of ECSA required can be calculated by multiplying the maximum number of concurrent tasks by 5 KB.

Concurrent tasks consist of the following:

- The number of Batch, TSO or IMS regions that have connected to WebSphere MQ, but not disconnected
- The number of CICS transactions that have issued a WebSphere MQ request, but have not terminated

The channel initiator typically requires ECSA usage of up to 160 KB, plus about 2.4 KB for each channel. The channel initiator does not use any CSA.

Private region storage usage

Every channel uses approximately 170 KB of extended private region in the channel initiator address space. Storage is increased by message size if messages larger than 32 KB are transmitted.

This increased storage is freed when:

- A sending or client channel requires less than half the current buffer size for 10 consecutive sends
- A heartbeat is sent or received

The storage is freed for re-use within the LE environment, but is not seen as free by the z/OS virtual storage manager.

This means that the upper limit for the number of channels is dependent on message size and arrival patterns as well as individual user system limitations on extended private region size. The upper limit on the number of channels is likely to be approximately 9000 on many systems as the extended region size is unlikely to exceed 1.6 GB. The use of message sizes larger than 32 KB reduces the

maximum number of channels in the system. For example, if messages that are 100 MB long are transmitted, and an extended region size of 1.6 GB is assumed, the maximum number of channels is 15.

Region sizes

The following table shows suggested values for region sizes. Two sets of values are given; one set is suitable for a test system, the other for a production system or a system that will become a production system eventually.

Table 19. Suggested definitions for JCL region sizes

Definition setting	Test system	Production system
Queue manager	REGION=7168K	REGION=0M
Channel initiator	REGION=7168K	REGION=0M

The region sizes suggested for the test system (REGION=7168K for both the queue manager region and channel initiator region) allow the queue manager and channel initiator to allocate 7 MB of private virtual storage below the 16 MB line (PVT) and up to 32 MB of extended private virtual storage above the 16 MB line (extended PVT).

These values are insufficient for a production system with large buffer pools and many active channels. The production region sizes chosen (REGION=0M) allow all available private storage above and below the 16 MB line to be used, although WebSphere MQ uses very little storage below the 16 MB line.

Note: You can use the z/OS exit IEALIMIT to override the region limits below the 16 MB line and IEFUSI to override the region limits above and below the 16 MB line.

Data storage

See the following sections for details of how to plan your data storage:

- **Logs and archive storage**

  "Logs and archive storage" on page 139 describes how to determine how much storage your active log and archive data sets require, depending on the volume of messages that your WebSphere MQ system handles and how often the active logs are off-loaded to your archive data sets.

- **DB2 storage**

  "DB2 storage" on page 135 describes how to determine how much storage DB2 will require for the WebSphere MQ data.

- **Coupling Facility storage**

  "Planning the size of your structures" on page 132 describes how to determine how large to make your Coupling Facility structures.

- **Page set and message storage**

  Chapter 16, "Planning your page sets and buffer pools", on page 123 describes how to determine how much storage your page data sets require, depending on the sizes of the messages that your applications will exchange, on the numbers of these messages, and on the rate at which they are created or exchanged.

Library storage

You need to allocate storage for the product libraries. The exact figures depend on your configuration, but an estimate of the space required by the distribution libraries is 80 MB. The target libraries require about 72 MB. Additionally, you require space for the SMP/E libraries.

You should refer to the program directory supplied with WebSphere MQ for z/OS for information about the required libraries and their sizes.

Two of the libraries, thlqual.SCSQMVR1 and thlqual.SCSQMVR2, must be in PDS-E format.

System LX usage

Each defined WebSphere MQ subsystem reserves one system linkage index (LX) at IPL time, and a number of non-system linkage indexes when the queue manager is started. The system linkage index is reused when the queue manager is stopped and restarted. Similarly, distributed queueing reserves one non-system linkage index. In the unlikely event of your z/OS system having inadequate system LXs defined, you might need to take these reserved system LXs into account.

Where to find more information

You can find more information about the topics discussed in this chapter from the following sources:

Table 20. Where to find more information about storage requirements

Topic	Were to look
System parameters	*WebSphere MQ for z/OS System Setup Guide*
Storage required to install WebSphere MQ	*WebSphere MQ for z/OS Program Directory*
IEALIMIT and IEFUSI exits	*MVS Installation Exits* , available from the zSeries Web site http://www.ibm.com/servers/eserver/zseries/zos/bkserv/
Latest information	WebSphere MQ SupportPac™ Web site http://www.ibm.com/software/mqseries

Chapter 16. Planning your page sets and buffer pools

This chapter contains the following sections:

Planning your page sets

When deciding on the most appropriate settings for page set definitions, there are a number of factors that should be considered.

Page set usage

In the case of short-lived messages, few pages are normally used on the page set and there is little or no I/O to the data sets except at startup, during a checkpoint, or at shutdown.

In the case of long-lived messages, those pages containing messages are normally written out to disk. This is performed by the queue manager in order to reduce restart time.

You should separate short-lived messages from long-lived messages by placing them on different page sets and in different buffer pools.

Number of page sets

Using several large page sets can make the role of the WebSphere MQ administrator easier because it means that you need fewer page sets, making the mapping of queues to page sets simpler.

Using multiple, smaller page sets has a number of advantages. For example, they take less time to back up, and I/O can be carried out in parallel during backup and restart. However, consider that this adds a significant overhead to the role of the WebSphere MQ administrator, who will be required to map each queue to one of a much greater number of page sets.

You are recommended to define at least five page sets, as follows:

- A page set reserved for object definitions (page set zero)
- A page set for system related messages
- A page set for performance-critical long-lived messages
- A page set for performance-critical short-lived messages
- A page set for all other messages

"Defining your buffer pools" on page 128 explains the performance advantages of distributing your messages on page sets in this way.

Size of page sets

You should allow enough space in your page sets for the expected peak message capacity. You should also specify a secondary extent to allow for any unexpected peak capacity, such as when a build up of messages develops because a queue server program is not running.

The size of the page set also determines the time taken to recover a page set when restoring from a backup, because a large page set takes longer to restore.

Note: Recovery of a page set also depends on the time the queue manager takes to process the log records written since the backup was taken; this is determined by the backup frequency. This is discussed in "Recovery procedures" on page 143.

Calculating the size of your page sets

For queue manager object definitions (for example, queues and processes), it is simple to calculate the storage requirement because these objects are of fixed size and are permanent. For messages however, the calculation is more complex for the following reasons:

- Messages vary in size.
- Messages are transitory.
- Space occupied by messages that have been retrieved is reclaimed periodically by an asynchronous process.

Page set zero

For page set zero, the storage required is:

```
   (maximum number of local queue definitions x 1010)
      (excluding shared queues)
+  (maximum number of model queue definitions x 746)
+  (maximum number of alias queue definitions x 338)
+  (maximum number of remote queue definitions x 434)
+  (maximum number of permanent dynamic queue definitions x 1010)
+  (maximum number of process definitions x 674)
+  (maximum number of namelist definitions x 12320)
+  (maximum number of message channel definitions x 2026)
+  (maximum number of client-connection channel definitions x 5170)
+  (maximum number of server-connection channel definitions x 2026)
+  (maximum number of storage class definitions x 266)
+  (maximum number of authentication information definitions x 1010)
```

Divide this value by 4096 to determine the number of records to specify in the cluster for the page set data set.

You do not need to allow for objects that are stored in the shared repository, but you do have to allow for objects that are stored or copied to page set zero (objects with a disposition of GROUP or QMGR).

The total number of objects that can be created is limited by the capacity of page set zero. There is an implementation limit on the number of local queues that can be defined, which is 524 287.

Page sets 01 to 99

For page sets 01 to 99, the storage required for each page set is determined by the number and size of the messages stored on that page set. (Messages on shared queues are not stored on page sets.)

Divide this value by 4096 to determine the number of records to specify in the cluster for the page set data set.

Calculating the storage requirement for messages

This section describes how messages are stored on pages. Understanding this will help you calculate how much page set storage you need to define for your

messages. To calculate the approximate space required for all messages on a page set you must consider maximum queue depth of all the queues that map to the page set and the average size of messages on those queues.

You must allow for the possibility that message "gets" might be delayed for reasons outside the control of WebSphere MQ (for example, because of a problem with your communications protocol). In this case, the "put" rate of messages might far exceed the "get" rate. This could lead to a large increase in the number of messages stored in the page sets and a consequent increase in the storage size demanded.

Each page in the page set is 4096 bytes long. Allowing for fixed header information, each page has 4057 bytes of space available for storing messages.

When calculating the space required for each message, the first thing you need to consider is whether the message will fit on one page (a *short message*) or whether it needs to be split over two or more pages (a *long message*). When messages are split in this way, you need to allow for additional control information in your space calculations.

For the purposes of space calculation, a message can be represented like this:

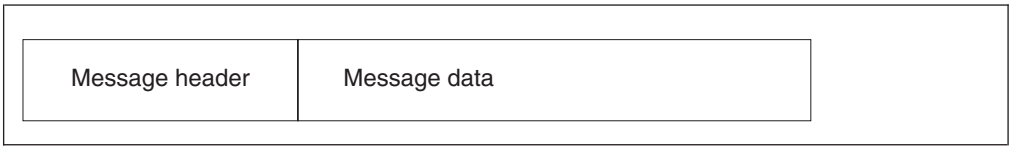

The message header section contains the message descriptor (352 bytes) and other control information, the size of which varies depending on the size of the message. The message data section contains all the actual message data, and any other headers (for example, the transmission header or the IMS bridge header).

A minimum of two pages are required for page set control information. This is typically less than 1% of the total space required for messages.

Short messages: A short message is defined as a message that will fit on one page.

For a short message the control information is 20 bytes long. When this is added to the length of the message header, the usable space remaining on the page is 3685 bytes. If the size of the message data is 3685 bytes or less, WebSphere MQ stores the messages in the next available space on the page, or if there is not enough space available, on the next page, as shown in Figure 19:

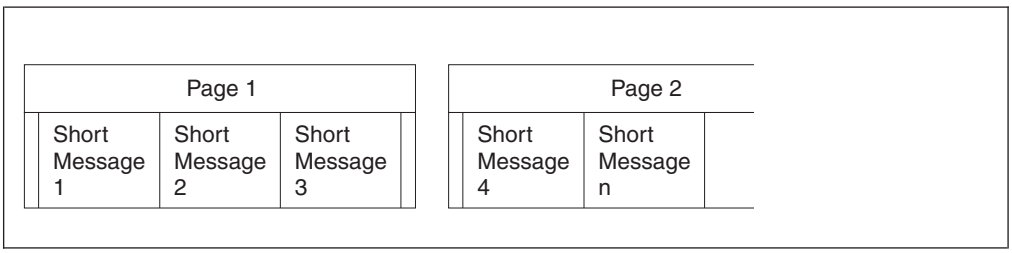

Figure 19. How WebSphere MQ stores short messages on page sets

If there is sufficient space remaining on the page, the next message is also stored on this page, if not, the remaining space on the page is left unused.

Long messages: If the size of the message data is greater than 3685 bytes, but not greater than 4 MB, the message is classed as a long message. When presented with a long message, WebSphere MQ stores the message on a series of pages, and stores control information that points to these pages in the same way that it would store a short message. This is shown in Figure 20:

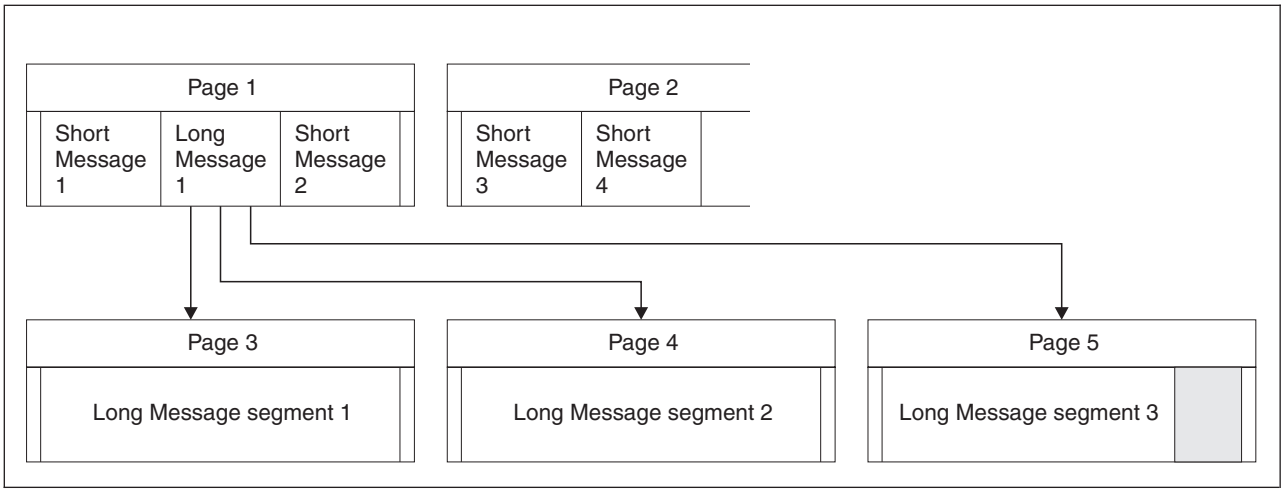

Figure 20. How WebSphere MQ stores long messages on page sets

Each segment of the long message is preceded by 8 bytes of control information, and the first segment also includes the message header portion of 352 bytes. This means that the first page contains 3697 bytes of the message data. The remaining message data is placed on subsequent pages, in 4049-byte segments. If this does not fill an exact number of pages, the remaining space in the last page is left unused.

The number of pages (n) used for a long message is calculated as follows:

```
     message data length + 352
n = -------------------------
                4049

rounded up to the nearest page
```

In addition to this, you need to allow space for the control information that points to the pages. The length of this (c) depends on the length of the message, and is calculated as follows:

```
c = 20 + (3n) bytes

(where n is the number of pages calculated above)
```

This means that the total page set space required for a long message is:

```
(n * 4096) + c bytes
```

Very long messages: Very long messages are messages with a size greater than 4 MB. These are stored so that each 4 MB uses 1037 pages. Any remainder is stored in the same way as a long message, as described above.

Enabling dynamic page set expansion

Page sets can be extended dynamically while the queue manager is running. A page set can have up to 123 extents, which can exist on multiple disk volumes.

Note: The maximum number of extents for a page set cataloged in an ICF catalog is between 119 and 123, depending upon the number of extents (1-5) allocated by direct access storage data management (DADSM) for each allocate or extend request.

In order to use this facility, your page sets must be allocated with secondary extent values defined. If you have existing page set definitions, they cannot be altered to add secondary extent definitions. You have to re-allocate each of your page sets with secondary extents, and then use the COPYPAGE function of CSQUTIL to copy the old versions of the page sets to the new ones.

Sample thlqual.SCSQPROC(CSQ4PAGE) (CSQ4PAGR when using reduced function WebSphere MQ) shows how to define the secondary extents. CSQUTIL is described in the *WebSphere MQ for z/OS System Administration Guide*.

How to determine an appropriate secondary extent value

You might decide the secondary extent value by considering how many times the page set should exceed its original value. For instance, if your page set primary allocation is 1000 units (records/pages, tracks, cylinders, KB, MB), and you want it to grow to be at most four times that size, then determine the secondary extent size from:

```
             ( maximum size   -  original size )
             -----------------------------------
                            119

In this case, 4000 - 1000  = 3000  = 25 or 26
              -----------    ----
                  119         119
```

You might not be able to use this much disk space, as described in "Number of extents available".

Note: If you define the size of your extent in records, IDCAMS rounds this up to map onto a physical boundary. The queue manager uses all this space for the secondary extent.

Number of extents available

The Data Facility Product (DFP) uses up to five non-contiguous areas of disk to satisfy the total space requirements of a primary or secondary extent. This means,

in the worst case of badly fragmented disk space, that you might only get around 22 times the secondary space allocated before you reach the maximum extent limit.

Multivolume data sets

If the definition of a page set allows it to utilize multiple volumes, the primary space is wholly contained on the first volume, and secondary extents are first allocated on the same volume, while space is available, and thereafter the next volume is used, while space is available, and so on. The process stops when you have used all the secondary extents, or no more disk space is available.

This behavior differs if the page set is a data set managed by Storage Management Subsystem (SMS), and you use a storage class that uses the GUARANTEED SPACE attribute. In this case, a primary extent is allocated on each volume when the page set is defined. Thereafter, secondary extents are allocated, as before, except that when the services of a new volume are required, the pre-allocated secondary extent is used.

Defining your buffer pools

You can use up to 16 buffer pools. However, you are recommended to use just the four buffer pools described in Table 21, except in the following circumstances:

- a particular queue is known to require isolation, perhaps because it exhibits significantly different behavior at various times.
  - such a queue might either require the best performance possible under the varying circumstances, or need to be isolated so as not to adversely impact the other queues in a buffer pool.
  - each such queue can be isolated into its own buffer pool. However, buffer pool 3 (as described below) might be the appropriate place.
- you want to isolate different sets of queues from each other for class of service reasons.
  - each set of queues might then require any or all of the three types of buffer pools 1, 2 and 3, as described in Table 21.

The following table shows suggested values for buffer pool definitions that affect the performance of queue manager operation, recovery, and restart. Two sets of values are given; one set is suitable for a test system, the other for a production system or a system that will become a production system eventually.

Table 21. Suggested definitions for buffer pool settings

Definition setting	Test system	Production system
BUFFPOOL 0	1 050 buffers	50 000 buffers
BUFFPOOL 1	1 050 buffers	20 000 buffers
BUFFPOOL 2	1 050 buffers	50 000 buffers
BUFFPOOL 3	1 050 buffers	20 000 buffers

You are recommended to reserve buffer pool zero for object definitions (in page set zero) and performance critical, system related message queues, such as the SYSTEM.CHANNEL.SYNCQ queue and the SYSTEM.CLUSTER.* queues. The remaining three buffer pools can be used for user messages, for example:

- Buffer pool 1 might be used for important long-lived messages.

  Long-lived messages are those that remain in the system for longer than two checkpoints, at which time they are written out to the page set. If you have a large number of long-lived messages, this buffer pool should be relatively small,

so that page set I/O is evenly distributed (older messages are written out to DASD each time the buffer pool becomes 85% full).

If the buffer pool is too large, page set I/O is delayed until checkpoint processing. This might affect response times throughout the system.

If you expect a small number of long-lived messages only, this buffer pool should be defined so that it is sufficiently large to hold all these messages.

- Buffer pool 2 might be used for performance-critical, short-lived messages.

  There is normally a high degree of buffer reuse, using a small number of buffers; however, you are recommended to make this buffer pool large to allow for unexpected message accumulation, for example, when a server application fails.
- Buffer pool 3 might be used for all other (usually performance non-critical) messages. Queues such as the dead-letter queue, SYSTEM.COMMAND.* queues and SYSTEM.ADMIN.* queues can also be mapped to buffer pool 3.

  Where virtual storage constraints exist and buffer pools need to be smaller, buffer pool 3 is the first candidate for size reduction.

Initially, you are recommended to define all buffer pools as shown in the table. Usage of buffer pools can be monitored by analysis of buffer pool performance statistics. In particular, you should ensure that the buffer pools are large enough so that the values of QPSTSOS, QPSTSTLA and QPSTNBUF remain at zero. (These performance statistics are described in the *WebSphere MQ for z/OS System Setup Guide*.)

Buffer pool zero and the buffer pool for short-lived messages (buffer pool 2) should be tuned so that the 15% free threshold is never exceeded (that is, QPSTCBSL divided by QPSTNBUF is always greater than 15%). If more than 15% of buffers remain free, I/O to the page sets using these buffer pools can be largely avoided during normal operation, although messages older than two checkpoints are written to page sets.

Note: The optimum value for these parameters is dependent on the characteristics of the individual system. The values given are only intended as a guideline and might not be appropriate for your system.

MQSeries SupportPac *Capacity planning and tuning for MQSeries for OS/390* (MP16) gives more information about tuning buffer pools.

Chapter 17. Planning your Coupling Facility and DB2 environment

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

This chapter discusses the following topics:

Defining Coupling Facility resources

If you intend to use shared queues, you must define the Coupling Facility structures that WebSphere MQ will use in your CFRM policy. To do this you must first update your CFRM policy with information about the structures, and then activate the policy.

Your installation probably has an existing CFRM policy that describes the Coupling Facilities available. The IXCMIAPU z/OS utility is used to modify the contents of the policy based on textual statements you provide. The utility is described in the *MVS Setting up a Sysplex* manual. You must add statements to the policy that define the names of the new structures, the Coupling Facilities that they are to be defined in, and what size the structures will be.

Planning your structures

A queue-sharing group requires a minimum of two structures to be defined. The first structure, known as the administrative structure, is used to coordinate WebSphere MQ internal activity across the queue-sharing group. No user data is held in this structure. It has a fixed name of *qsg-name*CSQ_ADMIN (where *qsg-name* is the name of your queue-sharing group). Subsequent structures are used to hold the messages on WebSphere MQ shared queues. Each structure can hold up to 512 shared queues.

Using multiple structures

A queue-sharing group can connect to up to 64 Coupling Facility structures. One of these structures must be the administration structure, but you can use up to 63 structures for WebSphere MQ data. You might choose to use multiple structures because:

- You have some queues that are likely to hold a very large number of messages and so will require all the resources of an entire Coupling Facility.
- You have a requirement for a very large number of shared queues, so they must be split across multiple structures because each structure can only contain 512 queues.
- RMF™ reports on the usage characteristic of a structure suggest that you should distribute the queues it contains across a number of Coupling Facilities.
- You want some queue data to held in a physically different Coupling Facility from other queue data for data isolation reasons.
- Recovery of persistent shared messages is performed using structure level attributes and commands, for example BACKUP CFSTRUCT. To simplify backup and recovery, you could assign queues that hold nonpersistent messages to different structures from those that hold persistent messages.

Planning your Coupling Facility and DB2 environment

When choosing which Coupling Facilities the structures should be allocated in, you should consider the following points:

- Your data isolation requirements.
- The volatility of the Coupling Facility (that is, its ability to preserve data through a power outage).
- Failure independence between the accessing systems and the Coupling Facility, or between Coupling Facilities.
- The level of Coupling Facility Control Code (CFCC) installed on the Coupling Facility (WebSphere MQ requires Level 9 or higher).

Planning the size of your structures

The administrative structure (*qsg-name*CSQ_ADMIN) must be large enough to contain 1000 list entries for each queue manager in the queue-sharing group. When a queue manager starts, the structure is checked to see if it is large enough for the number of queue managers currently *defined* to the queue-sharing group. Queue managers are considered as being defined to the queue-sharing group if they have been added by the CSQ5PQSG utility. You can check which queue managers are defined to the group with the MQSC DISPLAY GROUP command.

Table 22 shows the minimum required size for the administrative structure for various numbers of queue managers defined in the queue-sharing group. These sizes were established for a CFCC level 9 Coupling Facility structure; for higher levels of CFCC, they probably need to be larger.

Table 22. Minimum administrative structure sizes

Number of queue managers defined in queue-sharing group	Required storage
1	2560KB
2	3328KB
4	4608KB
8	7680KB
16	13568KB
31	24832KB

Note: Earlier versions of MQSeries for OS/390 and WebSphere MQ for z/OS check that the absolute size of the administrative structure is at least 10MB. You are strongly recommended to use a minimum administrative structure size of 10MB, even for small queue-sharing groups.

The size of the structures required to hold WebSphere MQ messages depends on the likely number and size of the messages to be held on a structure concurrently, together with an estimate of the likely number of concurrent units of work.

The graph in Figure 21 on page 133 shows how large you should make your CF structures to hold the messages on your shared queues. To calculate the allocation size you need to know

- The average size of messages on your queues
- The total number of messages likely to be stored in the structure

Find the number of messages along the horizontal axis. (Ticks are at multiples of 2, 5, and 8.) Select the curve that corresponds to your message size and determine the

required value from the vertical axis. For example, for 200 000 messages of length 1 KB gives a value between 256 and 512MB.

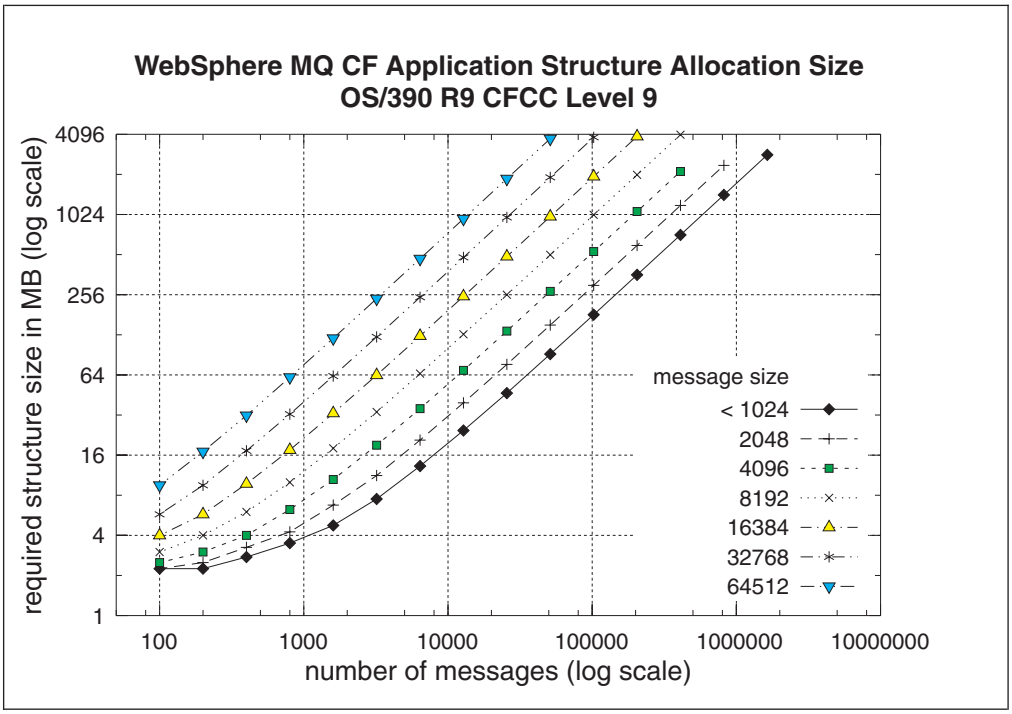

Figure 21. Calculating the size of a Coupling Facility structure

Your CFRM policy should include the following statements:

```
CFNAME(structure-name)
INITSIZE(value from graph in KB, that is, multiplied by 1024)
MAXSIZE(something larger)
FULLTHRESHOLD(85)
```

`INITSIZE` is the size in KB that XES allocates to the structure when the first connector connects to it. `MAXSIZE` is the maximum size that the structure can attain. `FULLTHRESHOLD` sets the percentage value of the threshold at which XES issues message IXC585E to indicate that the structure is getting full.

For example, with the figures determined above, you might include the following statements:

```
CFNAME(QSG1APPLICATION1)
INITSIZE(272144) /* 256 MB */
MAXSIZE(524288)  /* 512 MB */
FULLTHRESHOLD(85)
```

If the structure utilization reaches the threshold where warning messages are issued, intervention is required. You might use WebSphere MQ to inhibit **MQPUT** operations to some of the queues in the structure to prevent applications from

writing more messages, start more applications to get messages from the queues, or quiesce some of the applications that are putting messages to the queue.

Alternatively XES facilities can be used to alter the structure size in place. The following z/OS command:

```
SETXCF START,ALTER,STRNAME=structure-name,SIZE=newsize
```

alters the size of the structure to *newsize*, where *newsize* is a value that is less than the value of MAXSIZE specified on the CFRM policy for the structure, but greater than the current Coupling Facility size.

You can monitor the utilization of a Coupling Facility structure with the MQSC DISPLAY GROUP command.

If no action is taken and a queue structure fills up, an MQRC_STORAGE_MEDIUM_FULL return code is returned to the application. If the administration structure becomes full, the exact symptoms depend on which processes experience the error, but they might include the following problems:

- No responses to commands.
- Queue manager failure as a result of problems during commit processing.

Mapping shared queues to structures

The CFSTRUCT attribute of the queue definition is used to map the queue to a structure.

WebSphere MQ adds the name of the queue-sharing group to the beginning of the CFSTRUCT attribute. For a structure defined in the CFRM policy with name *qsg-name*SHAREDQ01, the definition of a queue that uses this structure will be:

```
DEFINE QLOCAL(myqueue) QSGDISP(SHARED) CFSTRUCT(SHAREDQ01)
```

Planning your DB2 environment

If you are using queue-sharing groups, WebSphere MQ needs to attach to a DB2 subsystem that is a member of a data-sharing group. WebSphere MQ needs to know the name of the data-sharing group that it is to connect to, and the name of a DB2 subsystem (or DB2 group) to connect to, to reach this data-sharing group. These names are specified in the QSGDATA parameter of the CSQ6SYSP system parameter macro (described in the *WebSphere MQ for z/OS System Setup Guide*).

WebSphere MQ uses the RRS Attach facility of DB2. This means that you can specify the name of a DB2 group that you want to connect to. The advantage of connecting to a DB2 group attach name (rather than a specific DB2 subsystem), is that WebSphere MQ can connect (or reconnect) to any available DB2 subsystem on the z/OS image that is a member of that group. There must be a DB2 subsystem that is a member of the data-sharing group active on each z/OS image where you are going to run a queue-sharing WebSphere MQ subsystem, and RRS must be active.

DB2 storage

You need to set up a set of DB2 tables that are used to hold WebSphere MQ data. These are automatically defined for you by WebSphere MQ when you set up your DB2 environment (as described in the *WebSphere MQ for z/OS System Setup Guide*).

For most installations, the amount of DB2 storage required is about 20 or 30 cylinders on a 3390 device. However, if you want to calculate your storage requirement, the following table gives some information to help you determine how much storage DB2 will require for the WebSphere MQ data. The table describes the length of each DB2 row, and when each row is added to or deleted from the relevant DB2 table. Use this information together with the information about calculating the space requirements for the DB2 tables and their indexes in the *DB2 for OS/390 Installation Guide*.

Table 23. Planning your DB2 storage requirements

DB2 table name	Length of row	A row is added when:	A row is deleted when:
CSQ.ADMIN_B_QSG	240 bytes	A queue-sharing group is added to the table with the ADD QSG function of the CSQ5PQSG utility.	A queue-sharing group is removed from the table with the REMOVE QSG function of the CSQ5PQSG utility. (All rows relating to this queue-sharing group are deleted automatically from all the other DB2 tables when the queue-sharing group record is deleted.)
CSQ.ADMIN_B_QMGR	Up to 3752 bytes	A queue manager is added to the table with the ADD QMGR function of the CSQ5PQSG utility.	A queue manager is removed from the table with the REMOVE QMGR function of the CSQ5PQSG utility.
CSQ.ADMIN_B_STRUCTURE	329 bytes	The first local queue definition, specifying the QSGDISP(SHARED) attribute, that names a previously unknown structure within the queue-sharing group is defined.	The last local queue definition, specifying the QSGDISP(SHARED) attribute, that names a structure within the queue-sharing group is deleted.
CSQ.ADMIN_B_SCST	342 bytes	A shared channel is started.	A shared channel becomes inactive.
CSQ.ADMIN_B_SSKT	254 bytes	A shared channel that has the NPMSPEED(NORMAL) attribute is started.	A shared channel that has the NPMSPEED(NORMAL) attribute becomes inactive.
CSQ.ADMIN_B_STRBACKUP	507 bytes	A new row is added to the CSQ.ADMIN_B_STRUCTURE table. Each entry is a dummy entry until the BACKUP CFSTRUCT command is run, which overwrites the dummy entries.	A row is deleted from the CSQ.ADMIN_B_STRUCTURE table.
CSQ.OBJ_B_AUTHINFO	3400 bytes	An authentication information object with QSGDISP(GROUP) is defined.	An authentication information object with QSGDISP(GROUP) is deleted.

Table 23. Planning your DB2 storage requirements (continued)

DB2 table name	Length of row	A row is added when:	A row is deleted when:
CSQ.OBJ_B_QUEUE	Up to 3620 bytes	• A queue with the QSGDISP(GROUP) attribute is defined. • A queue with the QSGDISP(SHARED) attribute is defined. • A model queue with the DEFTYPE(SHAREDYN) attribute is opened.	• A queue with the QSGDISP(GROUP) attribute is deleted. • A queue with the QSGDISP(SHARED) attribute is deleted. • A dynamic queue with the DEFTYPE(SHAREDYN) attribute is closed with the DELETE option.
CSQ.OBJ_B_NAMELIST	Up to 15127 bytes	A namelist with the QSGDISP(GROUP) attribute is defined.	A namelist with the QSGDISP(GROUP) attribute is deleted.
CSQ.OBJ_B_CHANNEL	Up to 14027 bytes	A channel with the QSGDISP(GROUP) attribute is defined.	A channel with the QSGDISP(GROUP) attribute is deleted.
CSQ.OBJ_B_STGCLASS	Up to 2856 bytes	A storage class with the QSGDISP(GROUP) attribute is defined.	A storage class with the QSGDISP(GROUP) attribute class is deleted.
CSQ.OBJ_B_PROCESS	Up to 3347 bytes	A process with the QSGDISP(GROUP) attribute is defined.	A process with the QSGDISP(GROUP) attribute is deleted.

Chapter 18. Planning your logging environment

The WebSphere MQ logging environment is established using the system parameter macros to specify options, such as whether to have single or dual active logs, what media to use for the archive log volumes, and how many log buffers to have. These macros are described in the *WebSphere MQ for z/OS System Setup Guide*.

This chapter discusses the following topics:

Planning your logs

You are recommended to have at least three log data sets, and to use dual logging and log archiving. The following table shows suggested values for log and bootstrap data set definitions that affect the performance of queue manager operation, recovery, and restart. Two sets of values are given; one set is suitable for a test system, the other for a production system or a system that will become a production system.

Table 24. Suggested definitions for log and bootstrap data sets

Definition setting	Test system	Production system
Log data set size[1]	10 000 records (40 MB approx.) Access Method Services rounds to nearest cylinder and allocates 10 080 records.	180 000 records (700 MB or 1 000 cylinders of 3390)
Number of active logs[1]	3	6
BSDS size	60 records (240 KB approx.)	120 records (480 KB approx.) A primary extent of this size should be sufficient for the maximum number of archive logs that can be recorded in the BSDS (1 000).
Note: 1. The optimum value for this definition is dependent on the characteristics of the individual system. The values given are intended as a guideline only and might not be appropriate for your system.		

Note: If you are using queue-sharing groups, ensure that you define the bootstrap and log data sets with SHAREOPTIONS(2 3).

Log data set definitions

Before setting up the log data sets, review the following section to decide on the most appropriate configuration for your system.

Should your installation use single or dual logging?

If your DASD type is 3390 or similar, you are recommended to use dual logging to ensure that you have an alternative backup source in the event of losing a data set. You should also use dual BSDSs and dual archiving to ensure adequate provision for data recovery.

If you use devices with in-built data redundancy (for example, Redundant Array of Independent Disks (RAID) devices) you might consider using single active logging. If you use persistent messages, single logging can increase maximum capacity by 10-30% and can also improve response times.

Dual active logging adds a small performance overhead. You might want to specify it on a test system used for benchmarking, in addition to a production system.

How many active log data sets do you need?

You should have sufficient active logs to ensure that your system is not impacted in the event of an archive being delayed.

In practice, you should have at least three active log data sets but it is preferable to define more. For example, if the time taken to fill a log is likely to approach the time taken to archive a log during peak load, you should define more logs. You are also recommended to define more logs to offset possible delays in log archiving. If you use archive logs on tape, allow for the time required to mount the tape.

How large should the active logs be?

Your logs should be large enough so that it takes at least 30 minutes to fill a single log during the expected peak persistent message load. If you are archiving to tape, you are advised to make the logs large enough to fill one tape cartridge, or a number of tape cartridges. (For example, a log size of 1000 cylinders on 3390 DASD will fit onto a 3490E non-compacted tape with space to spare.)

Note: When archiving to tape, a copy of the BSDS is also written to the tape. When archiving to DASD, a separate data set is created for the BSDS copy.

If the logs are small (for example, 10 cylinders) it is likely that they will fill up frequently, which could result in performance degradation. In addition, you might find that the large number of archive logs required is difficult to manage.

If the logs are very large, and you are archiving to DASD, you need a corresponding amount of spare space reserved on DASD for SMS retrieval of migrated archive logs, which might cause space management problems. In addition, the time taken to restart might increase because one or more of the logs has to be read sequentially at startup.

Active log placement

Ideally, each active log should be allocated on separate, low-usage DASD volumes. As a minimum, no two adjacent logs should be on the same volume.

When an active log fills, the next log in the ring is used and the previous log data set is copied to the archive data set. If these two active data sets are on the same volume, contention will result, because one data set is read while the other is written to. For example, if you have three active logs and use dual logging, you need six DASD volumes because each log is adjacent to both of the two other logs. Alternatively, if you have four active logs and you want to conserve DASD space, by allocating logs 1 and 3 on one volume and logs 2 and 4 on another, you need four DASD volumes only.

In addition, you should ensure that primary and secondary logs are on separate physical units. If you use 3390 DASD, be aware that each head disk assembly contains two logical volumes. The physical layout of other DASD subsystems such as RAMAC® arrays should also be taken into account.

Logs and archive storage

Active log data sets record significant events and data changes. They are periodically off-loaded to the archive log. Consequently, the space requirements for your active log data sets depend on the volume of messages that your WebSphere MQ handles and how often the active logs are off-loaded to your archive data sets. WebSphere MQ provides optional support for dual logging; if you use this your log storage requirement will be doubled.

If you decide to place the archive data sets on direct access storage devices (DASD), you need to reserve enough space on the devices. Space should also be reserved for the bootstrap data sets (BSDS). A typical size for each BSDS might be 500 KB. These are all separate data sets and should be allocated space on different volumes and strings if possible to minimize DASD contention and problems caused by any defects on the physical devices.

Because each change to the system is logged, the size of storage required can be estimated from the size and expected throughput of persistent messages (nonpersistent messages are not logged). You must add to this a small overhead for the header information in the data sets.

Additionally, CF structure backups are written to the active log of the queue manager where the BACKUP CFSTRUCT command is issued.

To arrive at the size of the log extents, an algorithm can be developed that depends on various factors including the message rate and size of persistent messages and how frequently you want to switch the log, and CF structure backup characteristics.

Figure 22 shows an approximate calculation for the number of records to specify in the cluster for the log data set.

```
Number of records = ((a * log switch interval) + S) / 4096

where a = (Number of puts/sec*(average persistent message size+500))
         + (Number of gets/sec*110)
         + (Number of units of recovery started*120)
         + (Number of syncpoints a second*240)
 and
       log switch interval = time period between successive log
                             switches required in seconds
 and
       S = the number of CF structure backups in the log switch
       interval multiplied by the average structure size
```

Figure 22. Calculating the number of records to specify in the cluster for the log data set

Each log data set should have the same number of records specified and should not have secondary extents. Other than for a very small number of records, AMS rounds up the number of records so that a whole number of cylinders is allocated. The number of records actually allocated is:

```
c = (INT (number of log records / b) + 1) * b

Where b is the number of 4096-byte blocks in each cylinder (180 for a
3390 device) and INT means round down to an integer
```

Planning your archive storage

This section describes the different ways of maintaining your archive log data sets.

Archive log data sets can be placed on standard-label tapes, or DASD, and can be managed by data facility hierarchical storage manager (DFHSM). Each z/OS logical record in an archive log data set is a VSAM control interval from the active log data set. The block size is a multiple of 4 KB.

Archive log data sets are dynamically allocated, with names chosen by WebSphere MQ. The data set name prefix, block size, unit name, and DASD sizes needed for such allocations are specified in the system parameter module. You can also choose, at installation time, to have WebSphere MQ add a date and time to the archive log data set name.

It is not possible to choose specific volumes for new archive logs. If allocation errors occur, off-loading is postponed until the next time off-loading is triggered.

If you specify dual archive logs at installation time, each log control interval retrieved from the active log is written to two archive log data sets. The log records that are contained in the pair of archive log data sets are identical, but the end-of-volume points are not synchronized for multi-volume data sets.

Should your archive logs reside on tape or DASD?

When deciding whether to use tape or DASD for your archive logs, there are a number of factors that you should consider:

- Review your operating procedures before making decisions about tape or disk. For example, if you choose to archive to tape, operators must be available to mount the appropriate tapes when they are required.
- During recovery, archive logs on tape are available as soon as the tape is mounted. If DASD archives have been used, and the data sets migrated to tape using hierarchical storage manager (HSM), there is a delay while HSM recalls each data set to disk. You can recall the data sets before the archive log is used. However, it is not always possible to predict the correct order in which they are required.
- When using archive logs on DASD, if many logs are required (which might be the case when recovering a page set after restoring from a backup) you might require a significant quantity of DASD to hold all the archive logs.
- In a low usage system or test system, it might be more convenient to have archive logs on DASD to eliminate the need for tape mounts.
- Both issuing a RECOVER CFSTRUCT command and backing out a persistent unit of work result in the log being read backwards. Tape drives with hardware compression perform very badly on operations that read backwards. You should plan sufficient log data on DASD to avoid reading backwards from tape.

Archiving to DASD offers faster recoverability but is more expensive than archiving to tape. If you use dual logging, you can specify that the primary copy of the archive log go to DASD and the secondary copy go to tape. This increases recovery speed without using as much DASD, and the tape can be used as a backup.

Archiving to tape

If you choose to archive to a tape device, WebSphere MQ can extend to a maximum of twenty volumes.

If you choose to off-load to tape, you should consider adjusting the size of your active log data sets so that each nearly fills a tape volume. This minimizes tape handling and volume mounts, and maximizes the use of tape resources. However, such an adjustment is not essential.

If you are considering changing the size of the active log data set so that the set fits on one tape volume, you must bear in mind that a copy of the BSDS is placed on the same tape volume as the copy of the active log data set. Adjust the size of the active log data set downward to offset the space required for the BSDS on the tape volume.

If you use dual archive logs on tape, it is typical for one copy to be held locally, and the other copy to be held off-site for use in disaster recovery.

Archiving to DASD volumes

WebSphere MQ requires that all archive log data sets allocated on non-tape devices (DASD) be cataloged. If you choose to archive to DASD, the CATALOG parameter of the CSQ6ARVP macro must be YES. (This macro is described in the *WebSphere MQ for z/OS System Setup Guide*.) If this parameter is NO, and you decide to place archive log data sets on DASD, you receive message CSQJ072E each time an archive log data set is allocated, although WebSphere MQ still catalogs the data set.

If the archive log data set is held on DASD, the archive log data sets cannot extend to another volume.

If you choose to use DASD, make sure that the primary space allocation (both quantity and block size) is large enough to contain either the data coming from the active log data set, or that from the corresponding BSDS, whichever is the larger of the two. This minimizes the possibility of unwanted z/OS X'B37' or X'E37' abends during the off-load process. The primary space allocation is set with the PRIQTY (primary quantity) parameter of the CSQ6ARVP macro.

Using SMS with archive log data sets

If you have MVS/DFP storage management subsystem (DFSMS) installed, you can write an Automatic Class Selection (ACS) user-exit filter for your archive log data sets, which helps you convert them for the SMS environment. Such a filter, for example, can route your output to a DASD data set, which in turn can be managed by DFSMS. You must exercise caution if you use an ACS filter in this manner. Because SMS requires DASD data sets to be cataloged, you must make sure the CATALOG DATA field of the CSQ6ARVP macro contains YES. If it does not, message CSQJ072E is returned; however, the data set is still cataloged by WebSphere MQ.

For more information about ACS filters, see the *DFP Storage Administration Reference* manual, and the *SMS Migration Planning Guide*.

Chapter 19. Planning for backup and recovery

Developing backup and recovery procedures at your site is vital to avoid costly and time-consuming losses of data. WebSphere MQ provides means for recovering both queues and messages to their current state after a system failure.

This chapter contains the following sections:

Recovery procedures

You should develop the following procedures for WebSphere MQ:
- Creating a point of recovery
- Backing up page sets
- Backing up CF structures
- Recovering page sets
- Recovering from out-of-space conditions (WebSphere MQ logs and page sets)
- Recovering CF structures

See the *WebSphere MQ for z/OS System Administration Guide* for information about these.

You should also be familiar with the procedures used at your site for the following:
- Recovering from a hardware or power failure
- Recovering from a z/OS component failure
- Recovering from a site interruption, using off-site recovery

General tips for backup and recovery

This section introduces some backup and recovery tasks. The queue manager restart process recovers your data to a consistent state by applying log information to the page sets. If your page sets are damaged or unavailable, you can resolve the problem using your backup copies of your page sets (provided that all the logs are available). If your log data sets are damaged or unavailable, it might not be possible to recover completely.

Periodically take backup copies

A *point of recovery* is the term used to describe a set of backup copies of WebSphere MQ page sets and the corresponding log data sets required to recover these page sets. These backup copies provide a potential restart point in the event of page set loss (for example, page set I/O error). If you restart the queue manager using these backup copies, the data in WebSphere MQ will be consistent up to the point that

these copies were taken. Providing that all logs are available from this point, WebSphere MQ can be recovered to the point of failure.

The more recent your backup copies, the quicker WebSphere MQ can recover the data in the page sets. The recovery of the page sets are dependent on all the necessary log data sets being available.

In planning for recovery, you need to determine how often to take backup copies and how many complete backup cycles to keep. These values tell you how long you must keep your log data sets and backup copies of page sets for WebSphere MQ recovery.

In deciding how often to take backup copies, consider the time needed to recover a page set. It is determined by:
- The amount of log to traverse
- The time it takes an operator to mount and remove archive tape volumes
- The time it takes to read the part of the log needed for recovery
- The time needed to reprocess changed pages
- The storage medium used for the backup copies
- The method used to make and restore backup copies

In general, the more frequently you make backup copies, the less time recovery takes, but the more time is spent making copies.

For each queue manager, you should take backup copies of:
- The archive log data sets
- The BSDS copies created at the time of the archive
- The page sets
- Your object definitions
- Your CF structures

To reduce the risk of your backup copies being lost of damaged, you should consider:
- Storing the backup copies on different storage volumes to the original copies.
- Storing the backup copies at a different site to the original copies.
- Making at least two copies of each backup of your page sets and, if you are using single logging or a single BSDS, two copies of your archive logs and BSDS. If you are using dual logging or BSDS, a single copy of both archive logs or BSDS will suffice.

Before moving WebSphere MQ to a production environment you should have tested and documented your backup procedures.

Backing up your object definitions

You should also create backup copies of your object definitions. To do this, use the MAKEDEF feature of the COMMAND function of the utility program (described in the *WebSphere MQ for z/OS System Administration Guide*).

You should do this whenever you take backup copies of your queue manager data sets, and keep the most current version.

Backing up your Coupling Facility structures

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

If you are using queue-sharing groups, you should take periodic backups of your CF structures. To do this, use the WebSphere MQ BACKUP CFSTRUCT command (described in the *WebSphere MQ Script (MQSC) Command Reference*). This command can only be used on CF structures that are defined with the RECOVER(YES) attribute.

It is recommended that you take a backup of all your CF structures about every hour, to minimize the time it takes to restore a CF structure.

You could perform all your CF structure backups on a single queue manager, which has the advantage of limiting the increased log utilization to a single queue manager. Alternatively, you could perform backups on all the queue managers in the queue-sharing group, which has the advantage of spreading the workload across the queue-sharing group. Whichever strategy you use, WebSphere MQ can locate the backup and perform a RECOVER CFSTRUCT from any queue manager in the queue-sharing group. The logs of all the queue managers in the queue-sharing group need to be accessed to recover the CF structure.

Do not discard archive logs you might need

WebSphere MQ might need to use archive logs during restart. You must keep sufficient archive logs so the system can be fully restored. WebSphere MQ might use an archive log to recover a page set from a restored backup copy. If you have discarded that archive log, WebSphere MQ cannot restore the page set to its current state. When and how you should discard archive logs is described in the *WebSphere MQ for z/OS System Administration Guide*.

Do not change the DDname to page set association

WebSphere MQ associates page set number 00 with DDname CSQP0000, page set number 01 with DDname CSQP0001, and so on up to CSQP0099. WebSphere MQ writes recovery log records for a page set based on the DDname that the page set is associated with. For this reason, you must not move page sets that have already been associated with a PSID DDname.

Recovering page sets

A key factor in recovery strategy concerns the period of time for which you can tolerate a queue manager outage. The total outage time might include the time taken to recover a page set from a backup, or to restart the queue manager after an abnormal termination. Factors affecting restart time include how frequently you back up your page sets, and how much data is written to the log between checkpoints.

To minimize the restart time after an abnormal termination, keep units of work short so that, at most, two active logs are used when the system restarts. For example, if you are designing a WebSphere MQ application, avoid placing an **MQGET** call that has a long wait interval between the first in-syncpoint MQI call and the commit point because this might result in a unit of work that has a long duration. Other common causes of long units of work are batch intervals of more than 5 minutes for the channel initiator, and in-doubt channels in the CICS mover.

You can use the DISPLAY THREAD command to display the RBA of units of work and to help resolve the old ones. For information about the DISPLAY THREAD command, see the *WebSphere MQ Script (MQSC) Command Reference* manual.

How often should a page set be backed up?

Frequent page set backup is essential if a reasonably short recovery time is required. This applies even when a page set is very small or there is a small amount of activity on queues in that page set.

If you use persistent messages in a page set, the backup frequency should be in the order of hours rather than days. This is also the case for page set zero.

To calculate an approximate backup frequency, start by determining the target total recovery time. This consists of:

1. The time taken to react to the problem.
2. The time taken to restore the page set backup copy.

   (For example, you can restore approximately 60 cylinders of 3390 data a minute from and to RAMAC Virtual Array 2 Turbo 82 (RVA2-T82) DASD using Access Method Services REPRO.)

   If you use SnapShot backup/restore, the time taken to perform this task is of the order of a few seconds. For information about SnapShot, see the *DFSMSdss™ Storage Administration Guide*.
3. The time the queue manager requires to restart, including the additional time needed to recover the page set.

   This depends most significantly on the amount of log data that must be read from active and archive logs since that page set was last backed up. All such log data must be read, in addition to that directly associated with the damaged page set.

   Note: When using *fuzzy backup* (where a snapshot is taken of the logs and page sets while a unit of work is active), it might be necessary to read up to three additional checkpoints, and this might result in the need to read one or more additional logs.

When deciding on how long to allow for the recovery of the page set, the factors that you need to consider are:

- The rate at which data is written to the active logs during normal processing:
  - Approximately 1.3 KB of extra data is required on the log for a persistent message.
  - Approximately 2.5 KB of data is required on the log for each batch of messages sent on a channel.
  - Approximately 1.4 KB of data is required on the log for each batch of messages received on a channel.
  - Nonpersistent messages require no log data. NPMSPEED(FAST) channels require no log data for batches consisting entirely of nonpersistent messages.

  The rate at which data is written to the log depends on how messages arrive in your system, in addition to the message rate. Messages received or sent over a channel result in more data logging than messages generated and retrieved locally.
- The rate at which data can be read from the archive and active logs.

  When reading the logs, the achievable data rate depends on the devices used and the overall load on your particular DASD subsystem. (For example, data rates of approximately 2.7 MB a second have been observed using active and archive logs on RVA2-T82 DASD.)

With most tape units, it is possible to achieve higher data rates for archived logs with a large block size. However, if an archive log is required for recovery, all the data on the active logs must be read also.

Recovering CF structures

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

At least one queue manager in the queue-sharing group must be active to process a RECOVER CFSTRUCT command. CF structure recovery does not impact queue manager restart time, as recovery is performed by an already active queue manager.

The recovery process consists of two logical steps that are managed by the RECOVER CFSTRUCT command:

1. locating and restoring the backup
2. merging all the logged updates to persistent messages that are held on the CF structure from the logs of all the queue managers in the queue-sharing group that have used the CF structure, and applying the changes to the backup.

The second step is likely to take much longer as a great quantity of log data might need to be read. The time taken can be reduced if you take frequent backups, or if you recover multiple CF structures at the same time, or both.

The queue manager performing the recovery locates the relevant backups on all the other queue managers' logs using the data in DB2 and the bootstrap data sets. The queue manager replays these backups in the correct time sequence across the queue sharing group, from just before the last backup through to the point of failure.

The time it takes to recover a CF structure depends on the amount of recovery log data that must be replayed, which in turn depends on the frequency of the backups. In the worst case, it takes as long to read a queue manager's log as it did to write it. So if, for example, you have a queue-sharing group containing six queue managers, an hour's worth of log activity could take six hours to replay. In general it takes less time than this, because reading can be done in bulk, and because the different queue manager's logs can be read in parallel. As a starting point, we recommend that you backup your CF structures every hour.

All queue managers can continue working with non-shared queues and queues in other CF structures while there is a failed CF structure. However, if the administration structure has also failed, all the queue managers in the queue-sharing group must be started before the RECOVER CFSTRUCT can be issued.

Backing up CF structures can require considerable log writing capacity, and can therefore impose a large load on the queue manager doing the backup. You should choose a lightly loaded queue manager for doing backups; for very busy systems, it might be sensible to add an additional queue manager to the queue-sharing group and dedicate it exclusively for doing backups.

Achieving specific recovery targets

If you have specific recovery targets to achieve, for example, completion of the queue manager recovery and restart processing in addition to the normal startup time within *xx* seconds, you can use the following calculation to estimate your backup frequency (in hours):

```
Formula (A)
                    Required restart time * System recovery log read rate
                          (in secs)                  (in MB/sec)
Backup frequency  = ------------------------------------------------------
   (in hours)              Application log write rate (in MB/hour)
```

Note: The examples given below are intended to highlight the need to back up your page sets on a frequent basis. The calculations assume that the majority of log activity is derived from a large number of persistent messages. However, there are situations where the amount of log activity is not easily calculated. For example, in a queue-sharing group environment, a unit of work in which shared queues are updated in addition to other resources might result in UOW records being written to the WebSphere MQ log. For this reason, the 'Application log write rate' in Formula (A) can only be derived accurately from the observed rate at which the WebSphere MQ logs fill.

For example, consider a system in which WebSphere MQ clients generate an overall load of 100 persistent messages a second. In this case, all messages are generated locally.

If each message is of user length 1 KB, the amount of data logged each hour is of the order:

```
100 * (1 + 1.3) KB * 3600 = approximately 800 MB

where
      100           = the message rate a second
      (1 + 1.3) KB  = the amount of data logged for
                      each 1 KB of persistent messages
```

Consider an overall target recovery time of 75 minutes. If you have allowed 15 minutes to react to the problem and restore the page set backup copy, queue manager recovery and restart must then complete within 60 minutes (3600 seconds) applying formula (A). Assuming that all required log data is on RVA2-T82 DASD, which has a recovery rate of approximately 2.7 MB a second, this necessitates a page set backup frequency of at least every:

```
  3600 seconds * 2.7 MB a second / 800 MB an hour  =  12.15 hours
```

If your WebSphere MQ application day lasts approximately 12 hours, one backup each day is appropriate. However, if the application day lasts 24 hours, two backups each day is more appropriate.

Another example might be a production system in which all the messages are for request-reply applications (that is, a persistent message is received on a receiver channel and a persistent reply message is generated and sent down a sender channel).

In this example, the achieved batch size is one, and so there is one batch for every message. If there are 50 request replies a second, the overall load is 100 persistent messages a second. If each message is 1 KB in length, the amount of data logged each hour is of the order:

```
50((2 * (1+1.3) KB) + 1.4 KB + 2.5 KB) * 3600 = approximately 1500 MB

where:
 50                 = the message pair rate a second
 (2 * (1 + 1.3) KB) = the amount of data logged for each message pair
 1.4 KB             = the overhead for each batch of messages
                      received by each channel
 2.5 KB             = the overhead for each batch of messages sent
                      by each channel
```

To achieve the queue manager recovery and restart within 30 minutes (1800 seconds), again assuming that all required log data is on RVA2-T82 DASD, this requires that page set backup is carried out at least every:

```
    1800 seconds * 2.7 MB a second / 1500 MB an hour = 3.24 hours
```

Periodic review of backup frequency

You are recommended to monitor your WebSphere MQ log usage in terms of MB an hour. You should perform this check periodically and amend your page set backup frequency if necessary.

Backup and recovery with DFHSM

The data facility hierarchical storage manager (DFHSM) does automatic space- and data-availability management among storage devices in your system. If you use it, you need to know that it moves data to and from the WebSphere MQ storage automatically.

DFHSM manages your DASD space efficiently by moving data sets that have not been used recently to alternate storage. It also makes your data available for recovery by automatically copying new or changed data sets to tape or DASD backup volumes. It can delete data sets, or move them to another device. Its operations occur daily, at a specified time, and allow for keeping a data set for a predetermined period before deleting or moving it.

All DFHSM operations can also be performed manually. The *Data Facility Hierarchical Storage Manager User's Guide* explains how to use the DFHSM commands. If you use DFHSM with WebSphere MQ, note that DFHSM:

- Uses cataloged data sets
- Operates on page sets and logs
- Supports VSAM data sets

WebSphere MQ recovery and CICS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

The recovery of CICS resources is not affected by the presence of WebSphere MQ. CICS recognizes WebSphere MQ as a non-CICS resource (or external resource manager), and includes WebSphere MQ as a participant in any syncpoint coordination requests using the CICS resource manager interface (RMI). For more information about CICS recovery, see the *CICS Recovery and Restart Guide*. For information about the CICS resource manager interface, see the *CICS Customization Guide*.

WebSphere MQ recovery and IMS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

IMS recognizes WebSphere MQ as an external subsystem and as a participant in syncpoint coordination. IMS recovery for external subsystem resources is described in the *IMS Customization Guide*.

Preparing for recovery on an alternative site

In the case of a total loss of a WebSphere MQ computing center, you can recover on another WebSphere MQ system at a recovery site. To do this, you must regularly back up the page sets and the logs. As with all data recovery operations, the objectives of disaster recovery are to lose as little data, workload processing (updates), and time as possible.

At the recovery site:

- The recovery WebSphere MQ queue manager **must** have the same name as the lost queue manager.
- The system parameter module used on the recovery queue manager should contain the same parameters as the lost queue manager.

The process for disaster recovery is described in the *WebSphere MQ for z/OS System Administration Guide*.

Example of queue manager backup activity

When you plan your queue manager backup strategy, a key consideration is retention of the correct amount of log data. The *WebSphere MQ for z/OS System Setup Guide* describes how to determine which log data sets are required, by reference to the system recovery RBA of the queue manager. WebSphere MQ determines the system recovery RBA using information about the following:

- currently active units of work
- page set updates that have not yet been flushed from the buffer pools to disk
- CF structure backups, and whether this queue manager's log contains information required in any recovery operation using them

You must retain sufficient log data to be able to perform media recovery. Whilst the system recovery RBA increases over time, the amount of log data that must be retained only decreases when subsequent backups are taken. CF structure backups are managed by WebSphere MQ, and so are taken into account when reporting the system recovery RBA. This means that in practice, the requirement for log data retention only reduces when page set backups are taken.

Figure 23 shows an example of the backup activity on a queue manager that is a member of a queue-sharing group, how the recovery RBA varies with each backup, and how that affects the amount of log data that must be retained. In the example the queue manager uses local and shared resources: page sets, and two CF structures, STRUCTURE1 and STRUCTURE2.

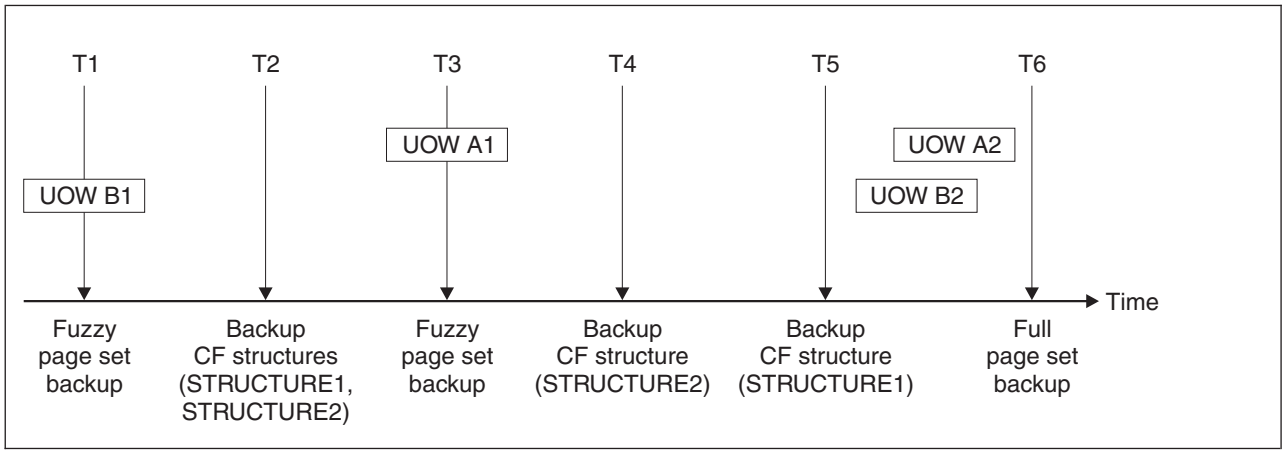

Figure 23. Example of queue manager backup activity

This is what happens at each point in time:

Point in time T1

A fuzzy backup is created of your page sets, as described in *WebSphere MQ for z/OS System Administration Guide*.

The system recovery RBA of the queue manager is the lowest of:

- the recovery RBAs of the page sets being backed up at this point
- the lowest recovery RBA required to recover the CF application structures. This relates to the recovery of backups of STRUCTURE1 and STRUCTURE2 created earlier.
- the recovery RBA for the oldest currently active unit of work within the queue manager (UOWB1)

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the fuzzy backup process.

Point in time T2
Backups of the CF structures are created. CF structure STRUCTURE1 is backed up first, followed by STRUCTURE2.

Log data retention is unchanged, because the same data as determined from the system recovery RBA at T1 is still required to recover using the page set backups taken at T1.

Point in time T3
Another fuzzy backup is created.

The system recovery RBA of the queue manager is the lowest of:

- the recovery RBAs of the page sets being backed up at this point
- the lowest recovery RBA required to recover CF structure STRUCTURE1, because STRUCTURE1 was backed up before STRUCTURE2
- the recovery RBA for the oldest currently active unit of work within the queue manager (UOWA1)

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the fuzzy backup process.

You can now reduce the log data retained, as determined by this new system recovery RBA.

Point in time T4
A backup is taken of CF structure STRUCTURE2. The recovery RBA for the recovery of the oldest required CF structure backup relates to the backup of CF structure STRUCTURE1, which was backed up at time T2.

The creation of this CF structure backup has no impact on log data retention.

Point in time T5
A backup is taken of CF structure STRUCTURE1. The recovery RBA for recovery of the oldest required CF structure backup now relates to recovery of CF structure STRUCTURE2, which was backed up at time T4.

The creation of this CF structure backup has no impact on log data retention.

Point in time T6
A full backup is taken of your page sets as described in *WebSphere MQ for z/OS System Administration Guide*.

The system recovery RBA of the queue manager is the lowest of:

- the recovery RBAs of the page sets being backed up at this point
- the lowest recovery RBA required to recover the CF structures. This relates to recovery of CF structure STRUCTURE2.
- the recovery RBA for the oldest currently active unit of work within the queue manager. In this case, there are no current units of work.

The system recovery RBA for this point in time is given by messages issued by the DISPLAY USAGE command, which is part of the full backup process.

Again, the log data retained can be reduced, because the system recovery RBA associated with the full backup is more recent.

Part 5. Planning to install WebSphere MQ

Chapter 20. WebSphere MQ Prerequisites

This chapter discusses the prerequisite products that you need to install before you can install and use full function WebSphere MQ for z/OS.

If you are using the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server, your prerequisites will be different from those stated here for the full function form of WebSphere MQ for z/OS Version 5 Release 3.1. See the WebSphere Application Server InfoCenter for more information.

Machine requirements

WebSphere MQ runs on any IBM zSeries or S/390® processor that is capable of running the required level of z/OS or OS/390, and that has enough storage to meet the combined requirements of the programming prerequisites, WebSphere MQ, the access methods, and the application programs.

Software requirements

This section lists the software requirements for WebSphere MQ. We recommend that you use one of the packaged offerings of z/OS from IBM (for example, ServerPac). This includes most of the products that you need to run WebSphere MQ for z/OS, and the integrated products have been verified by IBM.

The list of elements included in z/OS, and the recommended levels of nonexclusive elements are described in the *OS/390 Planning For Installation* manual.

WebSphere MQ for z/OS Version 5 Release 3.1 requires OS/390 Version 2.9 or later, together with the products included in z/OS. These include:

- C/C++
- optionally Cryptographic Services System SSL (if you want to use the Secure Sockets Layer (SSL) for channel security)
- optionally Cryptographic Services Security Level 3 (if you want to use SSL for channel security with US encryption strengths)
- DFSMS™/dfp
- High Level Assembler
- ICSS
- ISPF
- JES
- Language Environment® (previously known as LE/370)
- Security Server (previously known as RACF®)
- SMP/E
- optionally TCP/IP (see “Communications protocol and distributed queuing” on page 162)
- TSO/E
- UNIX Services (previously known as OpenEdition®)
- VTAM®

The following list gives the minimum levels required for the optional products that are not included in z/OS. You can also use any more recent versions of COBOL and PL/I compilers that can generate calls to WebSphere MQ conforming to the standard operating system linkage conventions. You might not need all these products.

Prerequisites

- CICS, Version 4.1
- COBOL:
  - IBM SAA® AD/Cycle® COBOL/370™
  - VS COBOL 2, Version 4
  - IBM COBOL for OS/390 (z/OS) and VM, Version 2.1
  - IBM COBOL for MVS and VM, Version 2.1
- IMS, Version 5.1
- IBM AD/Cycle PL/I for MVS and VM, Version 1.1

Additional requirements for some features

Some of the features of WebSphere MQ for z/OS Version 5 Release 3.1 have additional requirements. These are listed below. (These features are not available in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

Queue-sharing groups

- Coupling Facility, Level 9
- DB2, Version 6.1

CICS bridge (3270 transactions)

- CICS Transaction Server, Release 2 or later. Release 2 requires APAR PQ32659, Release 3 requires APAR PQ23961.

Internet Gateway

- Internet Connection Secure Server, Version 2.2
- Java, Version 1.1.1
- Web browser that supports HTML 3.2 or later

Java Support feature

IBM Developer Kit for OS/390, Java 2 Technology Edition, Version 1.3.1.

MQSeries Workflow (now called WebSphere Process Manager)

- MQSeries Workflow, Version 3.1

Non-IBM products

If you choose to use Unicenter TCPaccess Communication Server (formerly called SOLVE:TCPaccess) from Computer Associates instead of IBM's TCP/IP, the recommended minimum level is Version 5.2.

Clients

For WebSphere MQ for z/OS to support clients you need to install the client/server support code that is provided by the Client Attachment Feature of WebSphere MQ for z/OS.

Delivery

(This information does not apply to the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

WebSphere MQ for z/OS is supplied on 3480 cartridge only. One cartridge contains the product code together with three language features; U.S. English (mixed case), U.S. English (uppercase), Chinese, and Japanese, and the optional CICS mover and Internet Gateway features. The optional Client Attachment feature is supplied on a separate cartridge.

The Java Support feature, which is only required if you want to use Java and the Java Message Service, is supplied on a separate cartridge.

Chapter 21. Making WebSphere MQ available

This chapter gives an overview of what you need to do to make WebSphere MQ available to application programmers, and their applications. It contains the following sections:

- "Installing WebSphere MQ for z/OS"
- "Customizing WebSphere MQ and its adapters" on page 165
- "Verifying your installation of WebSphere MQ for z/OS" on page 166

If you are migrating from a previous version of WebSphere MQ for z/OS, see Chapter 22, "Migrating from previous versions", on page 167.

If you are using the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server, some of the steps needed to make WebSphere MQ for z/OS available might be different from those stated here for the full function form of WebSphere MQ for z/OS Version 5 Release 3.1. See the WebSphere Application Server InfoCenter for more information.

Installing WebSphere MQ for z/OS

WebSphere MQ for z/OS uses the standard z/OS installation procedure. It is supplied with a Program Directory that contains specific instructions for installing the program on a z/OS system. You must follow the instructions in the *WebSphere MQ for z/OS Program Directory.* They include not only details of the installation process, but also information about the prerequisite products and their service or maintenance levels.

SMP/E, used for installation on the z/OS platform, validates the service levels and prerequisite and corequisite products, and maintains the SMP/E history records to record the installation of WebSphere MQ for z/OS. It loads the WebSphere MQ for z/OS libraries and checks that the loads have been successful. You then have to customize the product to your own requirements.

Before you install and customize WebSphere MQ for z/OS, you must decide the following:

- Whether you are going to install one of the optional national language features.
- Which communications protocol and distributed queuing facility you are going to use.
- What your naming convention for WebSphere MQ objects will be.
- What command prefix string (CPF) you are going to use for each queue manager.

You also need to plan how much storage you require in your z/OS system to accommodate WebSphere MQ; Chapter 15, "Planning your storage requirements", on page 119 helps you plan the amount of storage required.

National language support

You can choose one of the following national languages for the WebSphere MQ operator messages and the WebSphere MQ operations and control panels (including the character sets used). Each language is identified by a language letter:

C Simplified Chinese
E U.S. English (mixed case)
K Japanese
U U.S. English (uppercase)

The samples, WebSphere MQ commands, and utility control statements are available only in mixed case U.S. English.

Communications protocol and distributed queuing

The distributed queuing facility provided with the base product feature of WebSphere MQ uses native z/OS communications (APPC or TCP/IP). It can either use APPC (LU 6.2), TCP/IP from IBM, or SOLVE:TCPaccess. The distributed queuing facility is also known as the channel initiator and the mover.

Alternatively, you can use CICS ISC for distributed queuing; this facility is also known as the CICS mover. You must install the CICS mover feature to use it. (This feature is retained for compatibility with previous releases but there will be no further enhancements to this function. Therefore, you are recommended to use the channel initiator for distributed queuing.)

You can enable both facilities and use them simultaneously on the same queue manager. However, the two types will have no knowledge of each other or each other's channels, and you must ensure that the channel names they use are distinct.

If you want to use clients, the client attachment feature is needed as well (but you can administer clients without it).

Whichever mover you choose, you must perform the following tasks to enable distributed queuing:

- Choose which communications interface to use. This can be:
  - APPC (LU 6.2)
  - IBM TCP/IP OpenEdition sockets
  - SOLVE:TCPaccess (native, IUCV, or OpenEdition sockets)
- Customize the distributed queuing facility and define the WebSphere MQ objects required.
- Define access security.
- Set up your communications. This includes setting up your TCPIP.DATA data set if you are using TCP/IP, LU names and side information if you are using APPC, and CICS definitions if you are using CICS. This is described in the *WebSphere MQ Intercommunication* manual.

Naming conventions

It is advisable to establish a set of naming conventions when planning your WebSphere MQ systems. The names you choose will probably be used on different platforms, so you should follow the convention for WebSphere MQ, not for the particular platform.

WebSphere MQ allows both uppercase and lowercase letters in names, and the names are case sensitive. However, some z/OS consoles fold names to uppercase, so do not use lowercase characters for names unless you are sure that this will not happen.

You can also use numeric characters and the period (.), forward slash (/), underscore (_) and percent (%) characters. The percent sign is a special character to RACF, so do not use it in names if you are using RACF as your External Security Manager. Do not use leading or trailing underscore characters if you are planning to use the Operations and Control panels.

Rules for naming WebSphere MQ objects are described in the *WebSphere MQ Script (MQSC) Command Reference* manual.

Choosing names for queue managers and queue-sharing groups

(Queue-sharing groups are not available in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

Each queue manager and queue-sharing group within a network must have a unique name. On z/OS the names of queue managers and queue-sharing groups can be up to four characters long. Each DB2 system and data-sharing group within the network must also have a unique name.

Queue manager and queue-sharing group names can use only uppercase alphabetic characters, numeric characters, and $, # or @; they must not start with a numeric character. For implementation reasons, queue-sharing group names less than four characters long are padded internally with @ symbols, so do not use names ending in @.

The queue manager name is the same as the z/OS subsystem name. You could identify each subsystem as a queue manager by giving it the name **QM***xx* (where *xx* is a distinguishing identifier), or you could choose a naming convention similar to **ADDX**, where **A** signifies the geographic area, **DD** signifies the company division, and **X** is a distinguishing identifier.

You might want to use your naming convention to distinguish between queue managers and queue-sharing groups. For example, you could identify each queue-sharing group by giving it the name **QG***xx* (where *xx* is the distinguishing identifier).

Choosing names for objects

Queues, processes, namelists, and clusters can have names up to 48 characters long. Channels can have names up to 20 characters long and storage classes can have names up to 8 characters long.

If possible, choose meaningful names within any constraints of your local conventions. Any structure or hierarchy within names is ignored by WebSphere MQ, however, hierarchical names can be useful for system management. You can also specify a description of the object when you define it to give more information about its purpose.

Each object must have a unique name within its object type. However, each object type has a separate name space, so you can define objects of different types with the same name. For example, if a queue has an associated process definition, it is a good idea to give the queue and the process the same name. It is also a good idea to give a transmission queue the same name as its destination queue manager.

You could also use the naming convention to identify whether the object definition is private or a global. For example, you could call a namelist **project_group.global** to indicate that the definition is stored on the shared repository.

Application queues: Choosing names that describe the function of each queue helps you to manage these queues more easily. For example, you could call a queue for inquiries about the company payroll **payroll_inquiry**. The reply-to queue for responses to the inquiries could be called **payroll_inquiry_reply**.

You can use a prefix to group related queues. This means that you can specify groups of queues for administration tasks like managing security and using the dead-letter queue handler. For example, all the queues that belong to the payroll application could be prefixed by **payroll_**. You can then define a single security profile to protect all queues with names beginning with this prefix.

You can also use your naming convention to indicate that a queue is a shared queue. For example, if the payroll inquiry queue mentioned above was a shared queue, you could call it **payroll_inquiry.shared**.

Storage classes and Coupling Facility structures

(This information does not apply to the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

The character set you can use when naming storage classes and Coupling Facility structures is limited to uppercase alphabetic and numeric characters. You should be systematic when choosing names for these objects.

Storage class names can be up to 8 characters long, and must begin with an alphabetic character. You will probably not define many storage classes, so a simple name is sufficient. For example, a storage class for IMS bridge queues could be called **IMS**.

Coupling Facility structure names can be up to 12 characters long, and must begin with an alphabetic character. You could use the name to indicate something about the shared queues associated with the Coupling Facility structure (that they all belong to one suite of applications for example). Remember that in the Coupling Facility itself, the structure names are the WebSphere MQ name prefixed by the queue-sharing group name (padded to four characters with @ symbols).

Choosing names for channels

To help you manage channels, it is a good idea if the channel name includes the names of the source and target queue managers. For example, a channel transmitting messages from a queue manager called QM27 to a queue manager called QM11 might be called **QM27/QM11**.

If your network supports both TCP and SNA, you might also want to include the transport type in the channel name, for example **QM27/QM11_TCP**. You could also indicate whether the channel is a shared channel, for example **QM27/QM11_TCP.shared**.

Remember that channel names cannot be longer than 20 characters. If you are communicating with a queue manager on a different platform, where the name of the queue manager might contain more than 4 characters, you might not be able to include the whole name in the name of the channel.

Using command prefix strings

Each instance of WebSphere MQ that you install must have its own *command prefix* string (CPF). You use the CPF to identify the z/OS subsystem that commands are intended for. It also identifies the z/OS subsystem from which messages sent to the console originate.

You can issue all MQSC commands from an authorized console by inserting the CPF before the command. If you enter commands through the system command input queue (for example, using CSQUTIL), or use the WebSphere MQ operations and control panels, you do not use the CPF.

To start a subsystem called CSQ1 whose CPF is '+CSQ1', issue the command +CSQ1 START QMGR from the operator console (the space between the CPF and the command is optional).

The CPF also identifies the subsystem that is returning operator messages. The following example shows +CSQ1 as the CPF between the message number and the message text.

```
CSQ9022I +CSQ1 CSQNCDSP ' DISPLAY CMDSERV' NORMAL COMPLETION
```

See the *WebSphere MQ for z/OS System Setup Guide* for information about defining command prefix strings.

Customizing WebSphere MQ and its adapters

WebSphere MQ requires some customization after installation to meet the individual and special requirements of your system, and to use your system resources in the most effective way. These are the tasks that you must perform when you customize your system:

1. Identify the z/OS system parameters
2. APF authorize the WebSphere MQ load libraries
3. Update the z/OS link list and LPA
4. Update the z/OS program properties table
5. Define the WebSphere MQ subsystem to z/OS
6. Create procedures for the WebSphere MQ subsystem
7. Create procedures for the channel initiator
8. Set up the DB2 environment
9. Set up the Coupling Facility
10. Implement your ESM security controls
11. Update SYS1.PARMLIB members
12. Customize the initialization input data sets
13. Create the bootstrap and log data sets
14. Define your page sets
15. Add the WebSphere MQ entries to the DB2 data-sharing group
16. Tailor your system parameter module
17. Tailor the channel initiator parameter module
18. Set up Batch, TSO, and RRS adapters
19. Set up the operations and control panels
20. Include the WebSphere MQ dump formatting member
21. Suppress information messages

These tasks are described in detail in the *WebSphere MQ for z/OS System Setup Guide*.

Using queue-sharing groups

(Queue-sharing groups are not available in the reduced function form of WebSphere MQ for z/OS supplied with WebSphere Application Server.)

Making WebSphere MQ available

If you want to use queue-sharing groups, you do not have to set them up when you install WebSphere MQ, you can do this at any time.

See the migration section in the *WebSphere MQ for z/OS System Setup Guide* for details of how to set up a new queue-sharing group. See the *WebSphere MQ for z/OS System Administration Guide* for details of how to manage your queue-sharing groups once you have set them up.

Verifying your installation of WebSphere MQ for z/OS

After the installation and customization has been completed, you can use the installation verification programs (IVPs) supplied with WebSphere MQ to verify that the installation has been completed successfully. The IVPs supplied are assembler language programs and should be run after WebSphere MQ has been customized to suit your needs. They are described in the *WebSphere MQ for z/OS System Setup Guide*.

Chapter 22. Migrating from previous versions

This chapter is for people who are planning to migrate from a previous version of WebSphere MQ for z/OS or MQSeries for OS/390. It contains the following sections:

What's new in WebSphere MQ for z/OS

This section describes the new function that has been added in WebSphere MQ for z/OS Version 5 Release 3 and WebSphere MQ for z/OS Version 5 Release 3.1. Most of the changes described were introduced in WebSphere MQ for z/OS Version 5 Release 3; by default, that is the release to which the information applies. Where changes have been made in WebSphere MQ for z/OS Version 5 Release 3.1, we identify them explicitly.

Support for WebSphere Application Server Version 5

WebSphere MQ for z/OS Version 5 Release 3.1 provides JMS support for WebSphere Application Server embedded messaging. You can use the embedded messaging either with the reduced function form of WebSphere MQ supplied with WebSphere Application Server, or with the full function WebSphere MQ for z/OS Version 5 Release 3.1 product itself.

In Chapter 14, "WebSphere MQ and WebSphere Application Server", on page 113, you'll find a general introduction to the reduced function form of WebSphere MQ, and an overview of the features of WebSphere MQ for z/OS that are not available to WebSphere Application Server users of that messaging.

Shared queue recovery

You can now use persistent messages on shared queues. Shared queue messages are recovered in a Coupling Facility (CF) list structure from periodic CF list structure backups, and from data about shared queue updates on the WebSphere MQ recovery logs of each queue manager in the queue sharing group.

This function introduces a new object called a *CF structure* (CFSTRUCT), and the following new MQSC commands:

DEFINE CFSTRUCT
: Creates a new CF structure.

ALTER CFSTRUCT
: Modifies the backup attributes of a CF structure.

DELETE CFSTRUCT
: Deletes a CF structure.

DISPLAY CFSTRUCT
: Displays the attributes for a specific CF structure.

BACKUP CFSTRUCT
: Starts a CF structure backup.

DISPLAY CFSTATUS
Displays the status of a CF structure, including the time and size of its latest backup, and usage information previously shown by the DISPLAY GROUP command.

RECOVER CFSTRUCT
Starts a CF structure recovery.

See the *WebSphere MQ Script (MQSC) Command Reference* for details of these commands.

Within a queue-sharing group, Version 5.2 queue managers can coexist with Version 5.3 or Version 5.3.1 queue managers, with some restrictions. Versions 5.3 and 5.3.1 queue managers can coexist without restriction. See the *WebSphere MQ for z/OS System Setup Guide* for more information.

Dynamic parameters

You can now dynamically change selected queue manager startup options while the queue manager is running. This function introduces commands like:

```
SET SYS IDFORE(256) CTHREAD(512) IDBACK(128)
```

to dynamically set selected parameters of the CSQ6SYSP, CSQ6LOGP and CSQ6ARVP startup parameter macros.

This function introduces the following new MQSC commands:

DISPLAY ARCHIVE
Displays archive information, including the tape unit report previously produced by the DISPLAY LOG command.

DISPLAY SYSTEM
Displays system information.

SET ARCHIVE
Sets archive parameter values.

SET SYSTEM
Sets certain system parameter values.

See the *WebSphere MQ Script (MQSC) Command Reference* for details of these commands.

Secure Sockets Layer (SSL) support

The Secure Sockets Layer (SSL) protocol provides out of the box channel security, with protection against eavesdropping, tampering, and impersonation.

SSL is an industry-standard protocol that provides a data security layer between application protocols and the communications layer, usually TCP/IP. The SSL protocol was designed by the Netscape Development Corporation, and is widely deployed in both Internet applications and intranet applications. SSL defines methods for data encryption, server authentication, message integrity, and client authentication for a TCP/IP connection.

SSL uses public key and symmetric techniques to provide the following security services:

Message privacy
SSL uses a combination of public-key and symmetric-key encryption to ensure message privacy. Before exchanging messages, an SSL server and an SSL client perform an electronic handshake during which they agree to use a session key and an encryption algorithm. All messages sent between the client and the server are then encrypted. Encryption ensures that the message remains private even if eavesdroppers intercept it.

Message integrity
SSL uses the combination of a shared secret key and message hash functions. This ensures that nothing changes the content of a message as it travels between client and server.

Mutual authentication
During the initial SSL handshake, the server uses a public-key certificate to convince the client of the server's identity. Optionally, the client might also exchange a public-key certificate with the server to ensure the authenticity of the client.

This function introduces a new object called an *authentication information object*, and the following new MQSC commands:

ALTER AUTHINFO
Modifies the attributes of an authentication information object.

DEFINE AUTHINFO
Creates a new authentication information object.

DELETE AUTHINFO
Deletes an authentication information object.

DISPLAY AUTHINFO
Displays the attributes for a specific authentication information object.

This function introduces the following new queue manager parameters:

SSLCRLNL
Allows access to a certificate revocation list. The SSLCRLNL attribute points to a namelist that contains AUTHINFO objects. Each AUTHINFO object points to an LDAP server.

SSLKEYR
Associates a key repository with a queue manager. The SSLKEYR attribute points to a RACF key ring. The key ring's default certificate is your queue manager certificate. The queue manager certificate key ring contains certificates of all certification authorities.

SSLTASKS
Sets the number of server subtasks to use for processing SSL calls.

This function introduces the following new channel parameters:

SSLCAUTH
Defines whether WebSphere MQ requires and validates a certificate from the SSL client.

SSLCIPH
Specifies the encryption strength and function (cipher specification), for example NULL_MD5 or RC4_MD5_US. The cipher specification must match at both ends of channel.

SSLPEER
Specifies the distinguished name (certificate identifier) filter for allowed partners.

As a result of the addition of SSL to WebSphere MQ, two of your load libraries, thlqual.SCSQMVR1 and thlqual.SCSQMVR2, must now be in PDS-E format.

See the *WebSphere MQ Script (MQSC) Command Reference* for details of the new commands and the new channel and queue manager attributes.

Configuration change notification

The configuration change notification function outputs details of new and changed WebSphere MQ objects to a configuration event queue to provide an audit trail or for use by system management programs.

This function introduces a new type of event, a configuration event, and a new queue, SYSTEM.ADMIN.CONFIG.EVENT. The event queue can be processed by user or vendor programs to track changes or record them in a central repository. One message is generated for each object that is created or deleted. Two messages are generated for each object that is changed.

A new queue manager attribute, CONFIGEV has been introduced, which lets you control these events.

A new command, REFRESH QMGR TYPE(CONFIGEV), generates an event for every current object.

Message grouping

Message grouping enables applications to retrieve messages within a group in order. Applications can process complete groups only, providing better performance. This enables different application instances to work concurrently on different sets of related messages

Message grouping lets you get messages in logical sequence, irrespective of their physical sequence. Each group message is physically stored by the queue manager with an extra message header, the MQMDE, which contains a 24-byte group ID and a 4-byte message sequence number. All messages in the same group have the same group ID, but different message sequence numbers. The last message in the group is indicated by a flag in the MQMDE. Message grouping uses this to get the messages in logical sequence, and to optionally wait for a complete group.

You can construct an application to put or get messages in groups, using additional put or get options as described in the *WebSphere MQ Application Programming Guide*. Although the message is physically stored by the queue manager with an extra message header, the MQMDE, the application can choose to provide or retrieve the information this header contains through additional fields in the MQMD version 2 structure, as described in the *WebSphere MQ Application Programming Reference*. The MQMD language declarations provided with Version 5.3 and later of WebSphere MQ have changed as follows:

- The PL/I and C language declarations contain the additional MQMD version 2 fields and a new MQMD2 structure initialized with version 2 values.
- You can configure the System/390 Assembler declaration through the DCLVER and VERSION parameters of the CMQMDA macro.

- The COBOL language declarations contained in CMQMDL and CMQMDV are for a version 1 MQMD, and additional copybooks CMQMD2L and CMQMD2V are provided that contain the MQMD version 2 structure.

For all languages, the default values for MQMD are those for a version 1 structure, and an additional, explicit MQMD1 structure is available (copybooks CMQMD1L and CMQMD1V in COBOL). This structure is provided for applications that have to deal with both version 1 and version 2 MQMDs.

When constructing an application that deals with group messages, ensure that, if the longer MQMD version 2 is passed to other programs, those other programs can cope with the additional structure length.

Queue manager improvements

You no longer need to IPL your system every time you install or upgrade a queue manager, even if you upgrade the early code. Also, a new configuration option, INDXTYPE, allows queue manager restart to be independent of the need to rebuild persistent message indexes. You can change the INDXTYPE of a private queue immediately, even if it contains messages, without waiting for a restart.

The DEFINE MAXSMSGS command has been superseded by the queue manager attribute MAXUMSGS.

Queue manager usability has improved in a number of ways:

- You can remove a page set, and reinstate it later.
- The process to reduce page set size has been simplified.
- The DISPLAY LOG command is extended to display the status of the logs.
- The tape unit report formerly produced by this command is now available with the new DISPLAY ARCHIVE command.
- The DISPLAY USAGE command is extended to display log data set information.
- Logging can be temporarily suspended using the SUSPEND QMGR command.

Queue manager performance has also improved, with the introduction of the EXPRYINT attribute, which enables more timely and efficient removal of expired messages, and the the introduction of the QINDXBLD attribute, which enables WebSphere MQ to build a queue index the first time the queue manager is used, instead of at queue manager restart.

Channels

Channel availability has improved by the introduction of the channel KeepAlive attribute. Channel KeepAlive is the time interval after which TCP/IP checks the health of an individual channel's connection. This means that WebSphere MQ applications can process messages more quickly after a TCP/IP communications failure. This attribute is applicable to all channel types except client connection channels.

The new attribute, LOCLADDR, can be used to configure a channel to use a particular IP address, port or port range for outbound communications. This attribute is applicable to all channel types except receiver and server-connection channels.

Another new channel attribute, BATCHHB, controls batch heartbeating. This attribute is applicable to sender, cluster-sender and cluster-receiver channels only.

The STOP CHANNEL, RESET CLUSTER, and REFRESH CLUSTER commands have extra options.

The value returned by CONNAME for the DISPLAY CHSTATUS CURRENT command is always the connection name, and not the remote queue manager or queue-sharing group name, which is now returned by the RQMNAME attribute.

Multiple channel exits are now supported. See the *WebSphere MQ Script (MQSC) Command Reference* for details of how the ALTER CHANNEL and DEFINE CHANNEL commands have changed to support this. You might need to change any cluster workload exits that you use (see the *WebSphere MQ for z/OS System Setup Guide* for more information).

Process definitions are no longer needed for channels that start automatically when messages arrive in the transmission queue.

The queue SYSTEM.CHANNEL.REPLY.INFO is no longer used.

WebSphere MQ for z/OS Version 5 Release 3.1 introduces two channel initiator parameters:

RCVTIME
Specifies approximately how long a TCP/IP channel waits to receive data from its partner.

RCVTMIN
Specifies the minimum time that a TCP/IP channel waits to receive data from its partner.

It also includes changes to the MQSC TRACE commands, START, DISPLAY, ALTER, and STOP, to improve the handling of channel initiator traces.

Context profiles

Context profiles can now be defined for each individual queue. If you use context security, you must either define a fully qualified context profile for each queue, using the form

```
hlq.CONTEXT.queuename
```

or alter your existing profile to be a generic context profile for all queues belonging to the specified queue manager or queue-sharing group, using the form

```
hlq.CONTEXT.**
```

If you leave your existing profiles unchanged, you will get security failure messages.

Utilities

A new utility, CSQJUFMT, has been implemented, which is an active log preformat utility. This utility is used to format active log data sets before they are used by a queue manager. If the active log data sets are preformatted by the utility, log write performance is improved on the queue manager's first pass through the active logs.

A new page set management function, PAGEINFO, has been added to the CSQUTIL utility to extract the page set recovery RBA from a page set.

There are new options for the FORMAT function of the CSQUTIL utility.

The CSQ1LOGP log print utility, CSQJU003 change log inventory utility, and CSQJU004 print log map utility, have been extended to support logical record sequence numbers (LRSNs). WebSphere MQ for z/OS Version 5 Release 3.1 introduces further enhancements.

See the *WebSphere MQ for z/OS System Administration Guide* for details of these utilities.

DB2 tables

Two new DB2 tables have been created: CSQ.ADMIN_B_STRBACKUP and CSQ.OBJ_B_AUTHINFO.

The channel table CSQ.OBJ_B_CHANNEL has increased in size, and now resides in a 32K tablespace.

Software prerequisites

Software prerequisites have changed. See“Software requirements” on page 157 for details of what you need.

Due to the added Secure Sockets Layer (SSL) support, two of the load libraries, thlqual.SCSQMVR1 and thlqual.SCSQMVR2, must be in PDS-E format, even if you do not use SSL.

In WebSphere MQ for z/OS Version 5 Release 3.1, the thlqual.SCSQSKL library is no longer used.

Samples

In WebSphere MQ for z/OS Version 5 Release 3.1, the sample programs have been updated to use more appropriate MQ API options and settings. A selection are supplied in executable, as well as source, form.

What to do when you migrate from a previous version

When you migrate from a previous version of MQSeries for OS/390, you can continue to use your existing subsystems with the new version, including their page sets, log data sets, object definitions, and initialization input data sets. You can continue to use your existing queues, including system queues such as the SYSTEM.CHANNEL.SYNCQ.

Do not cold start your queue managers when migrating from a previous version. If you do, you will lose all your messages and other information such as channel state.

However, there are some tasks that you need to perform when migrating from a previous version. Whether you need to perform each task depends on which of the new features you want to use, and which level of MQSeries you are migrating from. Generally, you need to perform more tasks if you are migrating from an earlier version of MQSeries than the previous version (that is, earlier than Version 5.2), you do not need to install the intervening versions.

A coexistence PTF is available for Version 5.2 to allow migration of a queue-sharing group of Version 5.2 queue managers to Version 5.3.1 in a phased manner, and interoperation of Version 5.2 and Version 5.3.1 queue managers within

a queue-sharing group. The *WebSphere MQ for z/OS System Setup Guide* describes a migration plan for an existing queue-sharing group that does not require an outage of the entire queue-sharing group.

The following list introduces the tasks that you might have to perform when migrating from Version 5.2 to Version 5.3.1. These tasks are described in the *WebSphere MQ for z/OS System Setup Guide*

- Apply the coexistence PTF to Version 5.2 queue managers in the queue-sharing group, if required.
- Update your DB2 tables.
- Check levels of prerequisite and corequisite software.
- Change the format of your SCSQMVR1 and SCSQMVR1 load libraries to PDS-E.
- Take your early code libraries out of your link list.
- Check and update your system parameter macros if you want to use some of the new functions.
- Alter your context profiles to conform to the new required definition.
- Update your channel object attributes.
- Change the active log and BSDS share options.
- Alter your security setup of your archive logs, active logs, and BSDS, if you want to be able to recover your CF structures.
- Set up your CF structure objects with the appropriate CFLEVEL.

See the *WebSphere MQ for z/OS System Setup Guide* for information about the additional tasks you might need to perform if you are migrating from Version 2.1 or Version 1.2.

Reverting to a previous version

After you have migrated to WebSphere MQ for z/OS Version 5 Release 3.1, you can revert to using Version 5.3, Version 5.2, Version 2.1, or Version 1.2, if exceptional circumstances so dictate. However, to do this you must apply a PTF to the previous version. This is described in the *WebSphere MQ for z/OS System Setup Guide*.

Part 6. Appendixes

Appendix A. Macros intended for customer use

The macros identified in this appendix are provided as programming interfaces for customers by WebSphere MQ.

Note: Do not use as programming interfaces any WebSphere MQ macros other than those identified in this appendix.

General-use programming interface macros

The following macros are provided to enable you to write programs that use the services of WebSphere MQ. The macros are supplied in library thlqual.SCSQMACS.

CMQA	CMQCXPA	CMQRMHA	CMQXCALA
CMQCDA	CMQDLHA	CMQTMA	CMQXCFBA
CMQCFA	CMQDXPA	CMQTMC2A	CMQXCFCA
CMQCFBSA	CMQGMOA	CMQWCRA	CMQXCDFA
CMQCFHA	CMQIIHA	CMQWDRA	CMQXCINA
CMQCFILA	CMQMDA	CMQWIHA	CMQXCVCA
CMQCFINA	CMQMDEA	CMQWPRA	CMQXPA
CMQCFSLA	CMQODA	CMQWXPA	CMQXQHA
CMQCFSTA	CMQPMOA	CMQXA	CMQXWDA
CMQCIHA			

Product-sensitive programming interface macros

The following macros are provided to enable you to write programs that use the services of WebSphere MQ. The macros are supplied in library thlqual.SCSQMACS.

CSQBDEF	CSQDQEST	CSQDQIST	CSQDQJST
CSQDQLST	CSQDQMAC	CSQDQMST	CSQDQPST
CSQDQSST	CSQDQWHC	CSQDQWHS	CSQDQ5ST
CSQDWQ	CSQDWTAS	CSQQDEFX	CSQQLITX

General-use programming interface copy files

The following COBOL copy files are provided to enable you to write programs that use the services of WebSphere MQ. The copy files are supplied in library thlqual.SCSQCOBC.

CMQCDL	CMQCFSTL	CMQIIHL	CMQRMHV
CMQCDV	CMQCFSTV	CMQIIHV	CMQTML
CMQCFBSL	CMQCFV	CMQMDEL	CMQTMV
CMQCFBSV	CMQCIHL	CMQMDEV	CMQTMC2L
CMQCFHL	CMQCIHV	CMQMDL	CMQTMC2V
CMQCFHV	CMQCXPL	CMQMDV	CMQWIHL
CMQCFILL	CMQCXPV	CMQODL	CMQWIHV
CMQCFILV	CMQDLHL	CMQODV	CMQV
CMQCFINL	CMQDLHV	CMQPMOL	CMQXV
CMQCFINV	CMQGMOL	CMQPMOV	CMQXQHL
CMQCFSLL	CMQGMOV	CMQRMHL	CMQXQHV
CMQCFSLV			

General-use programming interface include files

The following C include files are provided to enable you to write programs that use the services of WebSphere MQ. The files are supplied in library thlqual.SCSQC370.

CMQC CMQXC CMQCFC

The following PL/I include files are provided to enable you to write programs that use the services of WebSphere MQ. The files are supplied in library thlqual.SCSQPLIC.

CMQP CMQEPP CMQXP CMQCFP

Appendix B. Measured usage license charges with WebSphere MQ for z/OS

Note: This information does **not** apply when using the reduced function form of WebSphere MQ supplied with WebSphere Application Server.

Measured Usage License Charges (MULC) is a particular way of charging you for an IBM product that runs on a z/OS system, based on how much use you make of the product. To determine the product usage, the z/OS system records the amount of processor time that is used by the product when it executes.

z/OS and OS/390 can measure how much processing time is spent in doing work on behalf of the WebSphere MQ queue manager that is handling MQI calls, executing MQSC commands, or performing some other action to support the messaging and queuing functions used by your application programs. The amount of processing time is recorded in a file at hourly intervals, and the hourly records are totalled at the end of a month. In this way, the total amount of time that has been used by the WebSphere MQ for z/OS product on your behalf is computed, and used to determine how much you should pay for your use of the WebSphere MQ for z/OS product that month.

MULC is implemented as follows:

- When WebSphere MQ for z/OS is installed, it identifies itself to z/OS , and requests that the *System Management Facilities (SMF)* mechanism within z/OS is to automatically measure how much processor time is used by the WebSphere MQ for z/OS product.
- When enabled, the z/OS usage measurement facility collects usage figures for each hour of the day, and generates usage records that are added to a report file on disk.
- At the end of one full month, these usage records are collected by a program, which generates a report of product usage for the month. This report is used to determine the charge for the WebSphere MQ for z/OS product.

More details on MULC can be found in the *MVS Support for Measured License Charges* manual.

Appendix C. Notices

This information was developed for products and services offered in the United States. IBM may not offer the products, services, or features discussed in this information 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 information. The furnishing of this information does not give 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:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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 information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information 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 United Kingdom Laboratories,
Mail Point 151,
Hursley Park,
Winchester,
Hampshire,
England
SO21 2JN.

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 Programming License Agreement, or any equivalent agreement between us.

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.

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. 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 IBM's application programming interfaces.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AD/Cycle	AIX	CICS
COBOL/370	DB2	IBM
IBMLink	IMS	IMS/ESA
Language Environment	MQSeries	MVS
MVS/DFP	MVS/ESA	OpenEdition
OS/390	RACF	RAMAC
RMF	S/390	SAA
SupportPac	VTAM	WebSphere
z/OS		

Lotus and Lotus Notes are trademarks of International Business Machines Corporation and Lotus Development Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

Other company, product, or service names, may be the trademarks or service marks of others.

Index

D

E

F

G

H

Q

R

S

T

U

V

W

Sending your comments to IBM

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM.

Feel free to comment on what you regard as specific errors or omissions, and on the accuracy, organization, subject matter, or completeness of this book.

Please limit your comments to the information in this book and the way in which the information is presented.

To make comments about the functions of IBM products or systems, talk to your IBM representative or to your IBM authorized remarketer.

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

You can send your comments to IBM in any of the following ways:

- By mail, to this address:

  User Technologies Department (MP095)
  IBM United Kingdom Laboratories
  Hursley Park
  WINCHESTER,
  Hampshire
  SO21 2JN
  United Kingdom
- By fax:
  - From outside the U.K., after your international access code use 44–1962–816151
  - From within the U.K., use 01962–816151
- Electronically, use the appropriate network ID:
  - IBM Mail Exchange: GBIBM2Q9 at IBMMAIL
  - IBMLink™: HURSLEY(IDRCF)
  - Internet: idrcf@hursley.ibm.com

Whichever method you use, ensure that you include:

- The publication title and order number
- The topic to which your comment applies
- Your name and address/telephone number/fax number/network ID.

Printed in U.S.A.

GC34-6051-01